

Junos Space Network Director

Complete Software Guide for Network
Director Release 4.0

Published
2020-09-24

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Junos Space Network Director Complete Software Guide for Network Director Release 4.0
Copyright © 2020 Juniper Networks, Inc. All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement (“EULA”) posted at <https://support.juniper.net/support/eula/>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

About the Documentation | lii

Documentation and Release Notes | lii

Supported Platforms | lii

Documentation Conventions | lii

Documentation Feedback | iv

Requesting Technical Support | iv

Self-Help Online Tools and Resources | lvi

Creating a Service Request with JTAC | lvi

Guide 1 Network Director Quick Start Guide

1

Installing and Upgrading Junos Space Network Director

Network Director Installation Overview | 59

Setting up a Junos Space Appliance for Network Director | 59

Upgrading Junos Space Network Management Platform | 60

Installing Network Director | 61

Installing Network Director From Junos Space Store | 62

Installing Network Director by Manually Downloading the Network Director Application Image | 63

Upgrading Network Director | 64

Uploading DMI Schemas | 67

Preparing Devices for Management by Network Director | 68

Discovering Devices | 69

Next Steps | 71

2

Installing Data Learning Engine**Installing and Configuring Data Learning Engine for Network Director | 74**

Installing DLE | 74

What to Do Next | 78

Guide 2 Network Director User Guide

1

Working With Network Director**About Network Director | 81**

Understanding Network Director and the Management Life-Cycle Modes | 81

Benefits of Network Director | 82

Understanding Wireless Network Management in Network Director | 82

Understanding Cloud Analytics Engine and Network Director | 83

Understanding the Network Director User Interface | 85

Network Director Banner | 86

View Pane | 88

Displaying Devices in Various Network Views | 89

Filtering the Network Tree | 90

Expanding or Collapsing Nodes in the Network Tree | 90

Searching the Network Tree | 91

Tasks Pane | 91

Alarms | 93

Main Window or Workspace | 93

Tables in Network Director | 93

Moving and Resizing Columns | 93

Displaying the Column Drop-Down Menu | 94

Sorting on a Column | 94

Hiding and Exposing Columns | 95

Searching Table Contents | 95

Filtering Table Contents | 97

Accessing Network Director | 98

Logging In to Network Director | 98

Logging Out of Network Director | 99

Changing Your Password | 100

Understanding Network Director System Administration and Preferences | 101

Understanding Network Director User Administration | 101

Understanding the System Tasks Pane | 102

Audit Logs Overview | 103

Viewing Audit Logs From Network Director | 104

Managing Jobs | 105

Collecting Logs for Troubleshooting | 107

Setting Up User and System Preferences | 108

- Accessing the Preferences Page | 109

- Choosing Server Time or Local Time | 110

- Specifying Search Preferences | 110

- Enabling Import of Configuration Group Data from Ethernet Design | 110

- Specifying the Open Clos Server URL | 110

- Selecting the Approval Mode | 111

- Setting up Auto-resynchronization Preferences | 112

- Retaining Network Director Reports | 113

- Specifying Wireless Preferences | 113

- Changing Monitor Mode Settings | 113

 - Disabling Data Collection for Monitors | 114

 - Changing the Polling Interval | 116

 - Enabling and Disabling Collection for Managed Devices | 117

 - Specifying Database History Retention | 117

 - Specifying the Data Learning Engine (DLE) Settings | 118

- Changing Alarm Settings | 120

 - Configuring Global Alarm Notifications | 120

 - Retaining Alarm History | 120

 - Specifying Event History | 120

 - Enabling Alarms | 121

 - Changing the Severity of Individual Alarms | 134

 - Configuring Threshold Alarms | 134

 - Configuring Individual Alarm Notifications | 134

- Modifying Data Center Synchronization Interval Using the Virtualization Tab | 135

Getting Started with Network Director | 137

Getting Started with Junos Space Network Director | 137

Building Your Network | 137

Creating Profiles in Network Director | 138

Managing Software Images using Network Director | 139

Configuring Approval Modes for Device Configurations | 139

Resynchronizing Device Configuration | 140

Creating the Baseline Configuration | 140

Monitoring Your Network | 140

Setting up Network Traffic Analysis and Analyzing the Traffic | 141

Managing Network Faults and Notifications | 141

Generating Network Reports | 141

2

Working with the Dashboard

About the Dashboard | 144

Understanding the Dashboard | 144

Using the Dashboard | 145

Using Dashboard Widgets | 145

Dashboard Widget Reference | 147

Alarms Widget | 147

Alarms Widget Summary | 148

Alarms Widget Details | 148

Config Deployment Jobs Status Widget | 149

Config Deployment Jobs Status Widget Summary | 149

Config Deployment Jobs Status Widget Details | 150

Device & Port Latency Widget | 150

Device & Port Utilization Widget | 151

Using the Global Controls | 152

Interacting with the Heat Maps | 152

Viewing Active Flows on a Port | 153

Flow Analysis Details Window | 154

Viewing Traffic on a Device | 155

Equipment By Type Widget	158
Equipment By Type Widget Summary	158
Equipment By Type Widget Details	158
Port Status - Physical Widget	159
Port Status - Physical Widget Summary	159
Port Status - Physical Widget Details	160
Recent Flow Analysis Widget	160
Requirements for Flow Analysis	161
Recent Flow Analysis Main View	161
Flow Analysis Details Window	163
Simulate Flow Analysis Window	165
Top Talker - Wired Devices Widget	167
Top Talker - Wired Devices Widget Summary	168
Top Talker - Wired Devices Widget Details	168
Top Virtual Machines by Bandwidth Widget	168
Top vNetwork Hosts by Bandwidth Widget	169
Virtual Machines & Bare Metal Servers Widget	169
Requirements for Flow Analysis	170
Virtual Machines & Bare Metal Servers Widget Main View	171
Initiating Analysis on Selected Flows from the All Tab	171
Current Flows Window	172
Initiating Flow Analysis on All Flows on a VM or BMS from the Watchlist Tab	173
Add to Watchlist Window	174
Viewing the Results of Flow Analysis	176
Flow Analysis Results Window	176
Flow Analysis Details Window	177
Top Overlay Networks Widget	179
Top Overlay Networks Widget Summary	180
Top Overlay Networks Widget Details	180

Working in Build Mode

About Build Mode | 184

Understanding Build Mode in Network Director | 184

Discovering Devices | 184

Building the Logical, Location, and Custom Views | 185

Configuring Devices | 186

Deploying Device Configurations | 187

Importing Device Configurations | 187

Out-of-Band Configuration Changes | 188

Managing Devices | 188

Understanding the Build Mode Tasks Pane | 189

Understanding Network Configuration Profiles | 197

Assigning Profiles to an Interface, Device, or a Group of Devices | 201

Discovering Devices | 204

Discovering Devices in a Physical Network | 204

Specifying Target Devices | 205

Specifying Discovery Options | 208

Specifying Schedule Options | 209

Reviewing Device Discovery Options | 210

Viewing the Discovery Status | 210

Discovering Devices in a Datacenter Network | 212

Specifying the vCenter Targets | 213

Specifying the Virtual Network Credentials | 213

Reviewing the Virtual Network Discovery Options | 214

Viewing the Discovery Status | 214

Understanding the Device Discovery Process | 216

Benefits of the Device Discovery Process | 217

Troubleshooting Device Discovery Error Messages | 217

Setting Up Sites and Locations Using the Location View | 222

Understanding the Location View | 222

Setting Up the Location View | 224

Creating a Site | 228

How to Add or Edit a Location Site | 228

Creating or Editing a Site | 228

Configuring Buildings | 229

How to Add or Edit a Building | 229

Adding or Editing a Building for a Location | 230

Configuring Floors | 231

How to Add or Edit a Floor | 231

Adding or Editing a Building Floor for a Location | 232

Setting Up Closets | 233

How to Add or Edit a Closet | 233

Adding or Editing a Wiring Closet | 234

Assigning and Unassigning Devices to a Location | 234

How to Assign or Unassign Devices | 235

Assigning Devices | 235

Changing the Location of a Device | 236

How to Move a Device to a New Location | 237

Changing the Location of a Device | 237

Deleting Sites, Buildings, Floors, Wiring Closets, and Devices | 238

How to Delete a Location Object | 238

Deleting Sites | 239

Deleting Buildings | 239

Deleting Floors | 239

Deleting Closets | 239

Deleting Devices | 239

Configuring Outdoor Areas | 240

How to Configure an Outdoor Area | 240

Configuring an Outdoor Area | 241

Building a Topology View of the Network | 242

Understanding the Network Topology in Network Director | 243

Understanding the Topology View Tasks pane | 247

Setting Up the Topology View | 250

Managing the Topology View | 251

Viewing the Network Topology | 252

Refreshing the Topology | 254

Viewing Topology | 255

Viewing Topology Discovery Job | 256

Setting Up Locations | 257

Viewing the Alarm Details | 257

Discovering the Linux Hosts | 257

Displaying Device Connectivity | 257

Displaying QFabric Connectivity | 262

Displaying Virtual Chassis and Virtual Chassis Fabric Connectivity | 265

Displaying Virtual Network Connectivity | 268

Displaying Third-Party Device Details | 268

Uploading Floor Plans | 269

Uploading Topology Map | 270

Adding and Managing OUI Data in Network Director | 271

Creating Custom Device Groups | 273

Understanding Custom Device Groups | 273

Where Is the Custom Group Function Located in Network Director? | 274

How Do Custom Group Rules Work? | 274

What Happens When I Edit a Custom Group Rule? | 276

When Are Rules Executed? | 276

Creating Custom Device Groups | 276

Creating Custom Groups | 277

Creating a Custom Group | 277

Configuring Quick Templates | 282

Understanding Quick Templates | 282

Benefits of Quick Templates | 283

Configuring and Managing Quick Templates | 284

Creating a Quick Template | 285

Applying Templates to Devices | 286

Editing a Quick Template | 287

Deleting a Quick Template | 287

Cloning a Quick Template | 287

Using the Quick Template Details Window | 288

Viewing Deployed Quick Templates | 288

Configuring Device Settings | 290

Understanding Device Common Settings Profiles | 290

Creating and Managing Device Common Settings | 291

Managing Device Common Settings | 291

Creating a Device Common Settings Profile | 293

Specifying Basic Settings for Device Common Settings | 295

Specifying Management Settings for EX Switching Device Common Settings | 298

Specifying Management Settings for Wireless Device Common Settings | 302

Specifying Management Settings for Campus Switching ELS Device Common Settings | 306

Specifying Management Settings for Data Center Non-ELS Device Common Settings | 310

Specifying Management Settings for Data Center ELS Device Common Settings | 314

Specifying Protocol Settings for EX Switching Device Common Settings | 317

Specifying DNS Settings for Wireless Device Common Settings | 321

Configuring Wireless Dynamic Authorization Client (DAC) Settings | 321

Specifying DHCP Relay/DNS Settings for Campus Switching ELS Device Common Settings | 322

Specifying Protocol Settings for Data Center Switching Non-ELS Device Common Settings | 325

Specifying Protocol Settings for Data Center Switching ELS Device Common Settings | 327

Reviewing and Saving a Device Common Settings Configuration | 330

What to Do Next | 330

Assigning Device Common Settings to Devices | 331

Assigning Device Common Settings | 332

Editing the Assignments of the Device Common Setting | 334

Configuring Authentication, Authorization, and Access for Your Network | 335

Understanding Central Network Access Using RADIUS and TACACS+ | 335

Why Do I Want Remote Authentication ? | 336

Why Not Just Rely on Firewalls and Filters for Access Control? | 336

What About Using LDAP For Authentication? | 336

Where Is RADIUS Installed on the Network? | 337

How Is TACACS+ Installed on the Network? | 337

A Comparison of RADIUS and TACACS+ | 338

Creating and Managing RADIUS Profiles | 339

Managing RADIUS Profiles | 339

Creating RADIUS Profiles | 340

Specifying Settings for a RADIUS Profile | 341

What to Do Next | 344

Creating and Managing LDAP Profiles | 345

Managing LDAP Profiles | 346

Creating LDAP Profiles | 347

Specifying Settings for an LDAP Profile | 348

What to Do Next | 350

Understanding Access Profiles | 351

Creating and Managing Access Profiles | 352

Managing Access Profiles | 352

Creating an Access Profile | 354

Specifying Basic Settings for an EX Series Switching or Data Center Switching Access Profile | 356

Specifying RADIUS Accounting Settings for an EX Switching or Data Center Switching Access Profile | 358

Specifying Basic Settings for a Wireless Access Profile | 361

Specifying Server Group Settings for a Wireless Access Profile | 361

Specifying Basic Settings for a Campus Switching ELS Access Profile | 372

Specifying RADIUS and LDAP Settings for Campus Switching ELS | 372

Reviewing and Modifying the Access Profile Settings | 380

What To Do Next | 380

Understanding Authentication Profiles | 381

802.1X Authentication | 381

MAC RADIUS Authentication | 382

Captive Portal Authentication | 382

Last Resort Authentication | 382

Creating and Managing Authentication Profiles | 383

Managing Authentication Profiles | 383

Creating an Authentication Profile | 384

Specifying Authentication Settings for Switches | 386

Specifying Authentication Settings for Wireless | 390

What To Do Next | 394

Understanding Wireless Authorization Profiles | 395**Creating and Managing Wireless Authorization Profiles | 395**

Managing Authorization Profiles | 396

Creating a Wireless Authorization Profile | 397

Specifying Settings for a Wireless Authorization Profile | 399

What To Do Next | 403

Assigning Wireless Authorization Profiles to Controllers | 404

Assigning Authorization Profiles | 404

Editing Authorization Profile Assignments | 406

Configuring Interfaces and VLANs | 408**Understanding Port Profiles | 408**

Interface Settings Configured in the Port Profile | 409

Interface Settings Configured by Referencing Other Profiles | 410

Data Center Device Port Profile Settings | 410

Default Port Profiles | 410

Creating and Managing Port Profiles | 414

Managing Port Profiles | 415

Creating Port Profiles | 418

Specifying Settings for an EX Switching Port Profile | 418

Specifying Settings for a Campus Switching ELS Port Profile | 432

Specifying Settings for the Data Center Switching Non-ELS Port Profile | 446

Specifying Settings for a Data Center Switching ELS Port Profile | 460

	What to Do Next	475
Assigning and Unassigning Port Profiles from Interfaces		476
	Selecting Devices for Assignment	477
	Selecting Interfaces for Assignment	478
	Reviewing and Accepting the Assignments	481
	Editing Profile Assignments	482
	Unassigning a Port Profile from an Interface	483
Managing Auto Assignment Policies		484
Creating Auto Assignments		486
	Adding Port Profiles using the Select Port Profiles Page	487
	Adding Devices and Ports for Auto Assignment	487
	Viewing the Auto Assignment Policy Summary	488
Configuring Easy Config Setup		489
	Configuring Interface Settings	489
Understanding Port Groups		495
Creating and Managing Port Groups		495
	Managing Port Groups	496
	Creating Port Groups	497
	Specifying Settings for a Port Group	498
	What to Do Next	499
Understanding VLAN Profiles		499
Creating and Managing VLAN Profiles		502
	Managing VLAN Profiles	503
	Creating a VLAN Profile	504
	Specifying Basic EX Switching VLAN Settings	506
	Specifying Basic Wireless VLAN Settings	507
	Specifying Basic Campus Switching ELS VLAN Settings	508
	Specifying Basic VLAN Settings for Data Center Switching Non-ELS	510
	Specifying Basic VLAN Settings for Data Center Switching ELS	511
	Specifying Advanced VLAN Profile Settings for EX Series Switches	513
	Specifying Advanced VLAN Profile Settings for Wireless VLANs	515
	Specifying Advanced VLAN Settings for Campus Switching ELS	521
	Specifying Advanced VLAN Profile Settings for Data Center Switching Non-ELS	523
	Specifying Advanced VLAN Settings for Data Center Switching ELS	528

Reviewing and Saving the VLAN Profile Configuration | 530

What to Do Next | 531

Assigning a VLAN Profile to Devices or Ports | 531

Assigning a VLAN Profile | 532

Editing Profile Assignments | 535

Creating and Managing VLAN Pools | 535

Managing VLAN Pools | 536

Creating a VLAN Pool | 538

What To Do Next | 539

Configuring Firewall Filters (ACLs) | 540

Understanding Filter Profiles | 540

Creating and Managing Wired Filter Profiles | 542

Managing Wired Filter Profiles | 542

Creating a Wired Filter Profile | 543

Specifying Settings for an EX Series Switch Filter Profile | 544

Specifying Settings for a Campus Switching ELS Switch Filter Profile | 556

Specifying Settings for Creating a Data Center Switching Non-ELS Filter Profile | 570

Specifying Settings for a Data Center Switching ELS Filter Profile | 583

What to Do Next | 597

Creating and Managing Wireless Filter Profiles | 598

Managing Wireless Filter Profiles | 598

Creating a Wireless Filter Profile | 600

Specifying Settings for a Wireless Filter Profile (WLC) | 600

What To Do Next | 606

Assigning a Wireless Filter Profile to Controllers | 607

Configuring Class of Service (CoS) | 609

Understanding Class of Service (CoS) Profiles | 609

How Would I Use CoS (also known as QoS)? | 610

How Do I Create CoS Groups? | 610

How Is CoS Different From QoS? | 610

What Wireless Network Traffic Aspects Can I Control Using CoS? | 611

What CoS Parameters Can I Control? | 611

What Are the Default CoS Traffic Types? | 612

Data Center Switching CoS Configuration | 612

How Do I Implement Class of Service? | 613

Editing Discovered CoS Profiles | 613

Creating and Managing Wired CoS Profiles | 613

Managing Wired CoS Profiles | 614

Using the Default CoS Profiles for Switches | 615

Using the Default CoS Profiles for Data Center Switching | 615

Creating a Wired CoS Profile | 616

Specifying Settings for a Switching and Campus Switching ELS CoS Profile | 617

Specifying Settings for a Data Center Switching CoS Profile | 621

What to Do Next | 630

Creating and Managing Wireless CoS Profiles | 630

Managing Wireless CoS Profiles | 631

Creating a Wireless CoS Profile | 632

Specifying Settings for a Wireless CoS Profile | 633

What To Do Next | 635

Assigning a Wireless CoS Profile to Controllers | 635

Configuring Media Access Control Security (MACsec) | 637

Media Access Control Security Overview | 637

Configuring and Managing MACsec Profiles | 638

Creating a MACsec Profile | 639

Specifying Settings for a MACSsec Profile | 640

What to Do Next | 643

Assigning the MACsec Profiles | 644

Assigning a MACsec Profile to a Device | 644

Editing the MACsec Profile Assignments | 645

Configuring Link Aggregation Groups (LAGs) | 646

Understanding Link Aggregation | 646

Managing and Creating a Link Aggregation Group | 647

Link Aggregation Group Options | 649

Creating a Link Aggregation Group | 650

Managing ICCP Settings | 651

What To Do Next | 652

Understanding Multichassis Link Aggregation | 653

Creating and Managing Multichassis Link Aggregation Groups (MC-LAGs) | 654

Accessing the MC-LAG Page | 655

| ?

Creating an MC-LAG | 655

Selecting Peer Devices and Configuring Peer-to-Peer Link Settings | 656

Selecting Client Devices and Configuring Client-to-Peer Link Settings | 658

Saving MC-LAG Settings | 660

Deploying MC-LAG Configuration | 661

MC-LAG Automation Parameters | 661

Editing an MC-LAG | 663

Managing Peer Devices and Peer-to-Peer Link Settings | 664

Managing Client Devices and Client-to-Peer Link Settings | 665

Deleting an MC-LAG | 669

Managing an MC-LAG Created Through CLI Mode | 669

MC-LAG Peer Pairing | 669

Mapping Client Devices to Peer Devices | 670

Ports Mapping Between Peer-to-Peer and Client-to-Peer Devices | 670

Creating and Managing ESI Link Aggregation Groups (ESI-LAGs) | 670

Accessing the ESI-LAG Page | 671

Creating an ESI-LAG | 671

Selecting Peer Devices and Configuring Peer-to-Peer Link Settings | 672

Selecting Client Devices and Configuring Client-to-Peer Link Settings | 674

Saving ESI-LAG Settings | 675

Deploying ESI-LAG Configuration | 676

Editing an ESI-LAG | 676

Managing Peer Devices and Peer-to-Peer Link Settings | 677

Managing Client Devices and Client-to-Peer Link Settings | 678

Deleting an ESI-LAG | 679

ESI-LAG Automation Parameters | 680

Configuring Fibre Channel Gateways | 683

Configuring Fibre Channel Gateways | 683

Using the FC Gateway Service Profile to Configure FC Gateways | 683

Using a Combination of Profiles to Configure FC Gateways | 684

Creating and Managing FC Gateway Service Profiles | 684

Managing FC Gateway Service Profiles | 685

Creating FC Gateway Service Profiles | 686

Specifying Settings for an FC Gateway Service Profile | 687

What to Do Next | 688

Assigning an FC Gateway Service Profile to Ports | 688

Assigning an FC Gateway Service Profile | 689

Editing the Assignments of an FC Gateway Service Profile | 692

Creating Configurations for Fabrics | 693

Understanding Fabric Profiles | 693

Creating and Managing Fabric Profiles | 695

Managing Fabric Profiles | 695

Creating Fabric Profiles | 697

Specifying Settings for a Fabric profile | 697

What to Do Next | 701

Assigning a Fabric Profile to Devices and Ports | 701

Assigning a Fabric Profile | 702

Editing the Assignments of a Fabric Profile | 705

Creating and Managing Datacenter Fabrics | 707

Understanding Junos Fusion | 707

Understanding Junos Fusion Enterprise | 710

Understanding Junos Fusion Data Center | 713

Software Requirements for Junos Fusion | 715

Creating and Managing Fusion Configuration Templates | 716

 Create a Configuration Template for Junos Fusion Enterprise | 717

 Create a Configuration Template for Junos Fusion Data Center | 722

 Clone a Configuration Template | 726

 Apply Configuration Template to Devices | 726

 View Details about a Configuration Template | 731

 Delete a Configuration Template | 732

Managing Fusion Fabrics | 732

 Modify the Fusion Fabric | 734

 Edit Aggregation Device Details | 734

 Edit Satellite Device Details | 735

 Enable Uplink Failure Detection | 736

 View the Cabling Plan | 736

 View Fabric Connectivity | 736

 Replace Aggregation Device or Satellite Device in Junos Fusion | 737

Creating and Managing Satellite Software Upgrade Groups | 739

 Create a Software Upgrade Group | 740

 Edit a Software Upgrade Group | 740

 View Details of a Software Upgrade Group | 741

 Delete a Software Upgrade Group | 741

Understanding Layer 3 Fabrics | 742

User Privileges Required for the DHCP and File Server While Using Zero Touch Provisioning | 743

Managing Layer 3 Fabrics | 744

Creating Layer 3 Fabrics | 746

 Specifying the Fabric Requirements | 747

 Specifying the Device Details | 752

 Specifying Configuration Details | 753

 Viewing the Cabling Plan | 754

 Specifying Zero Touch Provisioning Details | 757

 Reviewing the Layer 3 Fabric Settings | 760

Editing Layer 3 Fabrics | 761

Viewing Layer 3 Fabric Connectivity | 764

Performing Layer 3 Fabric Connectivity Checks | 765

Setting Up Virtual Chassis Fabrics | 766

 Selecting a Provisioning Method | 767

 Adding Devices | 768

 Assigning Devices to a Virtual Chassis Fabric by Using a Graph | 768

 Assigning Devices to a Virtual Chassis Fabric by Using a Grid | 771

 Specifying Software Packages | 772

 Reviewing the Virtual Chassis Fabric Summary and Deploying Changes | 773

Managing Virtual Chassis Fabrics | 774

 Adding Spine Device(s) to the Virtual Chassis Fabric | 774

 Adding Leaf Node(s) to the Virtual Chassis Fabric | 775

 Replacing Spine or Leaf Devices | 776

 Removing Spine or Leaf Devices | 777

 Configuring Software Package for Upgrade | 778

 Deploying VCF Configuration Changes | 778

Understanding QFabric System Setup in Network Director | 779

Setting Up QFabric Systems | 780

 Managing Node Aliases | 780

 Managing Node Groups | 781

 Identifying CPE Switches | 782

 Reviewing the QFabric Summary and Deploying Changes | 782

Configuring Cloud-Based Datacenter Networks | 783

Understanding Cloud Networking | 784

Understanding Virtual Network Management | 785

Using OpenStack with VMware NSX | 787

Understanding the Build Mode Tasks Pane for Datacenter View	788
Cloud Infrastructure Requirements	791
Creating Data Centers Using Network Director	793
Create a Data Center	793
View the Data Center Creation Status	796
Assigning Network Devices to a Data Center	797
Managing Cloud Infrastructure	798
View the Cloud Infrastructure for a Data Center	799
Open VMware vSphere Web Client	800
Modify Cloud Infrastructure Details	800
Configure Orchestration Mode	800
Delete Cloud Infrastructure	802
Resynchronize Cloud Infrastructure	802
Viewing the Virtual Machine Inventory in a Cloud Infrastructure	803
View the Virtual Machine Inventory	803
View Connections Between VMs and the Physical Network	804
Viewing Overlay Networks	807
Viewing Virtual Tunnel End Point (VTEP) Details	808
Managing IP Connectivity	809
Adding an Autonomous System	810
Adding Devices	810
Creating Links	811
Viewing Data Center Connectivity	812
Viewing Bare Metal Server Details	818
Managing the Virtual Switch Inventory	819
View the Virtual Switch Inventory	819
Enable LLDP on Virtual Switches	820
Viewing the Hypervisor Servers in a Data Center	821
Managing Network Adapter Associations	822
Manage Network Adapter Associations	824
Configure Network Adapter Associations	824
Set Up the Devices for LLDP-Based Automatic Link Discovery	826
Export Network Adapter Associations	827
Delete Network Adapter Associations	827

Configuring Overlay Networks and Tenants | 828

VXLAN—EVPN Overlay Overview | 828

Create a Layer 3 Fabric based Underlay Network | 829

Creating and Managing Overlay Fabrics | 832

Creating an Overlay Fabric | 834

Modifying an Overlay Fabric | 835

Viewing Overlay Fabric Configuration Details | 835

Setting Up a VXLAN—EVPN-Based Data Center | 836

Creating and Managing Tenants | 837

Creating a Tenant | 838

Modifying a Tenant | 842

Viewing Tenant Configuration Details | 843

Configuring VRRP Profiles | 845

Understanding VRRP Profiles | 845

Creating and Managing VRRP Profiles | 846

Managing VRRP Profiles | 847

Creating VRRP Profiles | 848

Specifying VRRP Settings for an EX Switching, Campus Switching ELS, or Data Center Switching ELS or non-ELS | 848

Configuring Wireless Access Points and Radios | 851

Understanding Access Point Bias for Controllers | 852

How Do I Determine the Bias Settings I should Use? | 853

Example of a Layer 3 Network With Multiple Controllers | 854

A Third Option for Access Point Bias: Sticky | 854

How Do I Set the Controller Bias for an Access Point? | 855

How Can I Determine the Bias of an Access Point by Looking at a Controller? | 855

What About Controller Clusters? | 855

Understanding Wireless Radio Channels | 856

What WLAN Channels Are Available? | 856

What Channels Are not Available? | 857

How Do I Know Which Channels I Should Use? | 857

How Do I Avoid Co-Channel Interference? | 858

DFS Channels | **858**

802.11n Channels can be Wider and Work on Both Bands | **858**

How Are Channel Numbers Assigned? | **859**

How Do I Know What Channel an Access Point Is Using? | **859**

Understanding WMM Power Save and WLAN Client Battery Life | **859**

How Does WMM Power Save Extend Battery Life? | **860**

Where is WMM Defined? | **860**

How is WMM Power Save Implemented on Juniper Networks WLANs? | **861**

WMM Power Save is Disabled by Default | **861**

Why Should I Enable WMM Power Save in a Radio Profile? | **861**

Understanding Adaptive Channel Planner | **861**

Why Use Adaptive Channel Planner? | **862**

When Should I Use Adaptive Channel Planner? | **862**

Adaptive Channel Planner Improves Performance | **863**

Adaptive Channel Planner Resolves Interference | **863**

Adaptive Channel Planner Is Used by Dynamic Frequency Selection to Comply with Country Regulations | **863**

How Does Adaptive Channel Planner Work? | **863**

How Do I Configure Adaptive Channel Planner? | **864**

What Are Adaptive Channel Planner Results? | **864**

When Is Adaptive Channel Planner Most Beneficial? | **865**

What Happens When Severe Interference Is Detected? | **866**

Understanding Auto Tune Power Policy for Wireless Radios | **866**

How Does Wireless Transmit Power Work? | **867**

How Does Auto Tune Power Policy Work? | **867**

When is Auto Tune Power Policy Most Helpful? | **868**

How Do I Turn Off an Auto Power Policy? | **868**

What Changes Can I Make to an Auto Tune Power Policy? | **868**

Understanding Wireless Scanning | **869**

What Is the Difference Between Passive and Active Scanning? | **869**

What Channels Are Scanned? | **870**

How Does Scanning Work? | **870**

How Does Active Scanning Work? | **871**

What Additional Information Is Learned by an Active Scan? | **871**

What Happens to Scanned Information? | 872

CTS-to-self During Scanning | 872

What Is Spectral RF Scanning? | 872

Understanding Distributed Access Point Behavior on a Layer 3 Network | 873

What Is a Layer 3 Network? | 874

How Does an Access Point Find a Controller on a Layer 3 Network? | 875

Understanding How To Add Access Points to a Wireless Network By Using Network Director | 878

Understanding Radio Profiles | 879

RF Scanning | 879

Spectral Scanning | 880

Dynamic Frequency Selection (DFS) Channels | 880

RFID Asset Tracking | 880

WMM Power Save | 881

Countermeasures | 881

Limiting Client Power | 882

802.11n Channel Width | 882

Automatic Channel Tuning | 882

Automatic Power Tuning | 882

IEEE 802.11 | 882

Long and Short Preamble Length | 882

Understanding Auto AP Profiles | 883

How Are Specifically Configured Access Points Different from Access Points Configured with Auto AP? | 884

How Should I Use Auto AP Profiles? | 884

How Does Auto AP Work? | 885

Understanding WLAN Service Profiles | 885

SSID | 886

Beaconing the SSID Name | 887

Mapping WLAN Service Profiles to Additional Profiles | 887

SSID Encryption | 887

Authentication Used for Encryption Methods | 888

Associated Authentication Profile | 888

Associated Authorization Profile | 888

VLAN Use | 888

Bandwidth Limit for Client Sessions | **888**

Load Balancing Between Access Points | **888**

Using Proxy ARP | **889**

Restricting DHCP | **889**

Client Types | **889**

Call Admission Control Settings for Voice | **890**

Retry Count | **890**

Client Timeouts | **890**

802.11n Settings | **890**

Guard Intervals | **890**

Frame Aggregation | **891**

MAC Service Data Unit (MSDU) Length | **891**

MPDU Length | **891**

Maximum Bandwidth Used by a WLAN Service Profile's SSID | **891**

Maximum Transmission Unit Parameter | **891**

Client Probing of Idle Clients | **892**

Enable Pre-Shared Key (PSK) for WPA or WPA2 | **892**

Create a Pre-Shared Key (PSK) Phrase for WPA or WPA2 | **892**

Create a Pre-Shared Key (PSK) Raw Phrase for WPA or WPA2 | **892**

Enforce Data Rates | **892**

Retry Count for Sending Frames | **892**

WPA Encryption Type Used | **893**

Shared Key Authentication Values | **893**

Radio Transmit Rates Used | **893**

WMM Power Save | **893**

Understanding Wireless Mesh | **894**

Example Mesh Topology With One Access Point Wired | **895**

Why Use Mesh? | **896**

How Does Mesh Work? | **896**

Planning a Mesh Portal | **897**

How Do I Set Up and Configure Mesh? | **898**

How Do I Configure a Mesh Access Point? | **898**

Security Between a Mesh AP and the Mesh Portal AP | **898**

Understanding Wireless Encryption and Ciphers | 899

Wired Equivalent Privacy (WEP) was the Original Wireless Encryption | **900**

WPA Encryption Replaced WEP | **900**

WPA2 Is the Strongest Encryption Available | **901**

Security Ciphers for WPA and WPA2 | **901**

Which Encryption Method Should I Use? | **902**

Understanding PSK Authentication | 903

What Is PSK? | **903**

How Does PSK Work? | **903**

When Would I Use PSK Authentication? | **904**

Why Would I not Use PSK Authentication? | **904**

How Is WPA Encryption Different from WPA-PSK Encryption? | **905**

Understanding Web Portals | 905

Why Use a Web Portal on Your Wireless Network? | **906**

How Does MSS Support Web Portals? | **906**

How Does Web Portal WebAAA Work? | **906**

How Are Web Portals Created in Network Manager? | **907**

Understanding Local Switching on Access Points | 907

Why Use Local Switching? | **908**

How Does Local Switching Work? | **908**

When to Use Local Switching? | **909**

When not to Use Local Switching? | **910**

VLAN Profiles and Local Switching | **911**

How Do I Configure Local Switching? | **911**

Does a Web Portal Work with Local Switching Configured? | **911**

Does QoS Policy Enforcement Work with Local Switching Configured? | **911**

What Happens if the Controller or WAN Link Goes Down? | **911**

Is Local Switching Included in any Other Wireless Features? | **912**

Understanding Wireless Bridging | 912

Why Use Wireless Bridging? | **913**

How Does Wireless Bridging Work? | **914**

How is Wireless Bridging Configured? | **914**

Understanding Wireless Interference | 914**What Causes Wireless Radio Frequency Interference? | 915****Effects of Interference Seen by Clients | 915****You Can Monitor RF Interference with Network Director | 915****What Is RF Jamming? | 916****Understanding Rogue Access Points | 917****What is a Rogue Access Point? | 917****How Are Rogue Access Points and Rogue Clients Identified By Controllers? | 918****How are Rogue access points and Rogue Clients Classified as Rogue? | 918****You Can Change Some Rogue Classification Rules | 920****What Harm Can a Rogue Access Point Do? | 920****Section | ?****What Can I do To Prevent Rogue Access Points? | 921****How Do I Prevent a Benign Access Point From Being Classified as a Rogue? | 923****Understanding Rogue Clients | 923****What Defines a Rogue Client? | 924****How Are Rogue Clients Detected? | 925****What Can I do To Prevent Rogue Client Damage? | 925****How Do I Prevent a Benign Client From Being Classified as a Rogue? | 925****How Do I Make Sure An SSID Won't Be Classified as Rogue? | 925****How Can I Make Sure a Device is Classified as Rogue? | 925****Understanding an SSID Masquerade | 926****Fake SSID Attacks | 926****Fake BSSID Attacks | 926****Detecting Fake SSID Attacks and Fake BSSID Attacks | 927****Understanding Ad-Hoc Networks | 927****Why Are Ad-Hoc Networks a Security Risk? | 928****How Do I Detect an Ad-Hoc Network? | 928****Are All Ad-Hoc Networks Malicious? | 928****How Do I Know Whether an Ad-Hoc Network Is Malicious? | 929**

Understanding LLDP and LLDP-MED	930
Importing RingMaster Data to Network Director	931
Creating and Managing a Radio Profile	932
Managing Radio Profiles	933
Radio Profiles Need an Associated WLAN Profile	934
Creating a Radio Profile	935
Specifying Quick Setup Radio Profile Settings	935
Specifying Custom Radio Profile Setup Settings	939
Basic Settings for Custom Radio Profiles	940
RF Scanning Settings for Custom Radio Profiles	944
Auto Tune Settings for Custom Radio Profiles	945
Voice Configuration Settings for Custom Radio Profiles	948
802.11 Attributes Settings for Custom Radio Profiles	949
Adaptive Channel Planning Settings for Custom Radio Profiles	950
Snooping Mapping Settings for Custom Radio Profiles	951
What To Do Next	951
Assigning a Radio Profile to Radios	952
Selecting the Controllers for Radio Profile Assignment	954
Assigning the Radio Profile to Radios	954
Editing Radio Profile Assignments	956
What To Do Next	956
Specifying Custom Radio Profile Setup Settings	957
Basic Settings for Custom Radio Profiles	957
RF Scanning Settings for Custom Radio Profiles	961
Auto Tune Settings for Custom Radio Profiles	962
Voice Configuration Settings for Custom Radio Profiles	965
802.11 Attributes Settings for Custom Radio Profiles	967
Adaptive Channel Planning Settings for Custom Radio Profiles	968
Snooping Mapping Settings for Custom Radio Profiles	969
What To Do Next	969
WLAN Setup	970
Setting Up Wireless Radios	971
Creating a Radio Profile Using WLAN Setup	971
Assigning a Radio Profile to Radios By Using WLAN Setup	975

What To Do Next | 976

Configuring Wireless Mesh and Bridging | 976

Create a Mesh SSID and Radio Profile for Access Point Portal Radios | 977

Create an SSID and Radio Profile for Access Point Mesh Radios | 977

Configure the Mesh Access Points | 978

Physically Set Up the Mesh Access Points | 979

After the Mesh is Set Up | 979

Make Any Further Changes to Mesh Access Points From the Switch | 980

Creating and Managing Wireless Auto AP Profiles | 980

Managing Wireless Auto AP Profiles | 981

Creating a Wireless Auto AP Profile | 983

Specifying Basic Settings for a Wireless Auto AP Profile | 984

Specifying LLDP & Remote AP Settings for a Wireless Auto AP Profile | 985

Specifying Advanced Settings for a Wireless Auto AP Profile | 987

Reviewing a Wireless Auto AP Profile | 990

What To Do Next | 991

Assigning an Auto AP Profile to Controllers | 991

Controller Selection for Auto AP Assignment | 992

Assigning the Auto AP Profile to Controllers | 993

Reviewing and Assigning the Auto AP Profile Configuration | 994

Editing Auto AP Profile Assignments | 994

What To Do Next | 995

Understanding Bonjour | 995

Why Would I Use Bonjour? | 996

What Is Zero Configuration? | 996

How Do I Configure Bonjour on a Juniper Networks Network Using Network Director? | 996

Creating and Managing mDNS Profiles | 997

Managing mDNS Profiles | 997

Creating an mDNS Profile | 999

Specifying Settings for an mDNS Profile | 999

What to Do Next | 1001

Assigning an mDNS Profile to Devices | 1001

Creating and Managing an mDNS VLAN List | 1002**Managing an mDNS VLAN List | 1003****Creating an mDNS VLAN List | 1003****Specifying mDNS VLAN Members | 1004****Creating and Managing Local Switching Profiles | 1005****Managing Local Switching Profiles | 1006****Creating a Local Switching Profile | 1008****Specifying Local Switching Profile Settings | 1008****What To Do Next | 1010****Assigning a Local Switching VLAN Profile to Existing Access Points | 1011****Assigning a Local Switching Profile During Access Point Configuration | 1013****Creating and Managing Remote Site Profiles | 1014****Managing Remote Site Profiles | 1014****Creating a Remote Site Profile | 1016****Specifying Remote Site and Intrusion Detection Logging Settings | 1016****What To Do Next | 1024****Assigning Remote Site Profiles to Access Points | 1024****Creating and Managing RF Detection Profiles | 1026****Managing RF Detection Profiles | 1027****Creating an RF Detection Profile | 1028****Specifying RF Detection Profile Classification Settings | 1029****What To Do Next | 1032****Assigning RF Detection Profiles to Controllers | 1033****Configuring Link Layer Discovery Protocol (LLDP) on an Access Point | 1035****Configuring Wireless Controllers | 1037****Configuring a Controller | 1037****Configuring System Information for a Controller | 1038****Configuring Link Layer Discovery Protocol (LLDP) on a Controller | 1039****Configuring IP Services on a Controller | 1040****Configuring DSCP CoS Mapping on a Controller | 1042****Configuring RF Auto-tuning on a Controller | 1043****Configuring Load Balancing on a Controller | 1044****Configuring AAA 802.1X on a Controller | 1045**

Configuring AAA RADIUS on a Controller | 1047

Configuring AAA LDAP on a Controller | 1049

What To Do Next | 1050

Configuring Wireless Mobility and Network Domains | 1051

Understanding Mobility Domains | 1051

Creating a Mobility Domain for Wireless LAN Controllers | 1053

Configuring Security Settings for a Mobility Domain | 1056

Creating Network Domains for Wireless LAN Controllers | 1058

Configuring WLAN Service (SSIDs) | 1061

Understanding the Network Terms SSID, BSSID, and ESSID | 1061

An SSID is the Name of a Network | 1062

BSSIDs Identify Access Points and Their Clients | 1063

Ad-Hoc Networks Do Not Have a MAC Address | 1064

An ESS Consists of BSSs | 1064

Understanding Network Director SSID Configuration Using Profiles | 1064

Configuring an SSID with Network Director | 1065

SSID Access Information | 1066

SSID Authentication Information | 1067

SSID WLAN Information | 1067

SSID Radio Information | 1067

What Do I Do When all of Profiles are Complete? | 1067

Understanding Call Admission Control | 1068

What Radios Does Call Admission Control Affect? | 1069

How Do I Configure Call Admission Control? | 1070

Understanding Load Balancing for Wireless Radios | 1070

Why Would I Need Load Balancing? | 1071

When Would I Avoid Load Balancing? | 1071

How Load Balancing Works | 1071

Can I Group Access Points for Load Balancing? | 1071

Where Do I Configure Load Balancing in Network Director? | 1072

Understanding Wireless Encryption and Ciphers	1072
Wired Equivalent Privacy (WEP) was the Original Wireless Encryption	1073
WPA Encryption Replaced WEP	1073
WPA2 Is the Strongest Encryption Available	1074
Security Ciphers for WPA and WPA2	1074
Which Encryption Method Should I Use?	1075
Understanding the IEEE 802.11 Standard for Wireless Networks	1076
Differences Between 802.11 Standards	1076
802.11 Divides Each Frequency Band into Channels	1077
What Is 802.11i Security?	1078
What Is 802.1X?	1078
Understanding PSK Authentication	1079
What Is PSK?	1079
How Does PSK Work?	1079
When Would I Use PSK Authentication?	1080
Why Would I not Use PSK Authentication?	1080
How Is WPA Encryption Different from WPA-PSK Encryption?	1081
Understanding WLAN Service Profiles	1081
SSID	1082
Beaconing the SSID Name	1083
Mapping WLAN Service Profiles to Additional Profiles	1083
SSID Encryption	1083
Authentication Used for Encryption Methods	1083
Associated Authentication Profile	1084
Associated Authorization Profile	1084
VLAN Use	1084
Bandwidth Limit for Client Sessions	1084
Load Balancing Between Access Points	1084
Using Proxy ARP	1084
Restricting DHCP	1085
Client Types	1085
Call Admission Control Settings for Voice	1085
Retry Count	1086
Client Timeouts	1086

802.11n Settings | 1086**Guard Intervals | 1086****Frame Aggregation | 1086****MAC Service Data Unit (MSDU) Length | 1087****MPDU Length | 1087****Maximum Bandwidth Used by a WLAN Service Profile's SSID | 1087****Maximum Transmission Unit Parameter | 1087****Client Probing of Idle Clients | 1087****Enable Pre-Shared Key (PSK) for WPA or WPA2 | 1087****Create a Pre-Shared Key (PSK) Phrase for WPA or WPA2 | 1088****Create a Pre-Shared Key (PSK) Raw Phrase for WPA or WPA2 | 1088****Enforce Data Rates | 1088****Retry Count for Sending Frames | 1088****WPA Encryption Type Used | 1088****Shared Key Authentication Values | 1089****Radio Transmit Rates Used | 1089****WMM Power Save | 1089****Creating and Managing a WLAN Service Profile | 1090****Managing WLAN Service Profiles | 1090****Before You Create a WLAN Service Profile | 1092****Creating a WLAN Service Profile | 1092****Specifying WLAN Service Profile Quick Setup | 1093****Specifying WLAN Service Profile Custom Setup Settings | 1102****Specifying Basic Settings for Custom WLAN Profile Setup | 1103****Specifying WLAN Settings for Custom WLAN Profile Setup | 1111****Specifying Web Portal Settings Under Advanced WLAN Profile Setup | 1112****Specifying 802.11n and Client Type Settings Under Advanced WLAN Profile Setup | 1113****Specifying Voice Configuration Settings Under Advanced WLAN Profile Setup | 1115****Specifying Broadcast Settings Under Advanced WLAN Profile Setup | 1116****Specifying Client Timeouts Under Advanced WLAN Profile Setup | 1116****Specifying Rate Configuration Under Advanced WLAN Profile Setup | 1117****Specifying Device Detection Settings Under Advanced WLAN Profile Setup | 1119****What To Do Next | 1119**

Understanding Voice Clients and Voice Traffic | 1120**What Is Voice Over IP? | 1121****Voice over Wireless Phones | 1121****Private Branch Exchange (PBX) | 1121****How Is Voice Traffic Different From Data Traffic? | 1121****Latency, Jitter, and Packet Loss Affect Voice Traffic | 1121****Voice Traffic Is Susceptible to Roaming Issues | 1122****What Protocols Are Used for Voice? | 1122****What Is Different About Configuring for Voice Traffic? | 1122****Consider Using Call Admission Control (CAC) to Limit Clients per Access Point | 1123****Ensure That Network Equipment Supports Seamless Roaming | 1123****Support Quality of Service on all Hardware Used for Voice | 1123****Use Automatic Power Save Delivery to Preserve the Battery Life of Phones | 1123****Create a Unique WLAN Service Profile for Voice | 1123****Create a Unique Radio Profile for Voice | 1123****Configuring a Voice SSID with Network Director | 1124****Creating a CoS Profile Dedicated to Voice | 1124****Creating and Managing RF Snooping Filter Profiles | 1125****Managing Snooping Filter Profiles | 1126****Creating an RF Snooping Filter Profile | 1127****Specifying RF Snooping Settings | 1128****What To Do Next | 1132****Assigning RF Snooping Filter Profiles to Access Points | 1132****Assign an RF Snooping Profile to Access Points | 1133****What To Do Next | 1134****Managing Network Devices | 1135****Viewing the Device Inventory Page | 1136****Viewing Device Connectivity | 1139****Viewing Profiles Assigned to a Device | 1144****Viewing the Physical Inventory of Devices | 1146****Viewing Licenses With Network Director | 1147****Viewing a Device's Current Configuration from Network Director | 1150****Assigning Devices to Logical Category | 1150**

- Accessing a Device's CLI from Network Director | **1151**
- Accessing a Device's Web-Based Interface from Network Director | **1153**
- Deleting Devices | **1154**
- Rebooting Devices | **1155**
- Viewing Virtual Machines | **1155**
- Adding and Managing an Individual Access Point | **1156**
 - Managing Access Points | **1157**
 - Adding an Access Point to a Wireless Network | **1159**
 - Specifying Basic Access Point Settings | **1161**
 - Enabling Local Switching on an Access Point | **1163**
 - Configuring an Access Point as Remote | **1164**
 - Configuring Link Layer Discovery Protocol (LLDP) on an Access Point | **1165**
 - Configuring Bonjour on an Access Point | **1167**
 - Specifying Access Point Radio Settings | **1168**

4

Working in Deploy Mode

About Deploy Mode | **1172**

Understanding Deploy Mode in Network Director | **1172**

- Deploying Configuration Changes | **1173**
- Managing Software Images | **1174**
- Zero Touch Provisioning | **1175**
- Managing Devices | **1175**
- Managing Device Configuration Files | **1175**
- Managing Baseline Configuration | **1176**

Understanding the Deploy Mode Tasks Pane | **1177**

Deploying and Managing Device Configurations | **1180**

Deploying Configuration to Devices | **1180**

- Selecting Configuration Deployment Options | **1181**
- Using the Change Request Details Page | **1185**
- Creating a Change Request | **1186**
- Validating Configuration | **1186**
- Discarding the Pending Configurations | **1187**
- Viewing Pending Configuration Changes | **1187**

- Using the Pending Changes Window | **1187**
- Using the Configuration or Pending Configuration Window | **1188**
- Using the Deploy Configuration Errors/Warnings Window | **1189**
- Using the Configuration Validation Window | **1189**
- Deploying Configuration Changes to Devices Immediately | **1189**
- Scheduling Configuration Deployment | **1190**
- Specifying Configuration Deployment Scheduling Options | **1190**
- Editing Change Requests | **1191**
- Deleting Change Request | **1192**
- Resubmitting a Change Request | **1192**
- Performing a Rollback | **1193**
- Managing Configuration Deployment Jobs | **1194**
 - Selecting Configuration Deployment Job Options | **1194**
 - Viewing Configuration Deployment Job Details | **1195**
 - Canceling Configuration Deployment Jobs | **1196**
- Deploy Configuration Window | **1196**
- Importing Configuration Data from Junos OS Configuration Groups | **1198**
 - Enabling Import of Configuration Group Data | **1198**
 - Viewing Configuration Group Data | **1199**
 - Using the Configuration or Pending Configuration Window | **1200**
 - Deploying Configuration Group Changes to Devices Immediately | **1201**
 - Scheduling Configuration Group Change Deployment | **1201**
 - Specifying Configuration Deployment Scheduling Options | **1201**
 - Using the Deploy Configuration Errors/Warnings Window | **1202**
- Enabling High-Frequency Traffic Statistics Monitoring on Devices | **1202**
- Configuring Network Traffic Analysis | **1204**
- Approving Change Requests | **1206**
- Enabling SNMP Categories and Setting Trap Destinations | **1208**
 - Viewing Eligible Devices for Trap Forwarding | **1208**
 - Enabling Trap Forwarding | **1209**
 - Deploying SNMP Trap Configurations | **1210**
- Understanding Resynchronization of Device Configuration | **1214**
 - The Resynchronize Device Configuration Task | **1215**
 - How Resynchronization Works in NSOR Mode | **1216**

How Resynchronization Works in SSOR Mode | 1218

How Network Director Resynchronizes the Build Mode Configuration | 1220

Resynchronizing Device Configuration | 1220

The Resynchronize Device Configuration List of Devices | 1221

Resynchronizing Devices When Junos Space Is in NSOR Mode | 1222

Resynchronizing Devices When Junos Space Is in SSOR Mode | 1223

Resynchronizing Devices in Manual Approval Mode | 1224

Viewing the Network Changes | 1224

Viewing Resynchronization Job Status | 1225

Managing Device Configuration Files | 1225

Selecting Device Configuration File Management Options | 1226

Backing Up Device Configuration Files | 1227

Restoring Device Configuration Files | 1227

Viewing Device Configuration Files | 1228

Comparing Device Configuration Files | 1228

Deleting Device Configuration Files | 1229

Managing Device Configuration File Management Jobs | 1229

Creating and Managing Baseline of Device Configuration Files | 1230

Selecting Baseline Management Options | 1231

Baselining Device Configuration Files | 1231

Restoring Baseline Device Configuration Files | 1232

Viewing Baseline Configuration Files | 1233

Comparing Baseline Configuration with Current Configuration | 1233

Deleting Baseline | 1233

Managing Baseline Management Jobs | 1234

Deploying and Managing Software Images | 1235

Managing Software Images | 1235

Selecting Software Image Management Options | 1236

Adding Software Images to the Repository | 1236

Using the Device Image Upload Window | 1237

Viewing Software Image Details | 1237

Using the Device Image Summary Window | 1237

Deleting Software Images | 1238

Deploying Software Images | 1238

Specifying Software Deployment Job Options | 1239

Selecting Software Images To Deploy | 1240

Selecting Options for Software Deployment | 1241

Summary of Software Deployment | 1243

Managing Software Image Deployment Jobs | 1243

Selecting Software Image Management Options | 1244

Viewing Software Image Job Details | 1245

Using the Device Image Staging Window | 1245

Canceling Software Image Jobs | 1246

Managing Devices | 1248

Converting Automatically Discovered Access Points to Manually Configured Access Points | 1248

Enabling or Disabling Network Ports on Switches | 1250

Understanding Node Groups for a Qfabric System | 1251

Network Node Groups | 1251

Server Node Groups | 1252

Creating and Managing Node Groups for a QFabric System | 1252

Managing Node Groups | 1253

Creating Node Groups | 1254

Specifying Settings for a Node Group | 1254

Converting the QSFP+ Ports on QFX Series and QFabric Devices | 1256

Selecting Devices | 1256

Converting Ports | 1258

Reviewing and Deploying Port Conversions | 1259

Setting Up Zero Touch Provisioning for Devices | 1260

Understanding Zero Touch Provisioning in Network Director | 1260

Configuring and Monitoring Zero Touch Provisioning | 1261

Configuring Zero Touch Provisioning | 1262

Specifying the Server Details | 1263

Specifying the Software Image and Configuration Details | 1265

Reviewing and Modifying Zero Touch Provisioning Settings | 1266

What To Do Next | **1266**

Configuration Statements for Custom Configuration of DHCP Server | **1266**

Monitoring Zero Touch Provisioning Profiles | **1267**

Monitoring Devices and Traffic

About Monitor Mode | 1269

Understanding Monitor Mode in Network Director | **1269**

Scope and Monitor Tab Availability | **1270**

Monitors and Tasks | **1271**

Scope and Data Aggregation | **1272**

How Network Director Collects and Displays Monitoring Data | **1273**

How Network Director Displays and Stores Trend Data | **1273**

More About the Monitor Tabs | **1274**

The Summary Tab | **1274**

The Traffic Tab | **1275**

The Client Tab | **1275**

The RF Tab | **1275**

The Equipment Tab | **1275**

The Fabric Analysis Tab | **1276**

VC Equipment Summary Tab | **1276**

Understanding the Monitor Mode Tasks Pane | **1277**

Monitoring Traffic | 1283

Monitoring Traffic on Devices | **1283**

Monitoring Port Traffic Statistics | **1284**

Procedure for Monitoring Port Traffic Statistics | **1284**

Port on Device Window | **1285**

Port Traffic Stats Window | **1286**

Monitoring Traffic on Layer 3 VLANs | **1287**

Procedure for Monitoring Layer 3 VLAN Traffic Statistics | **1287**

L3 VLAN Traffic Stats Window | **1288**

Monitoring Routing Instances | 1289

- Procedure for Monitoring Routing Instances | 1290

- Show Routing Instances Window | 1290

- Show Interfaces Window | 1291

- Show Bridge Domains Window | 1292

- Show Connections | 1293

- Show Routing Tables | 1296

- Show MAC Table | 1299

Monitoring Port Utilization | 1300

- How to Access the Port Utilization Task | 1300

- Port Utilization Details Window | 1301

- Utilization for Device Window | 1301

- Device View | 1302

- Port View | 1302

- Utilization for Fabric Devices Window | 1303

- Device View | 1303

- Port View | 1304

Monitoring Tenant Details | 1304

- Viewing the List of Tenants | 1306

- View Port Details of Tenants | 1307

- View Endpoints | 1307

- View the Port Utilization Trend for a VXLAN Port | 1309

Monitoring Virtual Chassis Protocol Statistics | 1310

- Procedure for Monitoring Virtual Chassis Protocol Statistics | 1310

- Virtual Chassis Protocol Statistics Window | 1310

Viewing Congestion Events | 1312**Monitoring Client Sessions | 1314****Finding User Sessions | 1314**

- Procedure for Finding User Sessions | 1314

- Search User Session Window | 1315

Finding End Points | 1319

- Procedure for Finding End Points | 1319

- Find End Point Window | 1319

Refreshing End Point Information | 1320

Monitoring Client Sessions | 1321

Monitoring Radio Frequency | 1322

Monitoring RF 802.11 Packet Errors | 1323

Monitoring RF Interference Sources on One Radio | 1325

Monitoring RF Radio Interference Sources | 1325

RF Interference Sources Pie Chart for a Radio | 1326

Monitoring RF Interference Sources For Radios on One Access Point | 1328

Monitoring RF Interference Sources on Wireless Devices | 1329

Troubleshooting Excessive Wireless Interference | 1332

Monitoring RF Signal-to-Noise Ratio | 1334

Monitoring RF Throughput | 1336

Monitoring the Percentage of RF Packet Retransmissions | 1338

Procedure for Viewing RF Packet Transmission | 1338

Monitoring the RF Neighborhood | 1340

Procedure for Viewing a Radio's Neighbors | 1340

RF Neighborhood List | 1341

Monitoring the RF Spectrum of a Radio | 1343

Procedure for Viewing the Radio Spectrogram | 1343

Spectrogram Charts | 1344

Channel Spectrogram Chart | 1345

Swept Spectrum and Duty Cycle Charts | 1346

Monitoring Devices | 1347

Comparing Device Statistics | 1347

Procedure for Comparing Device Statistics | 1347

Compare Interfaces Window | 1348

Showing ARP Table Information | 1349

Procedure for Showing ARP Table Information | 1349

Show ARP Table Information Window | 1349

Viewing PoE Information | 1350

Procedure for Viewing PoE Information | 1350

Show PoE Information Window | 1351

Monitoring Backed-Up Wireless Access Points on Wireless LAN Controllers | 1352

 Procedure for Monitoring Backed-Up Wireless Access Points on Wireless LAN Controllers | 1352

 Backed-Up APs Window | 1353

Monitoring the Status of Aggregated Access Points and Radios | 1354

Monitoring the Status of Logical Interfaces | 1354

 Locating Information about Logical Interfaces | 1355

 Show Logical Interface Information Table | 1355

Monitoring the Status of Wireless Controllers, Access Points, and Radios | 1356

Monitoring the Status of a Virtual Chassis | 1357

Monitoring the Status of Virtual Chassis Members | 1358

Monitoring and Analyzing Fabrics | 1360

Monitoring Junos Fusion Fabric Systems and Components | 1360

Analyzing QFabric Devices | 1361

 Procedure for Analyzing a QFabric Device Manually | 1362

 Using the Fabric Health Check Tab | 1362

 Using the Topology Tab | 1363

Monitoring QFabric Devices and Components | 1365

Analyzing Virtual Chassis Fabrics | 1367

 Procedure for Analyzing a Virtual Chassis Fabric | 1367

 Using the Fabric Health Check Tab | 1367

 Using the Topology Tab | 1368

Monitoring Virtual Networks | 1371

Using Monitor Mode for Virtual Devices | 1371

 Current Active Alarms Monitor | 1372

 Status Monitor | 1373

 Hosts By % Bandwidth Utilization | 1374

 Top VMs By Bandwidth Utilization | 1374

 Host NIC Bandwidth Utilization | 1375

 Virtual Switch Summary By Version | 1375

 Virtual Machine Bandwidth Utilization Trend | 1375

Viewing vMotion History in Network Director | 1376

General Monitoring | 1378

Selecting Monitors To Display on the Summary Tab | 1378

Changing Monitor Polling Interval and Data Collection | 1379

Pinging Host Devices | 1379

Troubleshooting Network Connections Using Traceroute | 1381

Monitor Reference | 1383

802.11 Packet Errors Monitor | 1384

Access vs. Uplink Port Utilization Trend Monitor | 1385

AP Status Monitor | 1386

Current Sessions Monitor | 1388

Current Sessions by Type Monitor | 1388

Current Sessions by Type | 1389

Current Session Details | 1389

Current SSID Statistics Monitor | 1389

Current SSID Statistics Summary | 1389

SSID Statistics Details | 1390

Error Trend Monitor | 1391

Error Trend | 1391

Error Trend Details | 1392

Equipment Status Summary Monitor | 1393

Equipment Summary By Type Monitor | 1394

Equipment Summary By Type | 1394

Equipment Summary By Type Details | 1394

Node Device Summary Monitor | 1395

Percentage of Packets Retransmitted Monitor | 1396

Port Status Monitor | 1397

Port Status Summary | 1397

Port Status Details | 1397

Port Status for IP Fabric Monitor | 1399

Port Utilization Monitor | 1399

Power Supply and Fan Status Monitor | 1400

Power Supply and Fan Status | 1401

Power Supply and Fan Status Details | 1401

- QFabric Director Status Monitor | **1401**
- QFabric Interconnect Status Summary Monitor | **1402**
- QFabric VM Status Summary Monitor | **1403**
- Radio Status Monitor | **1403**
- Radio Technology Type Statistics Monitor | **1405**
 - Radio Technology Type Statistics Summary | **1405**
 - Radio Technology Type Statistics Details | **1406**
- Resource Monitor For Wireless LAN Controllers | **1407**
 - Resource Utilization Summary | **1407**
 - CPU and Memory Utilization Charts | **1407**
- Resource Utilization Monitor for Switches, Routers, Virtual Chassis, Virtual Chassis Fabrics, and QFabric Systems | **1408**
 - Resource Utilization Summary | **1408**
 - Resource Utilization Details | **1409**
- RF Interference Sources Monitor for Wireless Devices | **1409**
- RF Interference Sources Monitor For an Access Point | **1412**
- RF Throughput or Packet Throughput Level Monitor | **1413**
- Session Trends Monitor | **1415**
 - Session Trends | **1416**
 - Session Details | **1416**
- Signal-to-Noise Ratio Monitor | **1418**
 - Monitoring Signal-to-Noise Ratio | **1418**
 - Signal-to-Noise Ratio Details | **1419**
- SNR SSID Statistics Monitor | **1420**
 - SNR SSID Statistics Summary | **1421**
 - SNR SSID Statistics Details | **1421**
- Status Monitor for Junos Fusion Systems | **1421**
- Status Monitor for Layer 3 Fabrics | **1422**
- Status Monitor for QFabric Directors | **1423**
- Status Monitor for QFabric Systems | **1424**
- Status Monitor for QFabric Interconnects | **1425**
- Status Monitor for QFabric Nodes | **1425**
- Status Monitor for Switches and Routers | **1426**
- Status Monitor for Virtual Chassis | **1427**

- Status Monitor for Virtual Chassis Members | **1428**
- Status Monitor for Wireless Access Points | **1429**
- Status Monitor for Wireless LAN Controllers | **1430**
- Top APs by Session Monitor | **1431**
 - Top APs by Session Summary | **1431**
 - Top APs by Session Details | **1432**
- Top APs by Traffic Monitor | **1432**
 - Top APs by Traffic Summary | **1433**
 - Top APs by Traffic Details | **1433**
- Top Talker - Wireless Devices Monitor | **1433**
 - Top Talker-Wireless Devices Summary | **1434**
 - Top Talker-Wireless Devices Details | **1434**
- Top Talker - Wired Devices Monitor | **1435**
 - Top Talker - Wired Devices Summary | **1435**
 - Top Talker - Wired Devices Details | **1435**
- Top Users Monitor | **1436**
 - Top Users | **1436**
 - Top Session By User Details | **1436**
- Top Sessions by MAC Address Monitor | **1438**
 - Top Sessions | **1438**
 - Top Session by MAC Details | **1438**
- Traffic Trend Monitor | **1439**
- Unicast vs Broadcast/Multicast Monitor | **1440**
- Unicast vs Broadcast/Multicast Trend Monitor | **1441**
- User Session Details Window | **1442**
- Virtual Chassis Topology Monitor | **1443**
- VC Equipment Summary By Type Monitor | **1445**
 - VC Equipment Summary By Type | **1445**
 - VC Equipment Member Status By Type Details | **1445**

Using Fault Mode

About Fault Mode | 1448

Understanding Fault Mode in Network Director | 1448

What Are Events and Alarms? | 1448

Alarm Severity | 1449

Alarm Classification | 1449

Alarm State | 1451

Alarm Notifications | 1451

Threshold Alarms | 1451

Understanding the Fault Mode Tasks Pane | 1452

Using Fault Mode | 1454

Customizing Alarms | 1454

Searching Alarms | 1454

Changing Alarm State | 1457

Fault Reference | 1459

Alarm Detail Monitor | 1459

Finding Specific Alarms | 1460

Sorting Alarms | 1461

Reading Events | 1463

Investigating Event Attributes | 1464

Changing the Alarm State | 1464

Current Active Alarms Monitor | 1464

Alarms by Category Monitor | 1466

Alarms by Severity Monitor | 1466

Alarms by State Monitor | 1467

Alarm Trend Monitor | 1468

Working in Report Mode

About Report Mode | 1470

Understanding Report Mode in Network Director | 1470

Understanding the Report Mode Tasks Pane | 1472

Understanding the Scope of Reports You Can Create | 1474

Creating and Managing Reports | 1475

Managing Reports in Network Director | 1475

How to Locate and Manage Reports | 1476

Managing Report Definitions | 1476

Creating Reports | 1478

How to Create a Report Definition | 1478

Creating a Report Definition | 1480

Setting Report Options | 1482

Reviewing the Report Definition | 1483

Changing a Report Definition | 1484

Scheduling Reports | 1484

How to Create or Manage Schedules | 1485

Managing Schedules | 1485

Creating New Schedules | 1486

Editing Schedules | 1488

Deleting Schedules | 1488

Managing Generated Reports | 1489

Reviewing Generated Reports | 1489

Viewing Report Details | 1490

Exporting Reports | 1491

Deleting Generated Reports | 1491

Retaining Reports | 1491

Managing Reports on SCP Servers | 1492

How to Configure SCP Servers | 1492

Managing SCP Servers | 1493

Mailing Reports | 1494

How to Configure SMTP Servers | 1495

Managing SMTP Servers | 1495

Adding or Editing SMTP Server Settings | 1497

Report Reference | 1498

Active User Sessions Report | 1498

Alarm History Report | 1500

Alarm History Header | 1500

Alarm History Tables | 1501

Alarm Summary Report | 1503

Alarm Summary Header | 1503

Alarm Summary Charts | 1504

Audit Trail Report | 1505

Wireless Client Details Report | 1507

Client Devices Report | 1510

Device Inventory Report | 1511

Fabric Analyzer Report | 1513

Network Device Traffic Report | 1514

Network Device Traffic Report Header | 1515

Network Device Traffic Charts | 1515

Network Neighborhood Report | 1516

Network Neighborhood Report Header | 1517

Network Neighborhood Report Tables | 1517

Radio Traffic Report | 1518

Port Bandwidth Utilization Report | 1519

RF Interference Detail Report | 1521

RF Interference Summary Report | 1522

Rogue Summary Report | 1524

Wireless Security Alarms Report | 1525

Top Users by Data Usage Report | 1526

Top Users by Data Usage Header | 1526

Top Users of Data Table | 1527

Traffic and Congestion Summary Report | 1528

Working with Network Director Mobile

About Network Director Mobile | 1531

Overview of Network Director Mobile | 1531

Getting Started with Network Director Mobile | 1532

Network Director Mobile System Requirements | 1532

Logging Into Network Director Mobile | 1533

Understanding the Network Director Mobile User Interface | 1533

- Dashboard Mode | 1533

- Devices Mode | 1533

Configuring Network Director Mobile Settings | 1534

Working in the Network Director Mobile Dashboard Mode | 1535

Monitoring Network-Wide Activity Using Network Director Mobile | 1535

Network Director Mobile Dashboard Reference | 1535

- Network Summary Monitor | 1536

- Alarms Monitor | 1536

- Top Sessions Monitor | 1537

- Ports Monitor | 1538

- Session Count Monitor | 1538

- Session Trend Monitor | 1538

- RF Interferences Monitor | 1538

Working in the Network Director Mobile Devices Mode | 1540

Locating a Device and Viewing Device Properties Using Network Director Mobile | 1540

Monitoring Sessions on a Device Using Network Director Mobile | 1542

Monitoring Equipment Status on a Wireless LAN Controller Using Network Director Mobile | 1543

Working with Aruba Devices and Applications in Network Director

About Aruba Networks Integration in Network Director | 1548

Understanding Aruba Airwave Integration with Network Director | 1548

Managing Aruba Devices and Applications in Network Director | 1550

Viewing Aruba Wireless Device Inventory in Network Director | 1550

Linking to the Aruba Airwave Application | 1553

Launching the Aruba Airwave Application | 1554

Launching the Aruba Airwave Application from Build Mode | 1555

Launching the Aruba Airwave Application from Monitor Mode | 1555

Launching the Aruba Airwave Application from Fault Mode | 1556

Launching the Aruba Airwave Application from Reports Mode | 1556

Launching Aruba Airwave for an Individual Aruba Wireless Device or Device Group | 1556

Guide 3 Juniper Networks Data Center Switching Management Pack for VMware vRealize Operations (vROps) User Guide

Juniper Networks Data Center Switching Management Pack for vROps

Understanding Juniper Networks Data Center Switching Management Pack for vROps | 1561

Benefits of Juniper Networks Data Center Switching Management Pack for vROps | 1563

Adding and Configuring Juniper Networks Data Center Switching Management Pack for vROps | 1563

Adding the Juniper Networks Data Center Switching Management Pack for vROps | 1564

Specifying the Network Director Credentials in vROps | 1565

Specifying the VMware vCenter Details in vROps | 1566

Creating a Read-Only User in vROps | 1567

Adding a VMware vCenter in Network Director | 1568

Monitoring Juniper Networks Devices from vROps | 1569

Using the Juniper Infrastructure Overview Dashboard | 1570

View the Data Center Connectivity | 1571

View the Top Alerts for your Data Center | 1573

View Relationship Between Various Devices in the Data Center | 1574

Using the Juniper Network Fabric Monitoring Dashboard | 1575

View Data Center Fabric Details | 1576

View CPU and Memory Utilization History and Forecast of a Fabric | 1578

Using the Juniper Network Fabric Member Monitoring Dashboard | 1578

View Fabric Member Details | 1579

View CPU Utilization and Memory Utilization of a Fabric Member | 1580

Using the Juniper Top Network Fabrics Dashboard | 1580

Using the Juniper Top Network Fabric Members Dashboard | 1581

Managing Juniper Networks Data Center Infrastructure from vROps | 1582

Open the Juniper Networks Data Center Infrastructure View | 1582

View Data Center Details | 1583

Open Network Director from vROps | 1584

Performing Fault Management in vROps | 1586

Configuring Thresholds in vROps | 1587

Modifying the Polling Interval in vROps | 1588

About the Documentation

IN THIS SECTION

- Documentation and Release Notes | **iii**
- Supported Platforms | **iii**
- Documentation Conventions | **iii**
- Documentation Feedback | **iv**
- Requesting Technical Support | **iv**

Documentation and Release Notes

To obtain the most current version of all Juniper Networks[®] technical documentation, see the product documentation page on the Juniper Networks website at <https://www.juniper.net/documentation/>.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at <https://www.juniper.net/books>.

Supported Platforms

For the features described in this document, the following platforms are supported:

Documentation Conventions

[Table 1](#) defines notice icons used in this guide.

Table 1: Notice Icons

Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

Table 2 defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies guide names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS CLI User Guide</i> RFC 1997, <i>BGP Communities Attribute</i>

Table 2: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level. The console port is labeled CONSOLE.
< > (angle brackets)	Encloses optional keywords or variables.	stub <default-metric <i>metric</i>>;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (<i>string1</i> <i>string2</i> <i>string3</i>)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	community name members [<i>community-ids</i>]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	

GUI Conventions


Table 2: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> In the Logical Interfaces box, select All Interfaces. To cancel the configuration, click Cancel.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

We encourage you to provide feedback so that we can improve our documentation. You can use either of the following methods:

- Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the [Juniper Networks TechLibrary](#) site, and do one of the following:



- Click the thumbs-up icon if the information on the page was helpful to you.
- Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are

covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <https://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <https://www.juniper.net/customers/support/>
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>
- Download the latest versions of software and review release notes: <https://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <https://www.juniper.net/company/communities/>
- Create a service request online: <https://myjuniper.juniper.net>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: <https://entitlementsearch.juniper.net/entitlementsearch/>

Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit <https://myjuniper.juniper.net>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see <https://support.juniper.net/support/requesting-support/>.

Network Director Quick Start Guide

1

CHAPTER

Installing and Upgrading Junos Space Network Director

Network Director Installation Overview | **59**

Setting up a Junos Space Appliance for Network Director | **59**

Upgrading Junos Space Network Management Platform | **60**

Installing Network Director | **61**

Upgrading Network Director | **64**

Uploading DMI Schemas | **67**

Preparing Devices for Management by Network Director | **68**

Discovering Devices | **69**

Next Steps | **71**

Network Director Installation Overview

Junos Space Network Director is Juniper Networks' network management solution that empowers administrators to seamlessly manage Juniper Networks WLC Series Wireless LAN Controllers (WLCs), EX Series Ethernet Switches, EX Series switches with ELS, QFX Series switches, MX Series routers with ELS, QFabrics, Junos Fusion, and VMware virtual networks within a Junos Space application. Network Director provides a single interface for managing the entire management life cycle of the network, including configuration, device management, monitoring, fault management, and reporting.

This chapter describes how you can install Network Director, upgrade Network Director, bring your devices under Network Director management.

These installation steps are intended for network operators and administrators who install, configure, and manage Junos switching with EX Series, QFX Series, QFabric, and virtualized devices using Network Director.

Before you install Network Director, you must configure the Junos Space Appliance as a Junos Space node.

You can install Network Director in either Juniper Networks JA2500 Junos Space Hardware Appliance or a Junos Space Virtual Appliance. For details, see [“Setting up a Junos Space Appliance for Network Director” on page 59](#).

Setting up a Junos Space Appliance for Network Director

You can install Network Director in one of the following appliances:

- Juniper Networks JA2500 Junos Space Hardware Appliance—The JA2500 appliance is a dedicated hardware device that provides the computing power and specific requirements to run Network Director and the Network Director API as applications.

The JA2500 appliance has a 2-U, rack-mountable chassis with dimensions 17.81 in. x 17.31 in. x 3.5 in. (45.2 cm x 44 cm x 8.89 cm). The JA2500 appliance ships with a single AC power supply module; an additional power supply module can be installed in the power supply slot in the rear panel of the appliance. The JA2500 appliance can also be powered on by using one or two DC power supply modules. The appliance has six 1-TB hard drives arranged in a RAID 10 configuration. Two externally accessible cooling fans provide the required airflow and cooling for the appliance.

For details about the JA2500 appliance and instructions for installation, see [Juniper Networks JA2500 Junos Space Appliance](#).

- Junos Space Virtual Appliance—The Junos Space Virtual Appliance consists of preconfigured Junos Space Network Management Platform software with a built-in operating system and application stack that is easy to deploy, manage, and maintain. A Junos Space Virtual Appliance includes the same software and provides all the functionality available in a Junos Space physical appliance. However, you must deploy the virtual appliance on the VMware ESX or ESXi server, which provides a CPU, hard disk, RAM, and a network controller, but requires installation of an operating system and applications to become fully functional.

For information about installing Junos Space appliances in a fabric configuration and installing Junos Space Virtual Appliance on a VMware ESX or ESXi server, see [Junos Space Virtual Appliance](#).

Upgrading Junos Space Network Management Platform

You can install Network Director Release 4.0R1 newly only on Junos Space Network Management Platform Release 20.1R1. If you are using Junos Space Network Management Platform Release 20.1R1, you can skip this procedure and begin installation of Network Director.

If you are using a Junos Space Platform release that is earlier than the supported release, you need to upgrade Junos Space Platform before installing Network Director. To determine the Junos Space Platform release version and to upgrade Junos Space Network Management Platform, follow these steps:

1. Determine the installed Junos Space Platform version:
 - a. Log in to Junos Space by using the default username and password for Junos Space: **super** and **juniper123**.
Junos Space opens the dashboard.
 - b. Click the plus symbol (+) next to Administration to expand the Administration menu.
 - c. Click **Applications** to list all of the applications installed.
 - d. Note the version of the Junos Space Platform or the Network Application Platform. (Some earlier versions of the Network Management Platform were named Network Application Platform.) If the currently installed release is a supported one, you can skip the rest of this procedure; if not, you must upgrade Junos Space Platform to a supported release.
2. Upgrade Junos Space Network Management Platform to 20.1R1. For upgrade steps, see [Junos Space Network Management Platform Upgrade Instructions](#).

NOTE: Make sure that you are running Junos Space Platform Release 19.3R1 or Release 19.4R1. You can upgrade to Junos Space Platform Release 20.1R1 only from Junos Space Platform Release 19.3R1 or Release 19.4R1.

If the Junos Space Platform runs a version earlier than Release 17.1R1, you must first upgrade Junos Space Platform to Release 18.1R1. For details about upgrading Junos Space Platform to Release 18.1R1, see [Upgrading to Junos Space Network Management Platform Release 18.1R1](#).

Installing Network Director

IN THIS SECTION

- [Installing Network Director From Junos Space Store | 62](#)
- [Installing Network Director by Manually Downloading the Network Director Application Image | 63](#)

Before you begin:

- Configure a Junos Space Appliance or a Junos Space Virtual Appliance as a Junos Space node or as a specialized node used for fault monitoring and performance monitoring (FMPPM). For more information, see [Configuring a Junos Space Appliance as a Junos Space Node](#) for configuring a Junos Space Appliance or see [Configuring a Junos Space Virtual Appliance as a Junos Space Node](#).
- Upgrade Junos Space Platform to Release 20.1R1. For upgrade steps, see [“Upgrading Junos Space Network Management Platform” on page 60](#).
- Uninstall Connectivity Services Director or Edge Services Director before you install Network Director on your system. Network Director cannot be installed on a system that has Connectivity Services Director or Edge Services Director already installed. For steps to uninstall a Junos Space Application, see [Uninstalling a Junos Space Application](#).

You can install Network Director on Junos Space Network Management Platform by using one of the following methods:

NOTE: OpenNMS is disabled automatically when Network Director is installed in Junos Space Network Management Platform 20.1R1. This is applicable only for a fresh installation.

Installing Network Director From Junos Space Store

Starting Release 18.2R1, Junos Space Platform provides Junos Space store from where you can download and install or upgrade Network Director in a single operation. On the Junos Space store page, you can view the versions of Network Director that are compatible with the currently installed version of Junos Space Platform.

NOTE: Before you install or upgrade Network Director by using Junos Space store, you must configure the credentials to access Junos Space Store. For information see, *Configuring and Managing Junos Space Store*.

To install Network Director from Junos Space Store:

1. Click **Administration > Applications > Junos Space Store**.

The Junos Space Store opens. Junos Space Store lists all the applications that can be installed on the Junos Space Platform.

2. Click **Network Director**.

The right-side of the page lists the Network Director versions that can be installed on Junos Space Platform.

3. (Optional) Select the **Show only compatible version** check box to list only the compatible versions of Network Director that can be installed on the current installed version of Junos Space Platform.

4. Click **Next** to install Network Director.

The end user license agreement page appears.

5. Click **Accept and install** to install Network Director.

The Network Director installation job status appears. The status indicates each step that is completed while Network Director is getting installed.

Once installed successfully, Network Director is listed on the Applications page (**Administration > Applications**).

Installing Network Director by Manually Downloading the Network Director Application Image

Download the Junos Space Network Director Release 4.0R1 software image to the hard disk or to an SCP server. The SCP server where you download the Network Director image should be a Linux server. You can download the Network Director Software image from the [Network Director Download Software](#) page.

To install Junos Space Network Director:

1. Log in to Junos Space.
2. Click the add symbol (+) adjacent to the Administration and click **Applications**.

The Administration > Applications page opens.

3. Click add symbol (+) symbol to add the Network Director application.

The Add Application page opens.

4. You can upload the Network Director release image file by using HTTP or by using SCP:

To upload the image file using HTTP:

- a. Click **Upload via HTTP**.

The Upload Software via HTTP page opens.

- b. Click **Browse** to select the Network Director image file. You can either navigate to the local directory and select the Network Director software image, or copy and paste the download URL in the **Software File** if the image is not already downloaded to the local directory.

Your browser opens a dialog box to browse the Network Director image file.

- c. Click **Open** to download the image file.

- d. Click **Upload** to upload the image file.

A notification about the progress in the upload is displayed.

To upload the image file using SCP if you have a Linux server:

- a. Click **Upload via SCP**.

The Upload Software via SCP page opens.

Enter the following secure copy credentials to upload the image from a remote server to Junos Space.

- Enter the user name of the remote server.
- Enter the password of the remote server and reenter the password in the Confirm Password field.
- Enter the host IP address of the remote server.
- Enter the path of the remote server to which you have copied the Network Director image file.

- b. Click **Upload** to load the Network Director image file into Junos Space.

The Upload Application Job Information dialog opens.

5. Click **OK** to skip viewing the job results and to take you back to the Administration > Applications > Add Application page.

6. Select Network Director and click **Install**.

7. Click **OK** in the Application Configuration window dialog box.

You can check the Job Status page to view the progress of the installation job. Once the installation completes, Network Director appears on the Applications page. The Network Director application also appears in the Application Chooser (at the upper-left corner).

8. (Optional) Bookmark this page in your browser for future use.

You can use the bookmarked URL to log in to Network Director without logging in to Junos Space first.

Upgrading Network Director

You can upgrade to Network Director Release 4.0R1 from the following Network Director releases:

- Upgrading from Network Director 3.9R1
- Upgrading from Network Director 3.8R1

- Upgrading from Network Director 3.7R1
- Upgrading from Network Director 3.6R1

If you do not have a supported version of Network Director, upgrade to Release 3.6R1, 3.7R1, 3.8R1, or Release 3.9R1. For instructions on upgrading to Network Director Releases 3.6R1, 3.7R1, 3.8R1, or 3.9R1, see the following:

- [Network Director Release 3.6 Quick Start Guide](#)
- [Network Director Release 3.7 Quick Start Guide](#)
- [Network Director Release 3.8 Quick Start Guide](#)
- [Network Director Release 3.9 Quick Start Guide](#)

Before you start the upgrade, ensure that you have:

- Disabled monitoring for all categories in the Monitoring tab of the Preferences page. For more details, see [Disabling Data Collection for Monitors](#).
- Taken a back up of your database using the Junos Space backup feature. For more details, see [Executing the Data Back Up Procedure](#).
- Junos Space Release 20.1R1 running on your appliance. If your appliance is running an unsupported release of Junos Space Platform, you must upgrade Junos Space Platform before installing Network Director. For step-by-step instructions on upgrading Junos Space, see [“Upgrading Junos Space Network Management Platform” on page 60](#).

You can upgrade to Network Director Release 4.0R1 either by using the Junos Space Store option under **Administration > Applications** task or by manually downloading the Network Director software image.

To upgrade Network Director by using Junos Space Store, see [“Installing Network Director From Junos Space Store” on page 62](#).

To upgrade Network Director from a previous version by manually downloading the software image:

1. Download the Network Director Release 4.0R1 software image to the hard disk or to an SCP server. You can download the Network Director Software image from the [Network Director Download Software](#) page.
2. Log in to Junos Space Platform.
3. Click the add symbol (+) adjacent to the Administration and click **Applications**.
The Applications page opens.
4. Click add symbol (+) symbol to add the Network Director application.
The Add Application page opens.

5. Select Network Director from the list of installed applications and click **Upgrade Application** from the Actions menu.
6. In the Upgrade Application page, click either **Upload via HTTP** or **Upload via SCP** and navigate to the location where you saved the Network Director image.

To upload Network Director by using HTTP:

- a. Click **Upload via HTTP**.

The Upload Software via HTTP page opens.

- b. Click **Browse** to select the Network Director image file. You can either navigate to the local directory and select the Network Director software image, or copy and paste the download URL in the **File name** if the image is not already downloaded to the local directory.
- c. Click **Open** to download the image file.
- d. Click **Upload** to load the image file into Junos Space.

To upload Network Director by using SCP if you have a Linux server:

- a. Click **Upload via SCP**.

The Upload Software via SCP page opens.

Enter the following secure copy credentials to upload the image from a remote server to Junos Space.

- Enter the user name of the remote server.
- Enter the password of the remote server and reenter the password in the Confirm Password field.
- Enter the host IP address of the remote server.
- Enter the path of the remote server to which you have copied the Network Director image file.

Click **Upload** to load the image file into Junos Space.

The Upload Application Job Information dialog opens.

7. Click **OK** to skip viewing the job results.

8. Select Network Director and click **Upgrade**.

You can check the Job Status page to see the progress of the upgrade job. Once the upgrade completes, Network Director appears on the Applications inventory page. The new or upgraded application also appears in the Application Chooser (at the upper-left corner).

9. (Optional) Bookmark this page in your browser for future use.

You can use the bookmarked URL to log in to Network Director without logging in to Junos Space first.

Uploading DMI Schemas

Junos Space Network Management Platform interfaces with network devices using an open API called the Device Management Interface (DMI), which is a standard interface used by Juniper Networks devices. Each device type is described by a unique data model (DM) that contains all the configuration data of the device. The DMI schema lists all the possible fields and attributes for a type of device. The newer schemas describe the new features coming out with recent device releases. It is important that you load all your device schemas into Junos Space Network Management Platform; otherwise only a default schema is applied when you try to edit a device configuration by using the device configuration edit action in the Devices workspace.

Typically, when you perform a clean installation of Junos Space Platform, a schema (usually the latest one) is automatically set as the default for each device family. If an exact matching schema is not available, the default schema for the device family is used.

For the list of DMI schema that you can obtain and upload in Junos Space before you start working on Network Director 4.0R1 Release, see [Junos Space DMI Schema Requirements for Network Director](#).

If you cannot find the schema equivalent, use the latest DMI schema from the main release or contact the [Juniper Support](#). For example, for an EX4500 switch running Junos OS Release 13.2X51-D20, you must use the Junos OS Release 13.2X51-D20 schema. If this is not available, you can use the latest schema available from the Junos OS Release 13.2X51 releases.

You can download the schema from [Schema Repository](#).

To install or update a DMI schema on Junos Space Platform, see [Adding Missing DMI Schemas or Updating Outdated DMI Schemas by Using the Update Schema Menu](#).

Preparing Devices for Management by Network Director

To discover and manage devices, Network Director requires the following minimum device configuration as a prerequisite. Ensure that the device:

- Has a static management IP address. The address can be in-band or out-of-band, but must be reachable from the Junos Space server.
- Is enabled for SSH v2. On wireless LAN controllers, SSH is enabled by default. However, on EX Series switches you need to enable SSH. Issue the **set system services ssh protocol-version v2** command to enable SSH v2 on EX Series switches.
- Has a user ID with the superuser class configured. Junos Space and Network Director uses this user ID to authenticate the SSH connection with the device.
- Is enabled for SNMP with the appropriate read-only V1, V2, and V3 credentials created. You do not need to configure SNMP trap receivers; Network Director configures traps as a deployment task.

In addition, the following protocol ports must be open for Network Director communication:

- Port 22 for SSH connections. If you have changed the SSH port to a port other than port 22 on your Network Management Platform, you must change the SSH ports on your managed devices to the port that the Network Management Platform is using.
- Port 443 for virtualization and RingMaster import support. Use port 443 for outbound traffic to vCenter servers.

NOTE: If your RingMaster server uses any port other than port 443, then you must open that port from the Junos Space Network Management Platform server.

- Port 10162 for SNMP traps. Network Director receives traps from managed devices on this port. (After you install Network Director, use Network Director to configure SNMP on your devices to send traps to Network Director on this port.)
- Port 8889 for the management of wireless LAN controllers.
- Port 21 (TCP) and port 69 (UDP) for uploading the software image and configuration file to the FTP server.
- Ports 8774, 9696, 9292, 8777, 35357, and 8776 for accessing OpenStack and VMware NSX APIs.

You can verify whether a port is open by logging in to the Junos Space CLI and using the **nmap** command. For example, to determine whether port 8889 is open on a controller, issue this command:

```
root@space# nmap <IP address of controller> -p 8889
```

Discovering Devices

When you start Network Director for the first time, the system does not have any devices. The first step is to build your network. Even with large networks, Network Director has made this step relatively easy and straightforward. You will add devices to Network Director and the database by using a process called *device discovery*. Once a device is discovered, it shows in the interface and Network Director begins to monitor the device.

Network Director provides a wizard for device discovery. The following example shows the path for device discovery through the wizard. For an alternate path, you can get a step-by-step instruction from the help system.

In this example, we provide an IP address range, and Network Director populates the database with all supported devices within that range.

1. While in the **Build** mode, select **Logical View**, **Location View**, **Device View**, or **Custom Group View** from the View selector.

NOTE: Select **Datacenter View** if you want to view and manage data centers by using Network Director. You must set up a data center from the Datacenter View. Network Director automatically discovers and adds devices that are part of the data center set up to its inventory. There is no separate device discovery task for the Datacenter View. You can also add devices that are discovered from the other views to a data center by editing a data center in the Datacenter View.

2. To discover physical devices, click **Discover Devices** in the Tasks pane. Each mode has a Tasks Pane that displays the actions you can take while in that mode and that particular network view.
3. (Optional) Type a name for the discovery job. The default name is ND Discovery.
4. Click **Add** in the Device Targets window. You can add a single device IP address, a range of IP addresses, an IP subnet, or a hostname. In this example, we select an IP address range.
5. Provide the initial or the lowest IP address value and the ending or highest IP address value for the range and click **Add**. You can have up to 1024 devices in a range. After you click Add, the address range is listed in the Device Targets window.

6. Click **Next** or click **Discovery Options** to proceed to specify the device credentials and method of discovery.
7. Click **Add** in the Device Credentials window and enter the username and password assigned for administrative access.
8. Select **Ping**, **SNMP**, or both as the method of device discovery. Selecting both is the preferred method if the device is configured for SNMP.

If you select SNMP, the Add SNMP Settings dialog box is displayed. In this example, because we run SNMP version 2, we need to provide the community string. Click **Add** to save the setting.

NOTE: You cannot choose a method for device discovery for virtual network discovery.

9. Click **Next** or **Schedule Options** to proceed to schedule the time when discovery is run.

NOTE: Scheduling options are not available for virtual network discovery.

10. Indicate whether to run the device discovery now or set up a schedule to minimize network traffic. In this example, we set the schedule to run during off hours.
11. Click **Review** to review the settings before you exit the wizard.
12. Click **Finish** to complete the discovery setup and to save the settings.
13. Click **View Discovery Status** to view all scheduled and completed jobs. After a job completes, you can click **Show Details** to view further information on any unexpected results.

Next Steps

After your devices are up and synchronized, much of the function in Network Director is automatically enabled. However, there are a few additional tasks that you will need to perform to use all the features of Network Director. We suggest that you explore:

- Set up a Location View

Location View is one of seven different views, or perspectives, in your network. In Location View, you can manage devices based on a site. Here you define the buildings, floors, wiring closets, and outdoor areas. You can upload floor maps for easy reference and assign devices to a specific spot.

To set up a Location View:

1. Click **Build** in the Network Director banner.
2. Select **Location View** in the View pane to the left of the screen.
3. Click **Setup Locations** in the Tasks pane to start setting up your location, buildings, floors, racks, wiring closets, and outdoor areas.

- Enable Trap Forwarding and Alarms for Fault Management

A key component of Network Director is to diagnose problems with precision and ease. Network Director correlates multiple traps from the same device to a single alarm.

You must complete device discovery and the devices must be up before you can enable trap forwarding. Traps are not enabled by default; you need to enable them after device discovery.

1. Use **Set SNMP Trap Configuration** in the Tasks pane of Deploy mode to configure your device to send SNMP traps to Network Director.
2. Review the list of alarms in Preferences, located in the Network Director banner. All alarms are enabled by default, but you might want to disable those alarms that are not pertinent to your installation. You can also change the severity of an alarm by using Preferences.

- Set up users

After you install Network Director, there is only one username defined: *super* with the default password, *juniper123*.

You have the ability to set up users with different Network Director privileges. New Network Director users are set up in Junos Space and follow the roles and privileges as defined in Junos Space. For a complete discussion on how to properly set up users, see [Understanding Network Director User Administration](#).

- Configure Network Director API—The Network Director API is a set of Representational State Transfer (REST) APIs that enable network management functions and is installed when you install Network Director. To know more about configuring Network Director API, see [Setting Up the Network Director API](#).

- Learn what you can do with Network Director

There are two ways you can become familiar with the functions and features of Network Director:

- Use the extensive online help system that guides you through Network Director. Clicking the main Help icon provides a top-down view into the help system; clicking a Help icon on a pane or window provides context-sensitive information. Use the help system to familiarize yourself with Network Director and the different modes and panes in the interface.
- Refer to the [Network Director User Guide](#).

2

CHAPTER

Installing Data Learning Engine

Installing and Configuring Data Learning Engine for Network Director | 74

Installing and Configuring Data Learning Engine for Network Director

IN THIS SECTION

- [Installing DLE | 74](#)
- [What to Do Next | 78](#)

Data Learning Engine (DLE) enables Network Director to collect and analyze high-frequency statistics and sFlow data for devices that are managed by Network Director. Only the QFX Series devices support the analytics feature that is required for generating high-frequency statistics data. Network Director uses high-frequency statistics data to create network heat maps and to monitor latency in QFX devices and sFlow data to monitor network traffic in EX and QFX devices.

This topic contains the following sections:

Installing DLE

DLE runs on a dedicated CentOS server. You can install DLE either directly on a CentOS server or on a virtual machine (VM) that runs CentOS. Following are the system requirements to install DLE:

- The server or the VM on which you install DLE must have:
 - CentOS version 6.10, 64 bit or CentOS version 7.6, 64 bit
 - 16 GB RAM
 - 8 CPUs
 - 100 GB of hard disk space
- The Network Director server, the DLE server, and all the devices that are to be monitored using the analytics feature must be connected over a network, and have the following system time configurations:
 - Configured with the same time zone.
 - System clocks synchronized with a Network Time Protocol (NTP) server.

Before you install DLE make sure you have:

For CentOS Version 6.10

- Verified that CentOS version 6.10 is installed on the server as shown in the following example:

```
[root@user ~]# rpm --query centos-release  
centos-release-6-10.el6.centos.11.1.x86_64
```

Or

```
[root@user ~]# lsb_release -d  
Description: CentOS release 7.6 (Final)
```

For CentOS Version 7.6

- Verified that CentOS version 7.6 is installed on the server as shown in the following example:

```
[root@user ~]# rpm --query centos-release  
centos-release-7-6.1810.2.el7.centos.x86_64
```

Or

```
[root@user ~]# grep OS /etc/centos-release  
CentOS Linux release 7.6.1810 (Core)
```

- Synchronized the time on the DLE server, Network Director server, and the devices using a common NTP server as shown in the following example:

```
[root@user log]# ntpdate -u ntp.example.net  
13 Jan 12:51:02 ntpdate[11386]: adjust time server 192.0.2.1 offset -0.101819  
sec
```

NOTE: You can either specify the domain name/host name or IP address of the server.

NOTE: Juniper Networks recommends that you use an NTP server to synchronize the time between DLE, Network Director, and devices. However, if you do not use an NTP server, you need to synchronize the time manually.

- Verified that the network ports 8080, 4242, 50005, 8282, 8081, 50006, 50009, 9160, 7000, and 9042 are in listening mode by entering the `netstat -anp | grep <port number>` command as shown in the following example:

```
[root@user] # netstat -anp | grep 8282
```

```
tcp        0      0 0.0.0.0:8282          0.0.0.0:*            LISTEN
           1839/java
```

Alternatively, you can disable firewall on the DLE server to make sure that all the network ports are accessible as shown in the following example:

```
[root@user log]# service iptables stop
```

```
iptables: Setting chains to policy ACCEPT: filter      [ OK ]
iptables: Flushing firewall rules:                    [ OK ]
iptables: Unloading modules:                           [ OK ]
```

- Noted down the CentOS server IP address for configuring DLE in Network Director.

To install DLE:

1. Download the DLE RPM package version 15.1X53-D30 from the [Cloud Analytics Engine software download page](#) to your CentOS server.

The RPM file name has the following format:

dle-all-release-identifier.x86_64.rpm—for example, **dle-all-15.1X53-D30.1.x86_64.rpm**

2. Install the DLE RPM package on the CentOS server.

If you have downloaded the DLE RPM package to the tmp folder and you are installing the DLE package from the same (tmp) location, enter the command as shown in the following example:

```
[root@user tmp]# rpm -ivh dle-all-15.1X53-D30.1.x86_64.rpm
```

If you have downloaded the DLE RPM package to the tmp folder and you are installing the DLE package from a different location, enter the command as shown in the following example:

```
[root@user]# rpm -ivh /tmp/dle-all-15.1X53-D30.1.x86_64.rpm
```

A successful installation displays the output as shown in the following example:

```
warning: dle-all-15.1X53-D30.1.x86_64.rpm: Header V3 RSA/SHA256 Signature, key
  ID dc466ab6:
NOKEY
Preparing...                               ##### [100%]

1:dle-all                                  ##### [100%]
Starting Cassandra DB: [ OK ]
Starting KairosDB:      [ OK ]
Starting Data Learning Engine: [ OK ]
Starting cae monitor [ OK ]
```

3. Verify the status of DLE and the database processes by entering the service status commands as shown in the following example:

- [root@user]# service cassandra status

```
cassandra (pid 1483) is running...
```

- [root@user]# service kairosdb status

```
kairosdb (pid 1779) is running...
```

- [root@user]# service dle status

```
dle (pid 1862) is running...
```

4. Verify the DLE version installed on the CentOS server as shown in the following example:

```
[root@user]# rpm -qa |grep dle
```

```
dle-all-15.1X53-D30.1.x86_64.rpm
```

NOTE: You can view the DLE log file at `/opt/cae/dle/log/dle.log` file.

You can run the following commands to view the DLE log file:

```
[root@user tmp]# cd /opt/cae/dle/log
```

```
[root@user log]# tail -f dle.log
```

What to Do Next

After you have installed DLE on the CentOS server, you must perform the following operations to identify the applications that contribute to the traffic, traffic statistics, and the top applications:

- [Configuring the DLE IP in Network Director](#)
- [Enabling the high-freq statistics on the devices](#)
- [Enabling the network traffic analysis](#)
- [Viewing the traffic on a device](#)

Network Director User Guide

1

PART

Working With Network Director

About Network Director | **81**

Accessing Network Director | **98**

Understanding Network Director System Administration and Preferences | **101**

Getting Started with Network Director | **137**

About Network Director

IN THIS CHAPTER

- [Understanding Network Director and the Management Life-Cycle Modes | 81](#)
- [Understanding Wireless Network Management in Network Director | 82](#)
- [Understanding Cloud Analytics Engine and Network Director | 83](#)
- [Understanding the Network Director User Interface | 85](#)

Understanding Network Director and the Management Life-Cycle Modes

Junos Space Network Director enables unified management of Juniper Networks WLC Series Wireless LAN Controllers (WLCs), EX Series Ethernet Switches, EX Series switches with ELS, QFX Series switches, QFX Series switches with ELS, Data Center fabrics (Junos Fusion system, QFabric system, Virtual Chassis fabrics, and Layer 3 fabrics) and cloud-based and bare-metal-based data centers in your network. Providing full network life-cycle management, Network Director simplifies the discovery, configuration, visualization, monitoring, and administration of large networks. You can quickly deploy a network by using it, configure it optimally to improve network uptime and maximize resources, and respond agilely to the needs of applications and users.

The Network Director user interface is based on the network management life-cycle. The interface provides five main working modes that are aligned to the network management life-cycle, and a sixth mode for working with Network Director itself. Each mode provides access to different tasks:

- **Build mode**—Use Build mode to build your network in Network Director. You use Build mode to discover the devices in your network, to create and manage device configurations, and to manage devices. You can also organize your devices into hierarchical groupings based on logical relationships or by physical locations.
- **Deploy mode**—Use Deploy mode to deploy and manage changes to devices. In Deploy mode, you deploy the configurations you built in Build mode, install new software images on your devices, and manage device configuration files.
- **Monitor mode**—Use Monitor mode to gain visibility into your network performance and health. Monitor mode provides a host of information about your network such as the operational status of devices, traffic

patterns and trends, client session statistics, port capacity, and wireless signal throughput and interference patterns. You can also search for a user and view a history of the user sessions.

- **Fault mode**—Use Fault mode to gain visibility into unexpected network events and to manage faults or notifications.
- **Report mode**—Use Report mode to generate reports from the data that Network Director stores on network performance, status, and activity.

In addition to these modes, Network Director enables you to perform system-level tasks from the System button and the Preferences button. System-level tasks include viewing the Network Director user and system audit trail, managing jobs, and gathering logs for troubleshooting.

Benefits of Network Director

- Enables unified management of wired, wireless, and virtual infrastructures for the campus and data center.
- Improves operational efficiency by automating routine management tasks such as device and port provisioning.
- Supports flexible, large-scale deployment of devices. For example, Build mode enables you to apply configurations across multiple devices grouped by logical relationships, physical locations, or type.

RELATED DOCUMENTATION

[Understanding the Network Director User Interface | 85](#)

[Understanding Build Mode in Network Director | 184](#)

[Understanding Deploy Mode in Network Director | 1172](#)

[Understanding Monitor Mode in Network Director | 1269](#)

[Understanding Fault Mode in Network Director | 1448](#)

[Understanding Report Mode in Network Director | 1470](#)

[Understanding Network Director User Administration | 101](#)

Understanding Wireless Network Management in Network Director

Juniper Networks Junos Space Network Director offers wireless LAN management along with the wired network management for simplifying the wireless LAN (WLAN) operations, such as building, deploying, and monitoring a wireless network including the wired devices such as switches and wireless devices such as access points and wireless controllers.

Network Director integrates with Juniper RingMaster software that enable network administrators to plan, configure, and deploy a wireless network comprising hundreds of Juniper Network WLAN controllers and access points. Network Director can also integrate with Aruba Airwave application management platform that enable network administrators to view the Aruba wireless device inventory and launch the Aruba Airwave from within Network Director.

Network Director has the ability to launch both of these platforms from its user interface. To launch these platforms from within Network Director, you must specify URLs for these platforms in Network Director Preferences.

While you can manage the Juniper Networks wireless devices by using both RingMaster and Network Director, you can manage the Aruba wireless devices connected to Juniper switches by using only the Aruba Airwave application. You can import the RingMaster data to Network Director. Network Director uses this data to discover devices, build profiles, configure the Location View, build the Topology View (including a floor plan, if one is defined in the network plan), and to assign devices to various entities in Location View and Topology View.

For information about managing Juniper Networks wireless devices using RingMaster, see the latest *RingMaster User Guide*.

For information about integration of Aruba Airwave with Network Director, see chapter [“Understanding Aruba Airwave Integration with Network Director”](#) on page 1548.

For information about managing Aruba wireless devices by using Airwave management platform, see Aruba documentation.

RELATED DOCUMENTATION

[Setting Up User and System Preferences](#) | 108

[Linking to the Aruba Airwave Application](#) | 1553

[Launching the Aruba Airwave Application](#) | 1554

[Importing RingMaster Data to Network Director](#) | 931

[Network Director Documentation home page](#)

Understanding Cloud Analytics Engine and Network Director

Cloud Analytics Engine from Juniper Networks provides network data analysis to enable you to improve application performance and availability. It provides data collection, analysis, correlation, and visualization, helping you to better understand the behavior of workloads and applications across the physical and virtual infrastructure. Cloud Analytics Engine provides an aggregated and detailed level of visibility, tying

applications and the network together, and an application-centric view of network status, improving your ability to quickly roll out new applications and troubleshoot problems.

Cloud Analytics Engine provides these major capabilities:

- Application visibility and performance management, by analyzing application flows and workload placement.
- Capacity planning and optimization, by detecting hotspots and monitoring latency and microbursts.
- Troubleshooting and root cause analysis, by correlating overlay networks.

NOTE: Cloud Analytics Engine does not support underlay networks.

Network Director uses Cloud Analytics Engine to support the following features:

- High-frequency statistics gathering and reporting—Network Director uses high-frequency statistics to report on device and port latency and congestion events through the Device and Port Latency Heatmap dashboard widget and the View Congestion Events task in Monitoring mode.
- Flow analysis—Network Director provides information about application flows in your data center, such as the path the flow takes through the network, the latency experienced at each hop, and traffic statistics for each network device in the path. The Virtual Machines & Bare Metal Servers dashboard widget and the Recent Flow Analysis dashboard widget are the primary interfaces for this feature.

To use these features, you must:

- Install the following Cloud Analytics Engine components in your network:
 - QFX5100 switches running a Junos OS release that supports Cloud Analytics Engine.
 - One or more compute nodes with a Compute Agents (CA) installed. Compute Agent works with the Junos OS component of Cloud Analytics Engine to configure cloud analytics data collection on network devices and collect the requested data.
 - One or more Data Learning Engines (DLEs) installed. The DLE provides the API for integration with Network Director.

For more information about these components and installing them, see the Cloud Analytics Engine documentation included in the [QFX Series documentation](#).

- Provide Network Director with the IP address of the DLE or DLEs under **Preferences > Monitoring > Data Learning Engine Settings**.

In addition:

- For high-frequency statistics, you must enable the collection of the statistics on specific devices or ports, using the **Configuration Deployment > Enable High Frequency Stats** task in Deploy mode.

- For flow analysis, you must enable LLDP on the bare metal servers, on the servers hosting the virtual machines, and on the connecting switches. In addition, the switches must be discovered by using the SNMP option in Network Director.

RELATED DOCUMENTATION

[Device & Port Latency Widget | 150](#)

[Viewing Congestion Events | 1312](#)

[Virtual Machines & Bare Metal Servers Widget | 169](#)

[Recent Flow Analysis Widget | 160](#)

Understanding the Network Director User Interface

IN THIS SECTION

- [Network Director Banner | 86](#)
- [View Pane | 88](#)
- [Tasks Pane | 91](#)
- [Alarms | 93](#)
- [Main Window or Workspace | 93](#)
- [Tables in Network Director | 93](#)

Junos Space Network Director provides a simple to use, HTML5-based, Web 2.0 user interface that you can access through standard Web browsers. The user interface is task-oriented, using task-based workflows to help you accomplish administrative tasks quickly and efficiently. It provides you the flexibility to work with single devices or with multiple devices grouped by logical relationship, location, or device type. You can filter, sort, and select columns in tables, making looking for specific information easy.

Figure 1 illustrates the main components of the interface.

Figure 1: The Network Director User Interface Components

The screenshot shows the Junos Space Network Director interface. At the top is the 'Network Director Banner' with navigation options: Logical View, Deploy, Monitor, Fault, and Report. Below the banner, the interface is split into three main areas:

- View Pane (Left):** A tree view showing the network hierarchy, including 'My Network', 'Wireless Network', and various device groups like 'WLC800R-7.7' and 'WLC800R-8.0'.
- Tasks Pane (Middle-Left):** A list of tasks such as 'Device Discovery', 'View Discovery Status', 'Device Management', and 'Wired'.
- Main Window (Center):** A table titled 'Device Inventory for My Network' showing details for various devices. The table includes columns for Hostname, IP Address, Serial Number, Platform, OS Version, Device Family, Device Type, Connection State, and Config State. Above the table are three pie charts: 'Devices by Family', 'Connection State', and 'Configuration State'.

At the bottom, there is an 'Alarms' section with a status bar showing 'Critical', 'Major', 'Minor', and 'Info' counts. The date and time 'May 02 2014 10:38 AM EST' are displayed in the bottom right corner.

This topic describes:

Network Director Banner

Use the Network Director banner, shown in Figure 2, to select the working mode. You can also use the Network Director banner to perform other global tasks, such as setting up your preferences or accessing Junos Space. Table 3 describes the functions available to you on the banner.

Figure 2: Network Director Banner

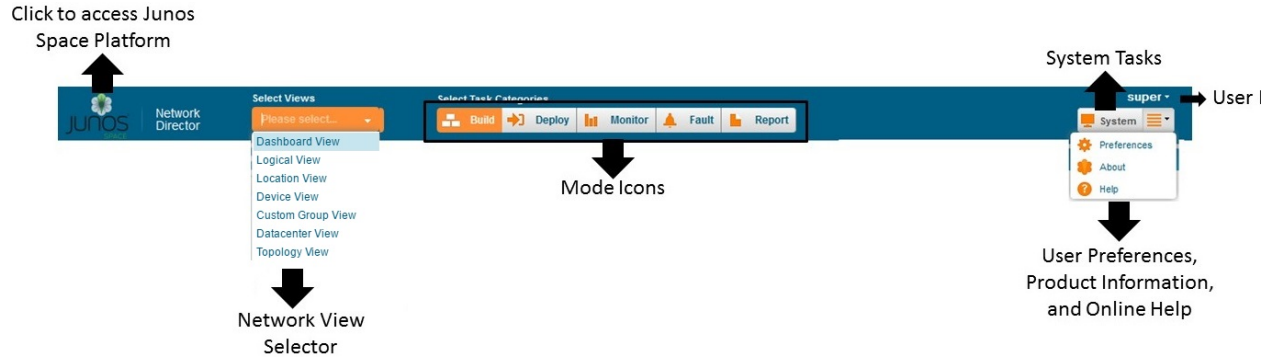



Table 3: Network Director Banner Functions

Item	Function
Accessing Junos Space Platform	Click to exit Network Director and open the Junos Space Network Application Platform. You can switch back and forth between Network Director and Junos Space without logging in again.
Network View Selector	<p>Select the network view that you want to work in. You can choose from one of the following views:</p> <ul style="list-style-type: none"> • Dashboard View • Logical View • Location View • Device View • Custom Group View • Datacenter View • Topology View <p>For more details, see “Displaying Devices in Various Network Views” on page 89.</p>
Mode Icons	<p>Select the mode you want to work in.</p> <p>NOTE: You might not have access to all the Network Director modes. What modes you have access to depends on your assigned user role.</p>
User Log out	<p>Displays the username using which you logged in to Network Director.</p> <p>Click the Down arrow next to the username and select Logout to log out of Network Director and Junos Space.</p>

Table 3: Network Director Banner Functions (*continued*)

Item	Function
System Tasks	<p>Access the system tasks such as viewing audit logs, jobs, and to collect troubleshooting logs.</p> <p>Click the Down arrow next to System and select Preferences to set your Network Director user and system preferences.</p>
System Preferences, Product Information, and Online Help 	<p>Click this button and select an appropriate option:</p> <ul style="list-style-type: none"> • Preferences—Enables you to set your Network Director user and system preferences. • Help—Open searchable help. This help icon is not context-sensitive—it always opens help to the first page. From here, you can browse or search the help. Context-sensitive help is available from the help icon provided on each pane or page. • About—Displays information about Network Director, such as the currently running version.

In addition to this, Network Director displays the date and time in the local time zone in the bottom right corner.

View Pane

IN THIS SECTION

- [Displaying Devices in Various Network Views | 89](#)
- [Filtering the Network Tree | 90](#)
- [Expanding or Collapsing Nodes in the Network Tree | 90](#)
- [Searching the Network Tree | 91](#)

In the View pane, Network Director provides you a unified, hierarchal view of your wired, wireless, and data center networks in the form of a expand tree that is expandable and collapsible. By selecting both a view and a node in the tree, you indicate the *scope* over which you want an operation or task to occur. For example:

- By selecting the Access node in Logical View, you indicate that the scope for a task is all access switches under the Access node.

- By selecting a floor node in Location View, you indicate that the scope for a task is all devices belonging to that floor.
- By selecting the EX4200 node in Device View, you indicate that the scope for a task is all EX4200 switches in your network.

You can perform the following actions in the View pane:

Displaying Devices in Various Network Views

Use the selection box in the Network Director banner to choose one of the following network views:

- **Dashboard View**—This is a customizeable view that provides information about your network, and is the default view that opens when you log in. You can select and add monitoring widgets to the Dashboard View based on your requirements. This is the default view that opens when you log in to Network Director.
- **Logical View**—Devices are organized by their logical relationships in the network. All switches appear in the Switching Network and are categorized by their role in the network: access, aggregation, or core. All wireless devices appear in the Wireless Network. Controllers appear under the mobility domains to which they belong and access points appear under the controller or cluster that manages them. All QFX Series switches that are part of a QFabric appear in the Fabric node.

Network Director builds most of this view for you as you discover devices. However, you need to manually assign switches to the access, aggregation, or core categories.


- **Location View**—Devices are organized by their physical locations. You build this view by creating sites, building, floors, aisles, racks, outdoor areas, and then assigning your switches, wireless controllers, access points, and QFabric systems to these locations.
- **Device View**—Devices are organized by device type: switches, wireless LAN controllers, and QFabric systems. Within each device type, devices are organized by device model. For example, all models of EX4200 switches are grouped together under one node in the tree.
- **View**—Displays devices that are part of your network
- **Custom Group View**—If you have defined one or more custom groups, Network Director displays these custom groups in this view. You can manually add devices to a custom group or define a rule to automatically add devices to the custom group once they are discovered in Network Director. The devices are grouped under each custom group.
- **Topology View**—Topology enables you to view all the discovered devices in your network, overlaid on a map where the devices are located across sites, buildings, floors, closets, aisles, and racks along with their physical interconnection with other devices in your network. Topology also provides visualization around physical and logical connectivity between various discovered interconnected devices.
- **Datacenter View**—Datacenter view enables you to view all the data centers in Network Director. Expand each data center to view servers that are part of the data center,

Filtering the Network Tree

To make it easier for you to focus on selected aspects of your network, you can apply predefined filters to your network tree so that only nodes and devices that meet the filter criteria are shown. For example, you can apply a filter so that only devices in a specific building are shown in the network tree in all views.

To apply filters:

1. Click the filter icon:



2. In the Filters dialog box, click **Show available filters**.

The Available Filters section of the dialog box appears.

3. Under Available Filters, click the tab for the view you want to use to define your filter. For example, if you want to filter on devices—that is, show only certain types of devices—click the **Device** tab.


The filters that you can apply are listed below the tab.

4. To select a filter, click its associated plus icon.

The filter appears in the Selected Filters section of the dialog box. You can repeat Steps 3 and 4 until you have selected all the filters you want apply.

5. Click **Apply**.

The Filters dialog box closes and the filters are applied. The filter icon changes appearance to indicate that filters have been applied:



To remove a filter, click the filter icon, click the trash can next to the filter in the Selected Filters list, and click **Apply**.

Expanding or Collapsing Nodes in the Network Tree

To expand a node in the network tree, select the node and then click the **Expand All** icon:



The node you selected and any child nodes under the selected node are expanded to show their contents.

Similarly, to collapse a node in the network tree, select the node and then click the **Collapse All** icon (next to the Expand All icon). The node you selected is collapsed and no nodes under it are shown.

Searching the Network Tree

To quickly find and select a device or device group, use the search function.

To perform a search, type three or more characters into the Search box and click the **Search** icon, as shown in [Figure 3](#).

Figure 3: Performing Search in the View pane



Network Director finds the first instance of a node whose name contains the characters. To find the next instance, click the right arrow.

Searches are not case-insensitive: a search on *wla115* and one on *WLA115* return the same results.

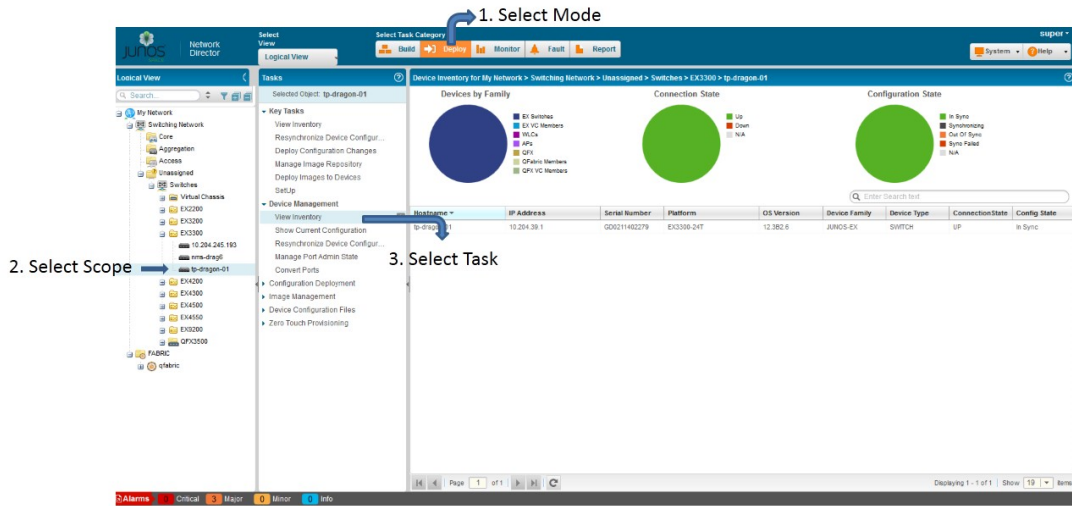
Tasks Pane

The Tasks pane is available in every mode and lists tasks specific to that mode. In addition to changing according to the mode selected, tasks listed in the Tasks pane can change as you select different scopes in the View pane. For example, some tasks are appropriate only at the device level and thus appear only when you have selected an individual device.

Clicking a task brings up task-specific content in the main window.

In general, to perform a task in Network Director, you navigate to the task as shown in [Figure 4](#). You select your mode, your scope, and then your task.

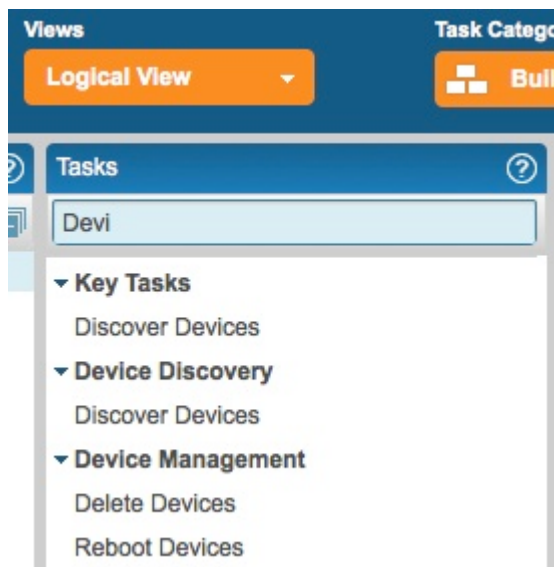
Figure 4: Navigating to a Task in a Tasks Pane



TIP: The location of the Tasks pane changes with mode. In Build and Deploy mode, it is adjacent to the View pane. In Monitor, Fault, and Report mode, it is located to the right of the main window.

Use the Search box in the Tasks pane to easily locate a task, as shown in Figure 5. To perform a search, type three or more characters into the Search box and press Enter.

Figure 5: Performing Search in the Tasks pane



Alarms

The Alarms bar that is displayed at the bottom of your browser window provides a quick summary of how many critical, major, minor, and info alarms are currently active in the network and is visible in every mode. To display more information about alarms, click the alarm count or the Alarms banner. You are automatically placed in Fault mode and the Fault mode monitors are displayed.

Main Window or Workspace

The main window or workspace displays the content relevant to the mode, scope, and task you have selected. When you log in to Network Director, this pane displays the Device Inventory page. The Device Inventory page is the default landing page for Build and Deploy modes. It contains a list of the devices for your current scope. It includes pie charts that permit you to see at a glance the connection states, configuration synchronization states, and device-type distribution for your devices.

Tables in Network Director

IN THIS SECTION

- [Moving and Resizing Columns | 93](#)
- [Displaying the Column Drop-Down Menu | 94](#)
- [Sorting on a Column | 94](#)
- [Hiding and Exposing Columns | 95](#)
- [Searching Table Contents | 95](#)
- [Filtering Table Contents | 97](#)

Tables are used throughout Network Director to display data. These tables share common features. By becoming familiar with these features, you can navigate and manipulate tabular data quickly and efficiently. The following sections describe:

Moving and Resizing Columns

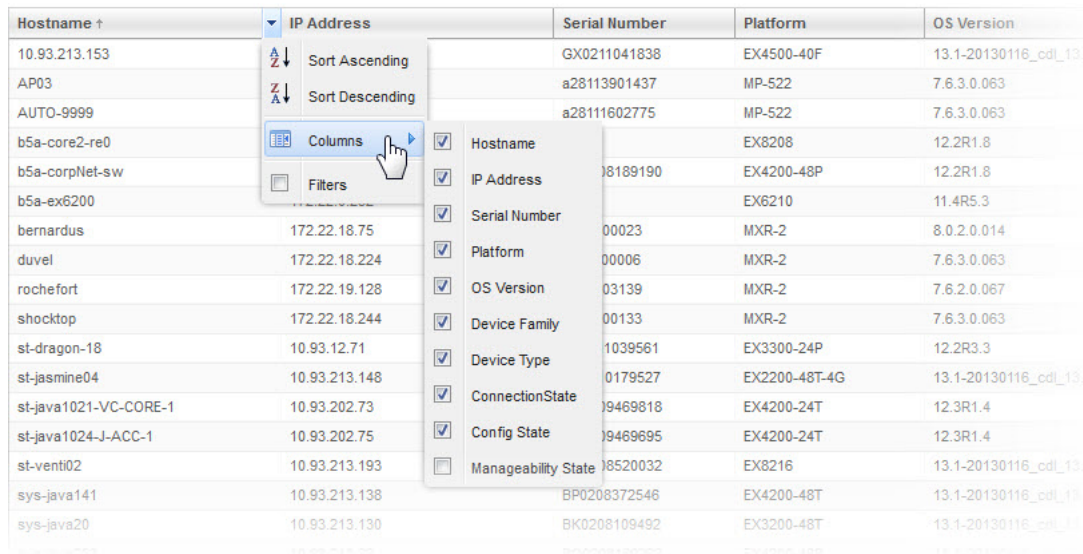
You can reposition and resize columns in a table. To move a column, drag and drop the column head to the new location. Network Director displays a green check mark when you mouse over a valid column location.

To resize a column, mouse over the edge of a column until the cursor becomes two vertical lines with outward arrows. Drag the column width to the new size.

Displaying the Column Drop-Down Menu

A drop-down menu is available from each column head, allowing you to perform additional operations on columns. To display the column drop-down menu, mouse over the column head. A downward arrow appears. By clicking the arrow, you display the drop-down menu, as shown in [Figure 6](#).

Figure 6: Column Drop-Down Menu



Hostname ↑	IP Address	Serial Number	Platform	OS Version
10.93.213.153		GX0211041838	EX4500-40F	13.1-20130116_cdi_13
AP03		a28113901437	MP-522	7.6.3.0.063
AUTO-9999		a28111602775	MP-522	7.6.3.0.063
b5a-core2-re0			EX8208	12.2R1.8
b5a-corpNet-sw		8189190	EX4200-48P	12.2R1.8
b5a-ex6200			EX6210	11.4R5.3
bernardus	172.22.18.75	00023	MXR-2	8.0.2.0.014
duvel	172.22.18.224	00006	MXR-2	7.6.3.0.063
rochefort	172.22.19.128	03139	MXR-2	7.6.2.0.067
shocktop	172.22.18.244	00133	MXR-2	7.6.3.0.063
st-dragon-18	10.93.12.71	1039561	EX3300-24P	12.2R3.3
st-jasmine04	10.93.213.148	0179527	EX2200-48T-4G	13.1-20130116_cdi_13
st-java1021-VC-CORE-1	10.93.202.73	9469818	EX4200-24T	12.3R1.4
st-java1024-J-ACC-1	10.93.202.75	9469695	EX4200-24T	12.3R1.4
st-venti02	10.93.213.193	8520032	EX8216	13.1-20130116_cdi_13
sys-java141	10.93.213.138	BP0208372546	EX4200-48T	13.1-20130116_cdi_13
sys-java20	10.93.213.130	BK0208109492	EX3200-48T	13.1-20130116_cdi_13

Sorting on a Column

You can sort the table on a column by clicking the column head—each click changes the direction of the sort. In addition, you can use the Sort Ascending and Sort Descending options in the drop-down menu.

When you sort on a column, a small arrow appears next to the column name to indicate that the table is being sorted by the column and the direction of the sort.

Network Director uses a lexical sort for tabular data that is not strict numeric data, which means that data such as IP addresses do not sort in numerical sequence, as shown in [Table 4](#).

Table 4: Numerical Sorts and Lexical Sorts

Numerical Sort	Lexical Sort
10.93.200.65	10.93.200.129
10.93.200.129	10.93.200.199
10.93.200.199	10.93.200.65

Hiding and Exposing Columns

You can customize your tables by hiding or exposing columns. This way, you can choose to see only relevant information.

To hide or expose columns, display the drop-down menu for any column head and mouse over the Columns option, as shown in [Figure 6](#). Select a column to expose it—clear a column to hide it.

As a general rule, Network Director displays all columns in a table by default. However, some tables have more columns than can fit easily within the page. In these tables, some columns are hidden by default.

Searching Table Contents

You can search for specific data in large tables by using search criteria.

To search for an item in a table, enter the search term in the text box. Select ANY for Network Director to search for the term in all columns in the table. Every table has a predefined default column that the system searches first; before it proceeds to search other columns.

You can also choose to search a particular column for a term. Network Director displays a list of all the columns in a table. To search a particular column for a term, select that column for the list.

NOTE: When you enter a search expression, note the following:

- You must add a back slash “\” if you want to use the following special characters in the search text:

+ ~ && || ! () { } [] ^ “ ~ * ? : \

- Field names are case-sensitive.

For example, if you have a few systems running on Junos OS 12.3 Release 4.5, then `os: 12.3R4.5` returns search results, whereas `OS: 12.3R4.5` does not return search results. This is because the field name that is indexed is `os` and not `OS`.

- If you want to search for a term that includes a space, enclose the term within double quotation marks.

For example, to search for all devices that are synchronized (that is, In Sync), enter “In Sync” in the Search field.

- You must append “*” if you want to search using partial keywords. Otherwise, the search returns 0 (zero) matches or hits.

You can filter search results by specifying one or more search terms. Network Director uses the AND operator for each search term that you enter. Network Director lists the search results in the table, depending on the search criteria that you specified.

For example, perform the following steps to search for an EX4200 switch that is running Junos OS Release 12.2:

1. Enter **EX4200** as the search term in the text box.
2. From the list that appears, select to search the Platform column.
Network Director lists all the EX4200 switches in your network.
3. Enter **12.2** as the search term after the comma separator in the text box.
4. From the list, select to search from the OS Version column.

Network Director lists all the EX4200 switches in your network that are running Junos OS Release 12.2.

Filtering Table Contents

For large tables, it is helpful to be able to sort data to show only relevant entries. When you mouse over the Filters option in the column drop-down menu, a fill-in box appears where you can type filter criteria. If you type a text string and click **Go**, entries that do not contain the text string (filter criterion) are removed from the table. A red asterisk appears on the column head to indicate that the column has been filtered. To restore all entries to the table, clear the Filters option.

For example, to filter the Device Inventory page so that only devices in the **192.168.1.0** subnet are displayed:

1. Mouse over the right side of the IP Address column head to expose the downward arrow.
2. Click the arrow to display the column drop-down menu.
3. Mouse over **Filters** to display the Filter field.
4. Type **192.168.1.** in the field and click **Go**.

Only the devices in the **192.168.1.0** subnet are shown.

RELATED DOCUMENTATION

[Understanding Network Director and the Management Life-Cycle Modes](#) | 81

[Network Director Documentation home page](#)

Accessing Network Director

IN THIS CHAPTER

- Logging In to Network Director | 98
- Logging Out of Network Director | 99
- Changing Your Password | 100

Logging In to Network Director

You connect to Network Director using your Web browser. The following Web browsers are supported: Internet Explorer versions 8.0 and 9.0, Mozilla Firefox version 3.6 or later, and Google Chrome version 17 and later. The minimum screen resolution is 1280 x 1024.

You can connect to Network Director one of two ways:

- Log in to Network Director directly by using the following URL:

```
https://<n.n.n.n>/networkdirector
```

where *n.n.n.n* is the IP address of the Junos Space Web interface. You can bookmark the login page for future use.

- Log in to Junos Space first by using the following URL:

```
https://<n.n.n.n>/mainui
```

where *n.n.n.n* is the IP address of the Junos Space Web interface.

You can then switch to the Network Director interface by selecting Network Director from the Applications list in the left pane of the Junos Space user interface.

The default username and password is the same for both Junos Space and Network Director:

- username—super
- password—juniper123

RELATED DOCUMENTATION

[Logging Out of Network Director | 99](#)

[Changing Your Password | 100](#)

[Understanding the Network Director User Interface | 85](#)

[Network Director Documentation home page](#)

Logging Out of Network Director

When you are finished using Network Director, log out to prevent unauthorized access. To log out of Network Director, click the username in the Network Director banner and select Logout from the list.

Network Director automatically logs you out if you have not performed any action, such as keystrokes or mouse clicks, for a set period of time. This automatic logout conserves server resources and protects the system from unauthorized access. By default, automatic logout occurs if a session has been idle for 60 minutes.

Network Director uses the same automatic logout period as Junos Space. To change the automatic logout period:

1. Click the System Platform icon.
2. Navigate to **Administration > Applications**.
3. Right-click **Network Management Platform** and select **Modify Application Settings..**
4. In the Modify Network Management Settings page, select **User**.

RELATED DOCUMENTATION

[Logging In to Network Director | 98](#)

[Changing Your Password | 100](#)

[Understanding the Network Director User Interface | 85](#)

[Network Director Documentation home page](#)

Changing Your Password

Any user, regardless of user role, can change his or her password.

Your username and password are the same in Junos Space and Network Director. To change your password, change it in Junos Space:

1. From Network Director, click the Junos Space icon in the Network Director banner.
2. Click the User Password icon in the Junos Space banner.
3. Follow the instructions to change your password.

RELATED DOCUMENTATION

[Logging In to Network Director | 98](#)

[Logging Out of Network Director | 99](#)

[Understanding the Network Director User Interface | 85](#)

[Network Director Documentation home page](#)

Understanding Network Director System Administration and Preferences

IN THIS CHAPTER

- [Understanding Network Director User Administration | 101](#)
- [Understanding the System Tasks Pane | 102](#)
- [Audit Logs Overview | 103](#)
- [Viewing Audit Logs From Network Director | 104](#)
- [Managing Jobs | 105](#)
- [Collecting Logs for Troubleshooting | 107](#)
- [Setting Up User and System Preferences | 108](#)

Understanding Network Director User Administration

Network Director uses the user administration features of the Junos Space platform on which it runs. Use Junos Space for tasks such as adding, deleting, and editing user accounts and roles, and changing user passwords. Refer to the Junos Space documentation for information about user administration.

When Network Director is installed, some additional user administration options are available in Junos Space, which are specific to Network Director:

- In addition to the Super Administrator role, the following predefined roles are available for Network Director users:

Network Director - Admin—Has access to all the Network Director tasks. This role is the system administrator role and has full privileges.

Network Director - Engineer—Has access to either all the device management tasks or only those device management sub-tasks to which the Engineer role is mapped. These users can also view the device monitoring and fault management tasks.

Network Director - Monitor—Has access to monitor the network status and performance or view the faults to determine the health of your network and take appropriate action.

Network Director - Configuration Approver—Has access to provide additional privileges to approve the configuration changes in addition to all the tasks that a Config Admin can perform.

Network Director - Image Admin—Has access to all the image management tasks.

Network Director - Config Admin—Has access to create, edit, delete, assign, deploy profiles, and manage fabrics (VCF, QFabric, and IP Fabric).

Network Director - Cloud Admin—Has access to create data center and perform other data center related tasks.

- You can create custom roles to grant users different access rights to the Network Director modes, group, dashboard widgets, and tasks. Users can access only those portions of the navigation hierarchy to which they are explicitly granted access through access privileges.

If you try to log in to Network Director using an account that does not have access rights to any Network Director modes, you will be redirected to Junos Space instead.

Access to Network Director system preferences is controlled by user access rights. For more information, see [“Setting Up User and System Preferences” on page 108](#).

RELATED DOCUMENTATION

[Understanding the Network Director User Interface | 85](#)

[Setting Up User and System Preferences | 108](#)

[Network Director Documentation home page](#)

Understanding the System Tasks Pane

The System Tasks pane provides tasks for viewing audit logs of Network Director user activities, for managing jobs, and for collecting troubleshooting logs.

To access the System Tasks pane, click **System** in the Network Director banner. The tasks are described in [Table 5](#).

Table 5: System Tasks

Task	Description
View Audit Logs	View a history of user activities on Network Director, including log in, log out, and task initiation and completion.

Table 5: System Tasks (continued)

Task	Description
Manage Jobs	View all jobs that are scheduled to run or have been run by Network Director. You can cancel jobs that are in progress or scheduled to run in the future.
Collect Jobs for Troubleshooting	Download a zip file containing logs and troubleshooting data from both Network Director and Junos Space.
Import RingMaster Data	Import data and configurations from RingMaster to Network Director.

RELATED DOCUMENTATION

[Viewing Audit Logs From Network Director | 104](#)

[Managing Jobs | 105](#)

[Collecting Logs for Troubleshooting | 107](#)

[Network Director Documentation home page](#)

Audit Logs Overview

Audit logs provide a record of login history and user-initiated tasks that are performed from the user interface. From the Audit Logs page, you can monitor user login-logout activity over time, track device management tasks, view services that were provisioned on devices, and so forth. Audit logging does not record non-user initiated activities, such as device-driven activities, and is not designed for debugging purposes.

Administrators can sort and filter on audit logs to determine which users performed what actions on what objects at what time. For example, an administrator can use audit log filtering to track the user accounts that were added on a specific date, track configuration changes across a particular type of device, view services that were provisioned on specific devices, or monitor user login-logout activity over time.

Over time, Network Director will archive a large volume of log entries. Such log entries might or might not be reviewed, but they must be retained for a period of time.

The audit logs can be saved to a local server (the server that functions as the active node for Network Director) or a remote network host or media.

RELATED DOCUMENTATION

[Viewing Audit Logs From Network Director | 104](#)

[Network Director Documentation home page](#)

Viewing Audit Logs From Network Director

Audit logs are generated for login activity and tasks that are initiated from the Network Director application. The Audit Logs page displays the logs for all user-initiated activities.

You can do the following on the Audit Logs page:

- Sort, filter, and search the log entries using the standard table manipulation features in Network Director.
- Obtain more information about a log entry by double-clicking the entry or by selecting the entry and clicking **Show Details**. The Audit Log Details window is displayed.
- For a user-initiated task that runs as a job, you can obtain more information about the job by clicking the job ID in the Job ID column.

To display the Audit Logs page:

1. Click **System** in the Network Director banner.
2. Select **View Audit Logs** from the Tasks pane.

The Audit Logs page is displayed with the fields listed in [Table 6](#).

Table 6: Audit Logs Page Fields

Field	Description
User Name	The login ID of the user that initiated the task
User IP	The IP address of the client computer from which the user initiated the task
Task	The name of the task that triggered the audit log
Time	The data and time when the user initiated the task
Result	The execution result of the task that triggered the audit log: <ul style="list-style-type: none"> • Success—Job completed successfully • Failure—Job failed and was terminated • Job Scheduled—Job is scheduled but has not yet started

Table 6: Audit Logs Page Fields (continued)

Field	Description
Description	A description of the audit log
Job ID	The job ID for any task that runs as a job

RELATED DOCUMENTATION

[Audit Logs Overview | 103](#)

[Managing Jobs | 105](#)

[Network Director Documentation home page](#)

Managing Jobs

Network Director enables you to view and manage jobs. You can view the status of completed jobs and cancel the jobs that are scheduled to execute at a later time or jobs that are in progress.

The Job Management page, accessible as a System task, enables you to view and manage all jobs. In addition, Network Director enables you to view special pre-filtered versions of this page from various other tasks, such as View Discovery Status, View Image Deployment Jobs, or View Baseline Mgmt Jobs. These pages contain the same fields (although some fields might be hidden) and have the same functionality as the Job Management page, but they list only those jobs relevant to particular tasks.

To display the Job Management page:

1. Click **System** on the Network Director banner.
2. Select **Manage Jobs** from the Tasks pane. The Job Management page appears.
3. To view the details of a job, select a row and click **Show Details** or double-click a row.
4. To cancel a scheduled job, select a job that is scheduled for a later time or a job that is in progress and click **Cancel**.

The fields in the Job Management page are described in [Table 7](#). To view any hidden column, keep the mouse on any column heading and select the down arrow and then click Columns. Select the check box to display the hidden columns.

NOTE: Details of jobs initiated from Network Director will be available only from Network Director. These jobs will not be listed in the Job Management pane in Junos Space platform and vice-versa.

Table 7: Job Management Page Fields

Field	Description
Job ID	The unique ID assigned to the job
Name	The name of the job
Percent	The percentage of completion of the job
State	The status of the job: <ul style="list-style-type: none"> • Success—Job completed successfully • Failure—Job failed and was terminated • Job Scheduled—Job is scheduled but has not yet started • In progress—Job is has started, but not completed • Cancelled—Job is cancelled
Job Type	The type of the job
Summary	Summary of the job scheduled and executed with status
Scheduled Start Time	The time when the job is scheduled to start
Actual Start Time	The actual time when the job started
End Time	The time when the job was completed
User	The login ID of the user that initiated the task
Recurrence	The recurrent time when the job will be restarted.

RELATED DOCUMENTATION

[Audit Logs Overview](#) | 103

Collecting Logs for Troubleshooting

Network Director enables you to collect logs and other data from both Network Director and Junos Space that can assist in managing and monitoring Network Director servers.

Network Director collects the logs and troubleshooting data into a compressed file that you can download. This file is named **troubleshoot_yyyy-mm-dd_hh-mm-ss.zip**—for example, **troubleshoot_2012-12-21_11-25-12.zip**. The date and time in the file name is the server Coordinated Universal Time (UTC) date and time.

To retrieve troubleshooting data and log files, follow these steps:

1. Click **System** on the Network Director banner.
2. From the Tasks pane, click **Collect Logs for Troubleshooting**. The Collect Logs for Troubleshooting page appears.
3. Click the **Download troubleshooting data and logs from Network Director and Junos Space** link.
Network Director begins collecting the logs and data. It can take a few minutes for Network Director to collect the information and create the zip file.
4. When the standard file download window for your browser opens, save the **troubleshoot_yyyy-mm-dd_hh-mm-ss.zip** file.
5. When you contact the Juniper Technical Assistance Center, describe the problem you encountered and provide the JTAC representative with the **troubleshoot.zip** file.

[Table 8](#) lists the files included in the **troubleshoot_yyyy-mm-dd_hh-mm-ss.zip** file.

Table 8: Log Files in the troubleshooting.zip File

Description	Location
Jboss log files	<code>/var/log/jboss/servers/server1</code>
MSS OS adapter log files	<code>/home/jmp/mssosadpater/var/errorLog/</code>
Daemon log files	<code>/opt/opennms/logs/daemon/</code>
Platform log files	<code>/var/log/platform</code>

Table 8: Log Files in the troubleshooting.zip File (continued)

Description	Location
Access Log Files	<code>/var/log/httpd</code>
Log files for Apache, NMA, Webproxy	<code>/var/log/httpd/</code>
Watchdog log file	<code>/var/log/</code>

RELATED DOCUMENTATION

[Managing Jobs | 105](#)

[Audit Logs Overview | 103](#)

[Viewing Audit Logs From Network Director | 104](#)

[Network Director Documentation home page](#)

Setting Up User and System Preferences

IN THIS SECTION

- [Accessing the Preferences Page | 109](#)
- [Choosing Server Time or Local Time | 110](#)
- [Specifying Search Preferences | 110](#)
- [Enabling Import of Configuration Group Data from Ethernet Design | 110](#)
- [Specifying the Open Clos Server URL | 110](#)
- [Selecting the Approval Mode | 111](#)
- [Setting up Auto-resynchronization Preferences | 112](#)
- [Retaining Network Director Reports | 113](#)
- [Specifying Wireless Preferences | 113](#)
- [Changing Monitor Mode Settings | 113](#)
- [Changing Alarm Settings | 120](#)
- [Modifying Data Center Synchronization Interval Using the Virtualization Tab | 135](#)

Depending on your privileges, the Preferences page displays either user settings or a combination of user settings and system settings. One or more of the following preference tabs appear when you open the Preferences page:

- **User**—All users can choose whether monitors and reports display the local time or the server time.
- **Search**—Network administrators can configure options for search indexing.
- **Config & Deploy**—Network administrators can:
 - choose to enable or disable import of configuration group data into Network Director.
 - specify the Auto Approval or Manual Approval mode for device configuration deployments.
- **Monitoring**—As a network administrator you can change the polling interval for data collection for Monitor mode monitors and enable or disable the internal processes used for data collection. You can also specify the IP address of the Data Learning Engine server, if installed, and the database record retention periods.
- **Fault**—As a network administrator you can enable or disable alarms. They can also set the retention period for alarms and the number of events per alarm.
- **Report**—Network administrators can specify the period of time for which Network Director reports are retained.
- **Wireless**—Network administrators can specify the Aruba Airwave application and RingMaster URLs.
- **Topology**—Network administrators can specify a retention period for the deleted links in Topology view.
- **Virtualization**—Network administrators can modify the synchronization time interval between Network Director and the cloud infrastructure.

This topic describes:

Accessing the Preferences Page

To open the Preferences page, click  in the Network Director banner and select **Preferences** as shown in [Figure 7](#).

Figure 7: Accessing the Preferences Page



The Preferences page opens with User Preferences as the default tab.

Choosing Server Time or Local Time

All users can specify whether Network Director displays local time or the server's time in monitors and reports on the User Preferences tab. The default setting is to display local time. To change the setting to display the server's time:

1. In the Preferences page, select **Use Server Time** from the list.
2. Click **OK** to save your changes or click **Cancel** to close Preferences.

Specifying Search Preferences

Network Director indexes the device inventory data periodically to enable users to perform efficient searches. You can specify a time interval after which Network Director initiates the next indexing on the Search tab. You can also specify to stop indexing while devices are imported into Network Director. If you are running short of system memory, selecting this option helps save some memory and speed up the discovery and import of new devices. By default, this option is selected and the search index update interval is set to 900 seconds.

Enabling Import of Configuration Group Data from Ethernet Design

For Network Director to be able to import configuration group data.

To enable the import of configuration group data:

1. In the Preferences window, select the **Config & Deploy** tab.
2. Select the **Enable migration from Ethernet Design** check box to enable import of configuration group data. By default this check box is not selected.
3. Click **Save** to save and close the preferences.

For detailed steps on importing configuration group data from Ethernet Design, see [“Importing Configuration Data from Junos OS Configuration Groups”](#) on page 1198.

Specifying the Open Clos Server URL

Layer 3 Fabrics can expand your data center network to thousands of ports. Network Director uses Layer 3 protocols and Open Clos architecture to achieve this. Open Clos is typically installed on a separate system. You must specify the URL of the Open Clos system in Network Director before you can create and manage Layer 3 Fabrics.

To specify the URL for the Open Clos server:

1. Open the Config & Deploy tab in the Preferences window.
2. Enter the URL to the system that is running Open Clos.
3. Click **OK**.

Selecting the Approval Mode

Use the Config & Deploy tab of System Preferences to configure the approval mode:

1. Select the **Manual Approval** mode if you want an approver to review and approve the changes before they are deployed.

By default, **Auto Approval** mode is selected. Use this mode if you want to deploy the configuration changes without a prior approval.

2. If you select the Manual Approval mode, add one or more approvers' e-mail addresses to notify the approvers every time a change request is submitted.

3. Specify the rollback limit, which is the number of change requests that can be rolled back.

The default value is 50. You can roll back a maximum of 1000 change requests.

4. Specify the time after which a change request elapses after the time it was created.

The minimum and maximum number of days that you can specify after which a change request elapses is 1 day and 365 days respectively. The change requests are highlighted in the following colors that indicate their overdue status.

- Red color—Indicates that the change request is in overdue status.
- Orange color—Indicates that the change request is due in less than 2 days.
- Green color—Indicates that the change request is not yet due.

5. Click **OK** to save the changes.

BEST PRACTICE: Configuring the approval mode must be a one-time operation. Do not change the approval mode frequently.

To change the approval mode from Auto Approval to Manual Approval, you must either deploy or discard the device configuration changes. You are unable to change the approval mode to Manual Approval, or from Manual Approval to Auto Approval if local changes are in pending deployment state. The message: **Do you want to retain the Change Request history?** is displayed when you change the approval mode. If you choose to retain the change request history, all the existing change requests are retained by the system. Hence, even if you switch to the Auto Approval mode, you can view the change requests that were created in Manual Approval mode.

NOTE: While configuring the Manual Approval mode, you can specify any number of approvers. If you specify more than one approver while configuring the Manual Approval mode, after any approver accepts or rejects a proposed change, the change request is not listed for the other approvers and they cannot approve or reject the same change request.

Setting up Auto-resynchronization Preferences

If you enable auto-resynchronization in Network Director, any configuration changes made on the physical device, including out-of-band CLI commits and change-request updates, automatically trigger resynchronization on the device.

To set up auto-resynchronization:

1. Select the **Config & Deploy** tab in the Preferences window.
2. Select the option **Purge unassigned system profiles after resynchronizing configuration**, which removes unassigned profiles generated by Network Director after resynchronization or deletion of a device.

NOTE: While upgrading Network Director, the profiles that are in unassigned state are not removed even if you select this option.

3. Specify the time interval in **Auto Resync TriggerWait Interval(sec)**. Network Director waits for this time interval before triggering auto-resynchronization.

The default time interval is 120 seconds.

4. Click **OK**.

Retaining Network Director Reports

By default, Network Director retains reports for 30 days. However, Network Administrators can change the retention period within the range 0 through 365 days. To change the setting, move the slider right or left on the Report tab of Preferences to the new setting. Click **OK** to save the setting.

Specifying Wireless Preferences

To manage Aruba wireless devices in your network, specify the URL to launch the Aruba's wireless management platform, the Aruba Airwave application.

Sites with Aruba Airwave application licenses can launch the Airwave application from within Network Director by supplying the Airwave application URL. For information about setting up preferences to launch the Aruba Airwave application, see ["Linking to the Aruba Airwave Application" on page 1553](#).

To manage Juniper wireless devices in your network, specify the URL to launch Juniper's wireless management platform, RingMaster.

Sites with RingMaster licenses can launch the RingMaster application from within Network Director by supplying the RingMaster URL. After typing the URL on the Wireless tab of Preferences, you can click **Launch RingMaster** in Build mode to load RingMaster into the main page. To enable launching RingMaster from within Network Director, simply type in the URL and click **OK** to save the setting.

Changing Monitor Mode Settings

IN THIS SECTION

- [Disabling Data Collection for Monitors | 114](#)
- [Changing the Polling Interval | 116](#)
- [Enabling and Disabling Collection for Managed Devices | 117](#)
- [Specifying Database History Retention | 117](#)
- [Specifying the Data Learning Engine \(DLE\) Settings | 118](#)

The Monitoring tab of Preferences has three tabs under it. These are:

- **Monitoring Settings**—Enables you to change the default polling interval for data collection for Monitor mode monitors. You can also disable or reenable the internal processes used for data collection on this sub-tab.
- **Client Session History**—Enables you to set the retention period for history records and the frequency that these records are checked for deletion.
- **Data Learning Engine Settings**—Enables you to specify the IP address of the Data Learning Engine (DLE) server or servers, which is a component of Cloud Analytics Engine that supports the flow path analysis and high-frequency statistics features in Network Director.
- **Device Settings**—Allows you to enable or disable data collection for one or more devices.

This section describes:

Disabling Data Collection for Monitors

Network Director internally gathers data for monitors by using a set of data collection processes. You can disable these data collectors if they do not pertain to your installation. For example, if you do not use Virtual Chassis, you can disable the data collection processes used for Virtual Chassis.

The data collection processes are divided into the following categories:

- Client
- Equipment
- FM
- RF
- Traffic
- Virtual

One data collector can be used by multiple monitors. Likewise, some monitors can be supported by multiple data collectors. These data collectors are enabled by default. To ensure proper data collection, if you enable the equipment data collectors, you must also enable the traffic collectors..

To disable or reenable a data collector:

1. Determine which monitors are used by the data collectors. Use [Table 9](#) to determine the relationship between the data collectors and the monitors.

Table 9: Monitor Mapping for Data Collectors

Monitor	Data Collector	Category
802.11 Packet Error	RFMonitorRadioStatsCollector	RF
AP Interference Source	RFMonitorIntSourcesCollector	RF
AP Status	EquipmentMonitorAPCollector	Equipment

Table 9: Monitor Mapping for Data Collectors (continued)

Monitor	Data Collector	Category
Current Sessions	Client Monitor Collector and SessionCountCollector	Client
Error Trend	PortTrafficMonitorCollector	Traffic
Logical Interfaces	EquipmentMonitorDeviceStatusCollector	Equipment
Find End Point	EquipmentMonitorEndPointCollector	Equipment
Percentage of Packets Retransmitted	RFMonitorRadioStatsCollector	RF
Port Status (physical)	EquipmentMonitorDeviceStatusCollector	Equipment
Radio Status	EquipmentMonitorAPCollector	Equipment
RF Neighborhood	RFMonitorRadioNeighborCollector	RF
RF Throughput or Packet Retransmitted	RFMonitorRadioStatsCollector	RF
Resource Utilization	EquipmentMonitorDeviceStatusCollector	Equipment
Session Trend	ClientMonitorCollector and SessionCountCollector	Client
Switch Status	EquipmentMonitorDeviceStatusCollector	Equipment
Traffic Trend	PortTrafficMonitorCollector	Traffic
Top Sessions by MAC Address	ClientMonitorCollector	Client
Top Users	ClientMonitorCollector	Client
Unicast vs Broadcast/Multicast	PortTrafficMonitorCollector	Traffic
Unicast vs Broadcast/Multicast Trend	PortTrafficMonitorCollector	Traffic
Virtual Chassis Topology	EquipmentMonitorVCStatsCollector and EquipmentMonitorDeviceStatusCollector	Equipment

Table 9: Monitor Mapping for Data Collectors (continued)

Monitor	Data Collector	Category
Virtual Chassis Protocol	EquipmentMonitorVCStatsCollector and EquipmentMonitorDeviceStatusCollector	Equipment
Virtual Chassis Statistics	EquipmentMonitorVCStatsCollector and EEquipmentMonitorDeviceStatusCollector	Equipment

2. Clear the check box to disable the collector or select to enable the collector.
3. Click **Save** and **Close** to save the configuration and to close the window.

Changing the Polling Interval

The frequency at which data is collected is determined by the polling interval. [Table 10](#) shows the default polling intervals used by each data collector.

Table 10: Default Polling Intervals

Collector	Polling Interval
ClientMonitorCollector	10 minutes
SessionCountCollector	10 minutes
EquipmentMonitorVCStatsCollector	30 minutes
EquipmentMonitorAPCollector	10 minutes
EquipmentMonitorEndPointCollector	1440 minutes
EquipmentMonitorDeviceStatusCollector	10 minutes
FMAAlarmCountCollector	10 minutes
RFMonitorIntSourcesCollector	15 minutes
RFMonitorRadioNeighborCollector	15 minutes
RFMonitorRadioStatsCollector	10 minutes
PortTrafficMonitorCollector	10 minutes

Table 10: Default Polling Intervals (continued)

Collector	Polling Interval
VirtualHostPMCollector	10 minutes
VirtualNICStatsCollector	10 minutes
VirtualMachineStatsCollector	10 minutes
VirtualMachineWeeklyStatsCollector	30 minutes

To change the polling interval:

1. Select the polling interval for a data collector in the Monitor Settings table.
2. Type the new interval level in whole minutes. For example, do not specify 1.5 minutes. Recommended intervals are 5, 10, or 20 minutes.
3. Click **OK** and then **Yes** to verify the change to the configuration.

Enabling and Disabling Collection for Managed Devices

By default all the devices that are discovered and managed by Network Director are enabled for data collection. You can disable or re-enable data collectors across all categories for devices that are managed by Network Director from this tab.

To enable or disable data collectors for devices

1. Open the **Device Settings** sub-tab in the Monitor tab.

All the devices that are managed by Network Director is displayed in the Device Settings section. The last column of the device table indicates the status of data collection as Enabled or Disabled.
2. Select the devices for which you want to enable or disable data collection and do one of the following:
 - Click **Enable** to enable data collection for the selected devices.
 - Click **Disable** to disable data collection for the selected devices.

Specifying Database History Retention

To keep the database manageable, the system periodically checks the age of the records and retires those that have past an expiration date. By default, Network Director ages database records off at 90 days and runs a database cleanup every 6 hours.

Use the Client Session History sub-tab to change the default values:

1. Select from the lists new values.

- Age of history records (in days) from 1 to 365 days.
- Cleanup job frequency (in hours) from 1 through 24 hours.

2. Click **OK** to save the changes.

Specifying the Data Learning Engine (DLE) Settings

The Data Learning Engine (DLE) is the component of Cloud Analytics Engine that enables Network Director to collect and analyze high-frequency statistics data from devices and to perform flow path analysis. For Network Director to use Cloud Analytics Engine, you must specify on what server or servers the Data Learning Engine (DLE) is running.

Each DLE supports up to a specific number of Compute Agents (CAs) running on network devices. If you have more CAs in a network than a single DLE can support, you might require multiple DLEs.

Use the Data Learning Engine Settings sub-tab under the Monitoring tab to specify which Data Learning Engine (DLE) server or servers Network Director uses. You can also change the default ports used by the DLE.

To configure DLE in Network Director:

1. Log in to Network Director.
2. Select **Preferences** from the list next to the **System** button in the Network Director banner.
The Preferences page is displayed.
3. Select the **Monitoring** tab and then select **Data Learning Engine Settings**.
4. In the **DLE IP Address** field, enter the IP address of the DLE server.

NOTE: Before you configure DLE in Network Director, make sure that there are no errors in the monitor.log file. The log file is stored in the `/var/log/jboss/server/server1` directory.

5. If you want to change the ports used by the DLE, click **View/Edit DLE Ports** to edit the ports and then click **OK**.

NOTE: If you change the default DLE ports (8282, 8081, and 50006), you must ensure that the new ports are open between DLE and Junos Space Network Management Platform or Network Director.

You can use the `netstat -anp | grep port-number` command to verify that the new ports are open (in listening mode) between DLE and Junos Space Network Management Platform or Network Director.

Table 11 describes the default DLE ports.

Table 11: Default DLE Port Descriptions

Port	Description
Flow Analysis API Port	Used by the flow path analysis feature and network traffic analysis feature to communicate with the DLE. Default value is 8282.
HFS API Port	Used by the high-frequency statistics feature to communicate with the DLE. Default value is 8081.
HFS Control Channel Port	Used by the high-frequency statistic feature for communication about threshold-related events with the DLE. Default value is 50006.

- Click **OK** to save the DLE settings.

The message **Preferences saved successfully** is displayed.

NOTE: After you configure the DLE settings, check whether the DLE connection state is **UP** in the DLE settings page.

- Click **Add Another** to add a new DLE server.

RELATED DOCUMENTATION

[Understanding Cloud Analytics Engine and Network Director | 83](#)

[Understanding Monitor Mode in Network Director | 1269](#)

Changing Alarm Settings

IN THIS SECTION

- [Configuring Global Alarm Notifications | 120](#)
- [Retaining Alarm History | 120](#)
- [Specifying Event History | 120](#)
- [Enabling Alarms | 121](#)
- [Changing the Severity of Individual Alarms | 134](#)
- [Configuring Threshold Alarms | 134](#)
- [Configuring Individual Alarm Notifications | 134](#)

Use the Fault tab to enable individual alarms, set the retention period for alarms, configure alarm notifications, configure threshold alarms, and to specify the number of events to keep for each alarm. The Fault tab has multiple sections, which you can expand and collapse by clicking the arrow next to the section title:

- Global Settings, for configuring Faults settings such as global alarm notifications and alarm data retention.
- Individual Alarms and Threshold Settings, for configuring settings for individual alarms and threshold alarms.

This section describes the following tasks that you can perform by using the Fault tab:

Configuring Global Alarm Notifications

You can configure global e-mail notifications to be sent when any alarm with notifications enabled is generated. To configure global e-mail notifications, enter the e-mail addresses to receive global alarm notifications in the Alarm Notifications Destinations field in the Global Settings section. Separate addresses with a comma (,). For information about enabling notification for an alarm, see [“Configuring Individual Alarm Notifications” on page 134](#).

Retaining Alarm History

Use the **No. of days to keep Alarm** field in the Global Settings section to specify the number of days to keep alarm history. The default retention time is 120 days; but you can specify a period of 7 through 1000 days. Specifying a longer retention time consumes more database resources. To change the alarm retention duration, type a new value and click **OK** and **Yes** to confirm the change.

Specifying Event History

Use the **Events/Alarm** field in the Global Settings section to specify the number of event entries that are kept in the alarm history. The default setting for events is 20. To change the setting, type a new value and click **OK** and **Yes** to confirm the change.

Enabling Alarms

Ensure all devices are configured to send traps to Network Director. This task is performed for the devices in Deploy mode through Set SNMP Trap Configuration.

Use the Individual Alarms and Threshold Settings section to disable and re-enable individual alarms or all alarms. Alarms appear on both tabs in the section: Alarm Settings and Threshold Settings. Fault alarms are preconfigured and initially enabled. To enable or disable alarms:

1. (Optional) Sort the alarms. By default, the list of alarms is sorted alphabetically within each category. You can also sort by description or alarm severity within a category by clicking a column heading.
2. Review the alarms and either select the check box in the heading to select all of the alarms or select the check box for the individual alarms you want to enable. For a full description of each of the alarms, see [Table 12](#).
3. Click **OK** and **Yes** to confirm the alarm change.

Table 12: Alarm Descriptions

Alarm Name	Description	Device Type
<i>AP and Radio (AP/Radio)</i>		
AP License Limit Exceeded	Generated when the number of wireless LAN access points (WLAs) exceed the number of licenses configured on a wireless switch. The trap occurs when a wireless switch receives a packet from an inactive access point. The switch is unable to attach the access point without exceeding the maximum (licensed) number of active access points.	Wireless LAN controller
AP Manager Changed Alarm	Generated when the access point's secondary link becomes the primary link.	Wireless LAN controller
AP Status Alarm	Generated when an access point changes state.	Wireless LAN controller
AP Tunnel Limit Exceeded	Generated when the number of tunnels on an access point exceeds the maximum number supported. This alarm is generated by the access point's Primary Access Manager (PAM) when the access point rejects a tunnel creation request because it has already created the maximum number of tunnels it can support.	Wireless LAN controller

Table 12: Alarm Descriptions (continued)

Alarm Name	Description	Device Type
M2U Conversion	Generated when multicast to unicast conversion is enabled on the access point, but cannot be performed.	Wireless LAN controllers
Radio Channel Changed	Generated when auto-tune changes a radio's channel.	Wireless LAN controller
Radio power changed	Generated when auto-tune changes a radio's power level.	Wireless LAN controller
Radio Status Alarm	Generated when a radio changes state. It also contains aggregate information about the access point in operational state, its security level and service availability.	Wireless LAN controller
WLC Tunnel Limit Exceeded	Generated when the wireless switch rejects a tunnel creation request because it has reached the maximum number of tunnels supported. When the trap event <code>trpzWsTunnelLimitType</code> equals the <code>platform-tunnel-limit</code> , the wireless switch has reached the maximum tunnel capacity. The actual tunnel limit varies by platform. When the trap <code>trpzWsTunnelLimitType</code> equals <code>ap-ws-tunnel-limit</code> , the wireless switch has reached the access point-to-switch tunnel's limit. The value of that limit depends on the current situation of the wireless switch (mobility domain, network domain, network and resiliency status).	Wireless LAN controller
<i>BFD</i>		
<code>BfdSessionDetectionTimeAlarm</code>	Generated when the threshold value for detection time is set and the BFD session detection-time adapts to a value greater than the threshold.	EX Series Switch
<code>BfdSessionTxAlarm</code>	Generated when the threshold value for transmit interval (in microseconds) is exceeded.	EX Series Switch
<i>BGP</i>		

Table 12: Alarm Descriptions (continued)

Alarm Name	Description	Device Type
BgpM2BackwardTransitionAlarm	Generated when the BGP FSM moves from a higher-numbered state to a lower-numbered state.	EX Series Switch
BgpM2EstablishedAlarm	Generated when the BGP Finite State Machine (FSM) enters the ESTABLISHED state.	EX Series Switch
<i>Chassis</i>		
FanFailureAlarm	Generated when the specified cooling fan or impeller has failed (is not spinning).	EX Series Switch
FEBSwitchoverAlarm	Generated when the Forwarding Engine Board (FEB) has switched over.	EX Series Switch
FRUCheckAlarm	Generated when the device has detected that a field-replaceable unit (FRU) has some operational errors and has gone into check state.	EX Series Switch
FRUFailedAlarm	Generated when a FRU has failed.	EX Series Switch
FRUInsertionAlarm	Generated when the system detects that the specified FRU is inserted into the chassis.	EX Series Switch
FRUOfflineAlarm	Generated when the specified FRU goes offline.	EX Series Switch
FRUOnlineAlarm	Generated when the specified FRU goes online.	EX Series Switch
FRUPowerOffAlarm	Generated when the specified FRU is powered off.	EX Series Switch
FRUPowerOnAlarm	Generated when the specified FRU is powered on.	EX Series Switch
FRURemovalAlarm	Generated when the system detects that the specified FRU was removed from the chassis.	EX Series Switch

Table 12: Alarm Descriptions (continued)

Alarm Name	Description	Device Type
HardDiskFailedAlarm	Generated when the hard disk for the specified routing engine has failed.	EX Series Switch
HardDiskMissingAlarm	Generated when the hard disk in the specified routing engine is missing from the boot device list.	EX Series Switch
PowerSupplyFailureAlarm	Generated when the specified power supply has failed (bad DC output).	EX Series Switch
RedundancySwitchOverAlarm	Generated when a graceful Routing Engine switchover (GRES) occurs on a switch with dual Routing Engines or on a Virtual Chassis.	EX Series Switch
TemperatureAlarm	Generated when the device has over heated.	EX Series Switch
<i>Client and User Session (ClientAndUserSession)</i>		
Client Association Failure	Generated when a client is unable to associate with an access point.	Wireless LAN controller
Client Authentication Failure	Generated when a client is unable to authenticate.	Wireless LAN controller
Client Authorization Failure	Generated when a client fails authorization.	Wireless LAN controller
Client Authorization Succeeded	Generated when a client authorizes.	Wireless LAN controller
Client Cleared	Generated when a client session is cleared.	Wireless LAN controller
Client Connectivity	Generated when a client session connects.	Wireless LAN controller
Client DeAssociated	Generated when a client de-association occurs.	Wireless LAN controller
Client DeAuthenticated	Generated when a client de-authenticates.	Wireless LAN controller

Table 12: Alarm Descriptions (continued)

Alarm Name	Description	Device Type
Client Disconnected	Generated when a client session disconnects administratively.	Wireless LAN controller
Client dot1x Failure	Generated when a client fails 802.1x.	Wireless LAN controller
Client Dynamic Authorization Changed	Generated when the authorization attributes for a user are dynamically changed by a authorized dynamic authorization client.	Wireless LAN controller
Client IP Address Changed	Generated when a client's IP address changes, normally when the client first connects to the network.	Wireless LAN controller
Client Roamed	Generated when a client roams from one location to another.	Wireless LAN controller
Dynamic Authorization Client Alarm	Generated when the authorization attributes for a user are dynamically changed by an authorized dynamic authorization client.	Wireless LAN controller
<i>Cluster/Modo</i>		
Cluster Sync Failure	Generated when the cluster configuration failed to apply.	Wireless LAN controller
Mobility Domain Failback	Generated when the mobility domain fails back to the primary seed.	Wireless LAN controller
Mobility Domain Failover	Generated when the mobility domain fails back to the secondary seed.	Wireless LAN controller
Mobility Domain Join	Generated when a member joins the mobility domain.	Wireless LAN controller
Mobility Domain Resiliency Status	Generated when a mobility domain seed changes resilient capacity.	Wireless LAN controller
Mobility Domain Timeout	Generated when a mobility domain member times out.	Wireless LAN controller

Table 12: Alarm Descriptions (continued)

Alarm Name	Description	Device Type
<i>Configuration (Configuration)</i>		
CmCfgChangeAlarm	Generated when the jnxCMCfgChgEventTable records a configuration management event.	EX Series Switch and wireless LAN controller
CMRescueChangeAlarm	Generated when a change is made to the rescue configuration.	EX Series Switch and wireless LAN controller
<i>Core and controllers (Controllers)</i>		
Device alarm	Generated when the device status changes (up to down or down to up).	EX Series Switch and wireless LAN controller
<i>CoS</i>		
CoSAlmostOutOfDedicatedQueuesAlarm	Generated when only 10% of CoS queues are available.	EX Series Switch
CoSOutOfDedicatedQueuesAlarm	Generated when there are no more available dedicated CoS queues.	EX Series Switch
<i>DHCP</i>		
JdhcpLocalServerDupClientAlarm	Generated when a DHCP client is detected changing interfaces.	EX Series Switch
JdhcpLocalServerIfLimitExceededAlarm	Generated when the client limit is reached on an interface.	EX Series Switch
Jdhcpv6LocalServerLimitExceededAlarm	Generated when the client limit is reached on an interface for DHCPv6.	EX Series Switch
<i>DOM</i>		
DomAlertSetAlarm	Generated when an interface detects Digital Optical Monitor (DOM) alarm conditions.	EX Series Switch
<i>Flow Collection (FlowCollection)</i>		

Table 12: Alarm Descriptions (continued)

Alarm Name	Description	Device Type
CollFlowOverloadAlarm	Generated when a collector PIC detects a hard or soft flow overload.	EX Series Switch
CollFtpSwitchOverAlarm	Generated when an FTP server switchover occurs.	EX Series Switch
CollMemoryUnavailableAlarm	Generated when a PIC is out of memory or the memory is unavailable.	EX Series Switch
CollUnavailableDestAlarm	Generated when a file transfer destination is unavailable.	EX Series Switch
CollUnsuccessfulTransferAlarm	Generated when a collector file is unable to transfer because the destination is unavailable.	EX Series Switch
<i>General</i>		
Authentication Failure Alarm	Generated when a protocol message is received that is not properly authenticated.	EX Series Switch and wireless LAN controller
Cold Start Alarm	Generated when a device is re-initializing and its configuration might have changed.	EX Series Switch and wireless LAN controller
Link Down Alarm	Generated when a link is down. The trap is generated when the ifOperStatus object for a communication link is about to enter the down state from another state other than notPresent. This other state is indicated by the included value of ifOperStatus.	EX Series Switch and wireless LAN controller
Link Up Alarm	Generated when a link comes up that was previously in the down state. The trap is generated when the ifOperStatus object for a communication link left the down state and transitioned into another state other than notPresent state. This other state is indicated by the included value of ifOperStatus.	EX Series Switch and wireless LAN controller

Table 12: Alarm Descriptions (continued)

Alarm Name	Description	Device Type
Warm Start Alarm	Generated when a device is re-initializing and its configuration has not changed.	EX Series Switch and wireless LAN controller
<i>Generic (GenericEvent)</i>		
GenericEventTrapAlarm	Generated by an Op script or event policies. This notification can include one or more attribute-value pairs. The pairs are identified by the <code>jnxEventAvAttribute</code> and <code>jnxEventAvValue</code> objects.	EX Series Switch
<i>L2ALD</i>		
L2aldGlobalMacLimitAlarm	Generated when the MAC limit is reached for the entire system. This trap is sent only once, when the limit is reached.	EX Series Switch
L2aldInterfaceMacLimitAlarm	Generated when the given interface reaches the MAC limit (<code>jnxl2aldInterfaceMacLimit</code>).	EX Series Switch
L2aldRoutingInstMacLimitAlarm	Generated when the MAC limit is reached for a given routing instance (<code>jnxl2aldRoutingInst</code>).	EX Series Switch
<i>L2CP</i>		
LacpTimeOutAlarm	Generated when LACP has timed out.	EX Series Switch
PortBpduErrorStatusChangeTrapAlarm	Generated when the port's BPDU error state (no-error or detected) changes.	EX Series Switch
PortLoopProtectStateChangeTrapAlarm	Generated when the port's loop-protect state (no-error or loop-prevented) changes.	EX Series Switch
PortRootProtectStateChangeTrapAlarm	Generated when the port's root-protect state (no-error or root-prevented) changes.	EX Series Switch
<i>MAC Forwarding Database (MACFDB)</i>		

Table 12: Alarm Descriptions (continued)

Alarm Name	Description	Device Type
MacChangedNotificationAlarm	Generated when MAC addresses of the monitored devices are learned or removed from the forwarding database (FDB).	EX Series Switch and wireless LAN controller
<i>Misc.</i>		
Counter Measures Alarm	Generated when counter measures are started against a rogue device.	Wireless LAN controller
Device Configuration Saved	Generated when the running configuration of the switch is written to the configuration file.	Wireless LAN controller
Multimedia Call Failure	Generated when a multimedia call fails.	Wireless LAN controller
<i>PoE (Power over Ethernet)</i>		
PoE Port ON-OFF Alarm	Generated when the PoE power is turned on or off.	EX Series Switch
PoE Power Usage High	Generated when Power over Ethernet (PoE) used is below or above the defined threshold.	EX Series Switch
<i>Passive Monitoring (PassiveMonitoring)</i>		
PMonOverloadSetAlarm	Generated when an overload condition is detected on a Passive Monitoring Interface.	EX Series Switch
<i>Ping</i>		
PingEgressJitterThresholdExceededAlarm	Generated when egress time jitter ($\text{jnxPingMaxEgressUs} - \text{jnxPingResultsMinEgressUs}$) exceeds the configured threshold ($\text{jnxPingCtlEgressJitterThreshold}$) causing the <code>egressJitterThreshold</code> bit to be set.	EX Series Switch and wireless LAN controller
PingEgressStdDevThresholdExceededAlarm	Generated when the standard deviation of the egress time ($\text{jnxPingResultsStddevEgressUs}$) exceeds the configured threshold ($\text{jnxPingCtlEgressTimeThreshold}$) and causes the <code>egress</code> bit to be set.	EX Series Switch and wireless LAN controller

Table 12: Alarm Descriptions (continued)

Alarm Name	Description	Device Type
PingEgressThresholdExceededAlarm	Generated when the egress time (jnxPingResultsStdevEgressUs) exceeds the configured threshold (jnxPingCtlEgressTimeThreshold) and the egress threshold bit is set in jnxPingCtlTrapGeneration.	EX Series Switch and wireless LAN controller
PingIngressJitterThresholdExceededAlarm	Generated when ingress time jitter (jnxPingResultsMaxIngressUs minus jnxPingResultsMinIngressUs) exceeds the configured threshold (jnxPingCtlIngressJitterThreshold) and the ingressJitterThreshold bit is set in jnxPingCtlTrapGeneration.	EX Series Switch and wireless LAN controller
PingIngressStddevThresholdExceededAlarm	Generated when the standard deviation of the ingress time (jnxPingResultsStdDevIngressUs) exceeds the configured threshold (jnxPingCtlIngressStddevThreshold) and the ingress StdDevThreshold bit is set in jnxPingCtlTrapGeneration.	EX Series Switch and wireless LAN controller
PingIngressThresholdExceededAlarm	Generated when the ingress time jitter (jnxPingResultsIngressUs) exceeds the configured threshold (jnxPingCtlIngressTimeThreshold) and the ingress threshold bit (jnxPingIngressThresholdExceeded) is set in jnxPingCtlTrapGeneration.	EX Series Switch and wireless LAN controller
PingRttJitterThresholdExceededAlarm	Generated when the round trip time jitter (jnxPingResultsMaxRttUs minus jnxPingResultsMinRttUs) exceeds the configured threshold (jnxPingCtlRttJitterThreshold) and the rttJitterThreshold bit is set in jnxPingCtlTrapGeneration.	EX Series Switch and wireless LAN controller

Table 12: Alarm Descriptions (continued)

Alarm Name	Description	Device Type
PingRttStdDevThresholdExceededAlarm	Generated when the standard deviation of the round trip time (jnxPingResultsStdDevRttUs) exceeds the configured threshold (jnxPingCtlRTTStdDev) and the rttStdDevThreshold bit is set in jnxPingCtlTrapGeneration.	EX Series Switch and wireless LAN controller
PingRttThresholdExceededAlarm	Generated when the round trip time (jnxPingCtlRttThreshold) exceeds the configured threshold (jnxPingCtlRttThreshold) and the rttThreshold bit is set in jnxPingCtlTrapGeneration.	EX Series Switch and wireless LAN controller
<i>RF Detect (RFDetect)</i>		
Adhoc user detected	Generated when RF detection sweep finds an ad hoc user or if a previously found ad hoc user disappears.	Wireless LAN controller
Client Blacklisted	Generated when an association, re-association, or de-association request is detected from a blacklisted transmitter.	Wireless LAN controller
DoS Attack Detected	Generated when RF detection finds a denial-of-service (DoS) attack occurring.	Wireless LAN controller
DoS Port Detected	Generated when RF detection finds a denial of service (DoS) attack occurring. This trap collects port and access point information instead of information about the listener.	Wireless LAN controller
RF Interference Detected	Generated when a new noise source appears. A given combination of noise source ID, listener, and channel triggers this trap. It is normally not triggered more than once every 15 minutes.	Wireless LAN controller
RF Detect Classification Changed	Generated when the RF detection classification rules change.	Wireless LAN controller
Rogue Device Detected	Generated when RF detection finds a rogue device.	Wireless LAN controller

Table 12: Alarm Descriptions (continued)

Alarm Name	Description	Device Type
Rogue Wired WLA Client Detected	Generated when a client is detected that connected through a rogue access point that is attached to a wired port.	Wireless LAN controller
Rogue WLA Client Detected	Generated when RF detection finds a suspect device.	Wireless LAN controller
Rogue WLA Interference Detected	Generated when RF detection finds an interfering rogue access point.	Wireless LAN controller
Spoofed MAC Detected	Generated when RF detection finds an access point using the MAC of the listener.	Wireless LAN controller
Spoofed SSID Detected	Generated when RF detection finds an access point using the SSID of the listener, and the access point is not in the mobility domain.	Wireless LAN controller
Suspected Device Detected	Generated when RF detection finds a suspect device.	Wireless LAN controller
Unauthorized AP Detected	Generated when RF detection discovers an unauthorized access point being used.	Wireless LAN controller
Unauthorized OUI Detected	Generated when RF detection finds an unauthorized OUI being used.	Wireless LAN controller
Unauthorized SSID Detected	Generated when RF detection finds an unauthorized SSID being used.	Wireless LAN controller
<i>RMon</i>		
RmonAlarmGetFailureAlarm	Generated when a GET request for an alarm variable returns an error. The specific error is identified by a varbind in jnxRmonAlarmGetFailReason.	EX Series Switch
<i>SONET</i>		
SonetAlarmSetAlarm	Generated when there is a notification of a recently set SONET or SDH alarm on an interface.	EX Series Switch

Table 12: Alarm Descriptions (continued)

Alarm Name	Description	Device Type
<i>SONET APS (SONETAPS)</i>		
APSEventChannelMismatchAlarm	Generated when the value of an instance of apsStatusChannelMismatches increments.	EX Series Switch
APSEventFEPLFAlarm	Generated when the value of an instance of apsEventFEPLFs increments.	EX Series Switch
APSEventModeMismatchAlarm	Generated when the value of an instance of apsEventModeMismatch increments.	EX Series Switch
APSEventPSBFAlarm	Generated when the value of an instance of apsStatusPSBFs increments.	EX Series Switch
APSEventSwitchoverAlarm	Generated when the value of an instance of apChanStatusSwitchover increments.	EX Series Switch
<i>Virtual Chassis (VirtualChassis)</i>		
VccpMemberAlarm	Generated when a member has completed transition from the down state to another state.	EX Series Switch
VccpPortAlarm	Generated when one of the member's communication links has completed transition from the down state to another state.	EX Series Switch
<i>VNetwork</i>		
HostConnectivityLostAlarm	Generated when all the uplink ports of a virtual switch residing in a host loses network connectivity.	Host
HostNetworkRedundancyLostAlarm	Generated when some uplink ports of a virtual switch residing in a host loses network connectivity. It indicates that there are one or more ports that still has network connectivity.	Host
VNetworkConnectivityLostAlarm	Generated when Network Director loses network connectivity with the vCenter server.	Virtual Network

Changing the Severity of Individual Alarms

You can change the severity of the alarms to match your corporate procedures and guidelines. For example, at your company a DoS attack might be considered a critical alarm, while Network Director has a default severity for DoS attacks as a major alarm. Alarms appear on both tabs in the Individual Alarms and Threshold Settings section: Alarm Settings and Threshold Settings.

To change the severity of an alarm:

1. Select the current severity in the **Severity** column. A list of the severity levels appear.
2. Select the new severity level for the alarm.
3. Click **OK** and **Yes** to confirm the change to the severity setting.

To configure alarm notifications, see [“Configuring Individual Alarm Notifications” on page 134](#).

Configuring Threshold Alarms

Threshold alarms are alarms that are generated when a monitored value crosses the configured threshold. They provide enhanced visibility into potential issues on the network. You configure and manage threshold alarms the same way as other alarms. You also have the option of setting the threshold level of individual threshold alarms.

To edit the threshold of threshold alarms:

1. Select the **Threshold Settings** tab in the Individual Alarms and Threshold settings section of the Fault tab.
2. Click **Edit Settings** in the Threshold Settings column of the alarm threshold you want to edit.
3. Set the threshold in the window that opens.
4. Click **Save** to save the new threshold.

To configure alarm notifications, see [“Configuring Individual Alarm Notifications” on page 134](#).

Configuring Individual Alarm Notifications

You can configure e-mail notifications to be sent when an individual alarm is generated. When you enable notification for an alarm, the notifications are sent to the e-mail addresses configured for the alarm and the addresses configured for global alarm notifications. Alarms appear on both tabs in the Individual Alarms and Threshold Settings section: Alarm Settings and Threshold Settings.

To configure e-mail notification for an alarm name:

1. Select the check box in the alarm's Notification column.
If you later want to disable notification for the alarm, clear the check box.
2. Click **Edit Notification** in the Notification column. The Alarm Notification Details window opens.
3. Type one or more e-mail addresses in the **Notification Email Addresses** box. Separate addresses with a comma (,).
You can later edit the addresses to send notifications to different addresses.
4. (Optional) Type a comment in the Comments box. This comment is included in the e-mail notification message.
5. Click **Save**.

RELATED DOCUMENTATION

[Understanding the Fault Mode Tasks Pane | 1452](#)

[Current Active Alarms Monitor | 1464](#)

[Alarms by State Monitor | 1467](#)

[Alarms by Severity Monitor | 1466](#)

[Alarms by Category Monitor | 1466](#)

[Network Director Documentation home page](#)

Modifying Data Center Synchronization Interval Using the Virtualization Tab

You can configure and manage data centers by using the Data Center View in Network Director. Network Director synchronizes data from the cloud infrastructure that is part of your data center every 24 hours. However, depending on the size of your data center network, you can modify this interval from the Virtualization tab in the Preferences page. Modify the interval by changing the **Periodic Cloud Infrastructure Synchronization Interval (hours)** field.

For data centers that use OpenStack as the cloud infrastructure provider, Network Director displays the links between hypervisors and switches only in the Data Center View. If you want to view the link details using the View Virtual Network Connectivity task, then you must select the **Synchronize links from Topology to Virtualization application for OpenStack** check box.

NOTE: Selecting the **Synchronize links from Topology to Virtualization application for OpenStack** check box might impact the performance of Network Director, while the synchronization is in process.

RELATED DOCUMENTATION

[Understanding the Network Director User Interface | 85](#)

[Enabling SNMP Categories and Setting Trap Destinations | 1208](#)

[Understanding Fault Mode in Network Director | 1448](#)

[Network Director Documentation home page](#)

Getting Started with Network Director

IN THIS CHAPTER

- [Getting Started with Junos Space Network Director | 137](#)

Getting Started with Junos Space Network Director

IN THIS SECTION

- [Building Your Network | 137](#)
- [Creating Profiles in Network Director | 138](#)
- [Managing Software Images using Network Director | 139](#)
- [Configuring Approval Modes for Device Configurations | 139](#)
- [Resynchronizing Device Configuration | 140](#)
- [Creating the Baseline Configuration | 140](#)
- [Monitoring Your Network | 140](#)
- [Setting up Network Traffic Analysis and Analyzing the Traffic | 141](#)
- [Managing Network Faults and Notifications | 141](#)
- [Generating Network Reports | 141](#)

This section describes a series of steps that you must perform after installing Network Director to manage and troubleshoot your network.

Building Your Network

The first step after you install and log in to Network Director is to build your network. Even with large networks, Network Director has made this step relatively easy and straightforward. The steps that you need to perform depend on whether your network contains legacy devices, or new devices, or a combination of both.

You add legacy devices, which already have some configurations, to Network Director by using a process called *device discovery*. Once such a device is successfully discovered, Network Director reads the device configurations and replicates these configurations in the form of profiles in Network Director. You can use device discovery to add Juniper Networks switches and Wireless LAN Controllers (WLCs) to Network Director. For more information on device discovery, see [“Discovering Devices in a Physical Network” on page 204](#) and [“Understanding the Device Discovery Process” on page 216](#).

With new devices or devices that are set to factory-default configuration, you can use the *zero touch provisioning (ZTP)* feature to provision the device. ZTP enables you to auto-discover, auto-upgrade, and load the requisite default configuration on Juniper Networks switches in your network automatically—without manual intervention. When you physically connect a switch that has the factory-default configuration to a network and boot the switch, the switch attempts to upgrade Junos OS automatically and autoinstall a configuration file from the network. For more information, see [“Configuring and Monitoring Zero Touch Provisioning” on page 1261](#).

Creating Profiles in Network Director

Profiles in Network Director are a group of feature-specific configurations that you can assign to devices. For example, you can create a CoS profile that combines all the supported class-of-service configurations for a particular device can family, and assign it to a port on a device.

- You can create a new profile for an interface or device by defining the custom configuration. You can use the Tasks pane in Build mode to manually create profiles. For more details, see .
- Network Director automatically creates profiles based on the configuration information read by the brownfield process. This is applicable when a device with supported configuration is discovered in Network Director. For more information, see *Brownfield Deployment in Network Director*.
- Network Director automatically creates profiles when a supported configuration of a device that is already discovered and managed by Network Director is modified outside Network Director (also known as out-of-band configuration changes). For more information, see [“Understanding Resynchronization of Device Configuration” on page 1214](#).

Following are some advantages of using profiles:

- Bulk provisioning—You can combine a group of configurations as a profile and apply it to one or more ports or devices in one go, thereby saving a lot of time and effort. Profiles ensure that the configurations are error-free as most configuration value ranges are set in the profile workflow. Network Director prompts the user if there are any errors. You must fix the errors before you can create a profile.
- Editing—For profiles that are already deployed on devices, if you want to make changes to the configuration values, you can modify the configuration values in the profile and redeploy the profile. Network Director updates the new configuration value on each device where the profile is deployed.
- Cloning—If you already have a set of profiles defined for your network and want to apply a different configuration for a set of devices or ports in your network, you use the clone feature. The clone feature

enables you to make a copy of any profile and make the necessary modifications. You can then apply these to devices and ports that require the different set of configuration.

For more information on profiles, see [“Understanding Network Configuration Profiles” on page 197](#).

Managing Software Images using Network Director

As a Network Administrator, you can store different versions of Junos OS software images in the Network Director image repository. You can then deploy these images on one or more managed devices manually or have the system deploy the images by using zero touch provisioning (ZTP).

For more information on managing and deploying software images, see [“Managing Software Images” on page 1235](#) and [“Deploying Software Images” on page 1238](#).

Configuring Approval Modes for Device Configurations

When you make configuration changes in Build mode, the changes are not deployed to devices automatically. You must manually deploy the changes to devices in Deploy mode. When you deploy configuration changes to a device, all pending configuration changes for that device are deployed. You can deploy the device configurations in the following two ways:

- **Auto Approval**—In this mode, the device configuration changes are approved automatically by the system and do not require explicit (manual) approval by a configuration approver before they can be deployed. This is the default approval mode.
- **Manual Approval**—In this mode, the device configuration changes must be explicitly approved by a configuration approver before the changes can be deployed to the device. An operator performs device configurations and creates a change request for that configuration and submits it for approval to one or more approvers. The approvers are notified by e-mail whenever a change request is created. If a configuration or a change to it is approved by an approver, then the operator is able to deploy it. If a configuration is rejected, the operator must make the necessary changes, resubmit the change request, and procure an approval before the configuration can be deployed.

NOTE: For manual approval, the **Network Director - Configuration Approver** role is available in Junos Space, which is specific to Network Director. A user with this role reviews device configurations and proposed changes to device configurations and can either approve or reject them.

For more information about deploying configuration to devices, see [“Deploying Configuration to Devices” on page 1180](#).

Resynchronizing Device Configuration

A network managed by Network Director has three repositories of information about the configuration of a network device—the configuration stored on the device itself, the device configuration record maintained by Junos Space, and the Build mode configuration maintained by Network Director.

When the configuration contained in all three repositories match, the device configuration state is shown as In Sync in Network Director. When the repositories do not match, the configuration state is shown as Out of Sync. A common cause for this state is out-of-band configuration changes—that is, configuration changes made to a device outside of Network Director.

When a device state is Out of Sync, you cannot deploy configuration changes on the device in Deploy mode. Use the Resynchronize Device Configuration task to resynchronize the three configuration repositories and change the device configuration state back to In Sync.

For more information about device resynchronization, see [“Understanding Resynchronization of Device Configuration” on page 1214](#).

Creating the Baseline Configuration

You can create a baseline of configuration and the Junos OS version of the devices on the Network Director server. By creating a baseline configuration file for a device you define a reference point to save the device configuration and its Junos OS version to a particular known state and later restore the configuration to that known state.

For more information about device resynchronization, see [“Creating and Managing Baseline of Device Configuration Files” on page 1230](#).

Monitoring Your Network

Network Director provides the visibility into your network status and performance by using the Monitor Mode.

Network Director monitors the devices it manages and maintains the information it collects from the devices in a database. You can view this data as easy-to-understand graphs and tables—known as monitoring widgets—to quickly visualize the state of your network, spot trends developing over time, and view important details.

For more information about the monitor mode, see [“Understanding Monitor Mode in Network Director” on page 1269](#).

You can also use the Dashboard widgets to monitor your network performance. For more information about the Dashboard widgets, see [“Understanding the Dashboard” on page 144](#).

Setting up Network Traffic Analysis and Analyzing the Traffic

The Network Traffic Analysis feature of Network Director monitors high-speed switched or routed networks. Once enabled, Network Director randomly samples network packets and sends the samples to a data learning engine (DLE) for analysis. Network traffic analysis uses packet-based sampling. Network Director samples one packet out of a specified number of packets from an interface enabled for network traffic analysis and sends the packet to the DLE. DLE uses this sampling information to create a picture of the network traffic, which includes the applications that contribute to the traffic, traffic statistics, and the top applications. You can enable network traffic analysis on all devices, except the wireless devices, that are managed by Network Director.

For more information about installing and configuring DLE, see [Installing and Configuring Data Learning Engine for Network Director](#).

Managing Network Faults and Notifications

In Fault mode, Network Director informs you of unexpected, significant events happening in your network. Examples of such events include link up or link down, power supply failure, client authentication failure, detection of an unauthorized access point, and so on.

Network Director receives information about events from its managed devices in the form of SNMP notifications. A single event can often generate multiple SNMP notifications. To simplify management of events, Network Director correlates these notifications, creating high-level alarms of different severity levels for the events. For example, a power supply failure might generate a number of notifications. Network Director correlates these notifications and raises a single power supply failure alarm for the device. Network Director also automatically clears an alarm if it receives notification from the device that the error condition has been resolved.

To tailor Network Director fault management to your organization's requirements, you can enable or disable the receipt of specific alarms and change the default severity level of alarms.

For more information about the fault mode in Network Director, see [“Understanding Fault Mode in Network Director” on page 1448](#).

Generating Network Reports

Use the Report mode in Network Director to create standardized reports from the monitoring and fault data collected by Network Director. An essential part of the network management life cycle, reporting provides administrators and management insight into the network for maintenance, troubleshooting, trend and capacity analysis, and provides records that can be archived for compliance requirements.

Network Director provides reports in PDF and HTML formats that use graphs and tables to clearly convey data. Reports are also available in CSV format for importing into spreadsheets.

For more information about managing reports in Network Director, see [“Managing Reports in Network Director”](#) on page 1475.

RELATED DOCUMENTATION

[Understanding Build Mode in Network Director | 184](#)

[Understanding Deploy Mode in Network Director | 1172](#)

[Understanding Monitor Mode in Network Director | 1269](#)

[Understanding Fault Mode in Network Director | 1448](#)

[Understanding Report Mode in Network Director | 1470](#)

2

PART

Working with the Dashboard

[About the Dashboard](#) | 144

[Using the Dashboard](#) | 145

[Dashboard Widget Reference](#) | 147

About the Dashboard

IN THIS CHAPTER

- [Understanding the Dashboard | 144](#)

Understanding the Dashboard

The Dashboard is a customizable page to view information about the network, and is the default page that opens when you log in. You select monitoring widgets to display on the Dashboard that show various information about the network. The Dashboard is a view. To open a different view, select a view from the Views list in the Network Director banner.

RELATED DOCUMENTATION

[Using Dashboard Widgets | 145](#)

[Network Director Documentation home page](#)

Using the Dashboard

IN THIS CHAPTER

- [Using Dashboard Widgets | 145](#)

Using Dashboard Widgets

The Dashboard is a customizable page for viewing information about the network. You select monitoring widgets to display on the Dashboard that show various information about the network. The Dashboard is the default view that opens when you log in. When a different view is selected, select **Dashboard View** from the Select View list in the Network Director banner to open the Dashboard.

To select what appears on the Dashboard:

1. To add a monitor to the Dashboard:
 - a. Select **Add Widgets**. Thumbnails of the available widgets appear.
 - b. To add a widget to the Dashboard, mouse over the widget's thumbnail, then click the **Add** button that appears on the widgets.
 - c. When you are finished adding widgets, click **Done**. The new widgets appear on the Home page.
2. To refresh a widget's data, click the **Refresh** button in its title bar.
3. To see additional information for a widget, click the **Maximize** button in the widget's title bar.
4. To remove a widget from the Dashboard, click the Close button (**X**) in its title bar.
5. To open online help for a widget, click the Help button (?) in its title bar.
6. To move a widget, click its title bar and drag it to the new location.

RELATED DOCUMENTATION

[Understanding the Dashboard | 144](#)

[Network Director Documentation home page](#)

Dashboard Widget Reference

IN THIS CHAPTER

- [Alarms Widget | 147](#)
- [Config Deployment Jobs Status Widget | 149](#)
- [Device & Port Latency Widget | 150](#)
- [Device & Port Utilization Widget | 151](#)
- [Equipment By Type Widget | 158](#)
- [Port Status - Physical Widget | 159](#)
- [Recent Flow Analysis Widget | 160](#)
- [Top Talker - Wired Devices Widget | 167](#)
- [Top Virtual Machines by Bandwidth Widget | 168](#)
- [Top vNetwork Hosts by Bandwidth Widget | 169](#)
- [Virtual Machines & Bare Metal Servers Widget | 169](#)
- [Top Overlay Networks Widget | 179](#)

Alarms Widget

IN THIS SECTION

- [Alarms Widget Summary | 148](#)
- [Alarms Widget Details | 148](#)

The Alarms widget provides summary and detailed information about network alarms.

This topic describes:

Alarms Widget Summary

The summary view of the Alarms widget displays summary information about network alarms and their location. The number of active alarms of each severity is shown in colored circles on the left side of the widget. The distribution of alarms by site is shown on a map. The alarms count for each site is shown as a pie chart. The color of each pie chart segment indicates severity level. The colored circles to the left of the map also serve as the legend for the color coding.

Mouse over a pie chart to see more information about the alarms for that site.

Alarms Widget Details

To open the Alarms widget details page, click the **Maximize** button in the widget's title bar. The Alarms widget details window displays detailed information active alarms. The top of the page contains a larger view of the widget. The bottom of the page contains a table of detailed information about active alarms. [Table 13](#) describes the columns in this table. Click an alarm severity level circle to filter the table to show only alarms of that severity. To close the details page, click the **Minimize** button in the title bar.

Table 13: Alarm Widget Details Table

Column	Description
Name	The alarm name.
ID	A system and sequentially-generated identification number.
Assigned To	If assigned to an individual, it shows the name of the person assigned; otherwise, it shows System to mark that the alarm is still unassigned.
Severity	The severity of the alarm. Severity levels are: <ul style="list-style-type: none"> • Critical—A critical condition exists; immediate action is necessary. • Major—A major error has occurred; escalate or notify as necessary. • Minor—A minor error has occurred; notify or monitor the condition. • Info—An informational message; no action is necessary.
Entity ID	The identification of the entity responsible for causing this alarm. The Entity ID is the key for correlation of events into an alarm. The Entity ID could be a MAC address of a radio or an IP address of the device.
Reporting Device IP	The IP address of the device or network element that generated the alarm. The SNMP agent is located at the source IP. In most cases, the source IP is the IP address of the switch or controller.
Reporting Device	The hostname of the reporting device.

Table 13: Alarm Widget Details Table (continued)

Column	Description
Creation Date	The date and time the alarm was first reported.
Last Updated	The date and time that the information for the alarm was last modified.
Updated By	Either the system or the last user who modified the alarm.
Acknowledged	Indicates if the alarm has been acknowledged.

RELATED DOCUMENTATION

[Understanding the Dashboard | 144](#)

[Using Dashboard Widgets | 145](#)

[Network Director Documentation home page](#)

Config Deployment Jobs Status Widget

IN THIS SECTION

- [Config Deployment Jobs Status Widget Summary | 149](#)
- [Config Deployment Jobs Status Widget Details | 150](#)

The Config Deployment Jobs Status widget provides summary and detailed information about the status of configuration deployment jobs.

This topic describes:

Config Deployment Jobs Status Widget Summary

The Config Deployment Jobs Status widget displays summary information about the status of configuration deployment jobs. The information appears in a table. The vertical axis lists the job statuses. The horizontal axis shows the times when job status data was collected. You can do the following tasks:

- Select a time period to view from the **Deployment Trend** list.
- Click the **Refresh** button to refresh the information displayed.

Config Deployment Jobs Status Widget Details

To open the Config Deployment Jobs Status widget details page, click the **Maximize** button in the widget's title bar. The Config Deployment Jobs Status widget details window displays detailed information about the status of configuration deployment jobs. The page shows the same summary information table as the widget. It also shows a table of detailed configuration job status information. To close the details page, click the **Minimize** button in the title bar.

RELATED DOCUMENTATION

[Understanding the Dashboard | 144](#)

[Using Dashboard Widgets | 145](#)

[Network Director Documentation home page](#)

Device & Port Latency Widget

The Device & Port Latency widget provides a graphical view of latency on devices. The heat map represents each device as a color-coded box. The color coding indicates the level of latency on a device. Cooler colors (for example, green) indicate lower latency, while hotter colors (for example, red) indicate higher latency.

The Device & Port Latency widget can show information only for devices that support Cloud Analytics Engine and that have the high-frequency traffic statistics feature enabled in Network Director. For information about the Cloud Analytics Engine, see [“Understanding Cloud Analytics Engine and Network Director” on page 83](#).

To use the Device & Port Latency Heat Map widget, you must do the following first:

- Configure the Data Learning Engine (DLE) settings under **Preferences > Monitoring > Data Learning Engine Settings**. The DLE is a component of Cloud Analytics Engine. For information on configuring the DLE settings, see [“Specifying the Data Learning Engine \(DLE\) Settings” on page 118](#).
- Enable high-frequency traffic statistics on the devices you want to monitor. For information, see [“Enabling High-Frequency Traffic Statistics Monitoring on Devices” on page 1202](#).

You can do the following with the heat map:

- Select the time period to view from the list in the title bar.
- Select how to organize the heat map by clicking the Settings icon and then selecting an option from the **Group Devices By** list. Each option opens a different view of the heat map, with device boxes grouped according to your selection.
- Select a filter for which devices to show by clicking the Settings icon and then selecting an option from the **Show** list.
- Drill down into the heat map's hierarchy by clicking one of the device container names (for example, a site or building). To move back up the hierarchy, click the navigation arrows above the heat map.
- Mouse over a device box to view detailed latency and congestion information about the device in a pop-up window.
- Click a device box to view detailed latency and congestion information about the device's ports. In this view, each port is represented by a box that is color-coded to show its level of latency and congestion. Mouse over a port box to see more information about the port. To move back up the hierarchy, click the navigation arrows above the heat map.
- Slide the circular controls along the bar under the heat map to filter the devices or ports shown in the heat map by degree of latency.

RELATED DOCUMENTATION

[Understanding the Dashboard | 144](#)

[Using Dashboard Widgets | 145](#)

Device & Port Utilization Widget

IN THIS SECTION




- [Using the Global Controls | 152](#)
- [Interacting with the Heat Maps | 152](#)
- [Viewing Active Flows on a Port | 153](#)
- [Flow Analysis Details Window | 154](#)
- [Viewing Traffic on a Device | 155](#)

The Device & Port Utilization Heatmap widget provides a graphical view of device port utilization percentage. The heat map represents each device as a color-coded box. The color coding indicates the overall level of port utilization on a device. Cooler colors (for example, green) indicate lower port utilization, while hotter colors (for example, red or dark red) indicate higher port utilization.

You can view the utilization level for each port on a device by clicking the box representing the device. A heat map is displayed that represents each port on the device as a color-coded box, with the color coding representing the level of port utilization.

Using the Global Controls

Use the controls in the upper right corner to make global changes to how the device and port heat maps are displayed. You can:

- Select the time period over which device utilization and port utilization are shown.
- Display information about the devices or the ports in either graphical heat map or tabular format by clicking either  (graphical) or  (tabular).
- Select how to organize the heat map by clicking the Settings icon (), and then selecting an option from the **Group Devices By** list. Each option creates a different view of the heat map, with device boxes grouped according to your selection.

Interacting with the Heat Maps

You can interact with the device and port heat maps as follows:

- If you have grouped the devices by location, you can drill down into the heat map's hierarchy by clicking one of the device container names (for example, a site or building). To move back up the hierarchy, click the navigation arrows above the heat map.
- Mouse over a device box to see detailed device-level port utilization information in a pop-up window. In the pop-up window, click the **View top 5 ports** link to view the top five ports that use the most bandwidth on the device.
- Click on a device box to display a heat map of the ports on the device. In this port-level heat map, each port is represented by a box that is color-coded to show its level of utilization. To return to the device view, click the navigation arrows above the heat map.
- Slide the bandwidth utilization control to filter and view devices based on utilization.
- Mouse over a port box to display information about the port—such as port name, status, speed, and percent utilization—in a pop-up window.
 - For ports on devices that support Cloud Analytics Engine, you can view any existing flow analysis results on flows through the port by clicking **View active flows through this link**. See [“Viewing Active Flows on a Port” on page 153](#) for more information.

- For ports on devices that are configured for traffic analysis, you can view the traffic analysis data by clicking **Analyze Traffic on the Port**.
- Slide the circular controls along the bar under the heat map to Filter the devices or ports shown in the heat map by degree of port utilization.

Viewing Active Flows on a Port

For devices that support Cloud Analytics Engine, you can view the results of the most recent flow analysis traces on application flows on the port by mousing over the port and clicking **View active flows through this link**. The Current Active Flows window is displayed.

For information about how to integrate Network Director with Cloud Analytics Engine, see [“Understanding Cloud Analytics Engine and Network Director” on page 83](#). For information about how to start flow analysis on a flow, see [“Virtual Machines & Bare Metal Servers Widget” on page 169](#).

The Current Active Flows window lists only application flows for which flow analysis traces exist—there might be other active application flows on the port that are not shown. Each flow is uniquely defined by source IP address and TCP/UDP port, destination IP address and TCP/UDP port, and transport protocol. [Table 14](#) describes the fields in this window.

Table 14: Fields in the Current Active Flows Window

Field	Description
Source IP Source Port	Source IP address and source TCP/UDP port for the flow. In the case of a flow between two VMS, the IP address is the source VTEP address. If the port is associated with a well-known service, the service name is also shown.
Destination IP Destination Port	Destination IP address and destination TCP/UDP port for the flow. In the case of a flow between two VMS, the IP address is the destination VTEP address. If the port is associated with a well-known service, the service name is also shown.
Protocol	Either TCP or UDP.
Bandwidth	Bandwidth used by the flow. This is a count of the number of packets through the port for the flow up to this point in time. For a value to be displayed in this field, flow analysis must have been performed on flow with the Capture Bandwidth option enabled.
Flow Analysis	Click View Results to see the results of the most recent flow analysis trace. The Flow Analysis Details window opens. NOTE: The View Results link is not available for VM to VM flows.


Flow Analysis Details Window

The Flow Analysis Details window provides detailed information about a flow trace.

The Flow Analysis Details window is divided into three sections:

- The flow path diagram—This diagram shows the path taken by a probe through the network. By default, the path shown is the path taken by the probe that experienced the highest per-hop latency in the trace. You can change this diagram to reflect the path taken by a different probe by selecting the probe from the top Latency Trend chart.
- Latency Trend charts—These charts show the change in latency experienced by the probes during the trace. The bars in the top chart are grouped by completed probes, with each bar in a probe group representing the latency experienced by the probe at a hop. By clicking a probe group, you can change the flow path diagram and the Analysis Results section to reflect the results of that probe. For traces of long duration, the bar chart shows only a portion of the trace results.

The bottom area chart graphs the highest latency experienced by each probe over the entire duration of the trace. You can use the provided controls to focus on a portion of the trace—the portion you choose is reflected in the top bar chart. By default, the focus is on the portion of the trace that had the highest latency. If the trace is ongoing, a rotating circle appears at the end of the plotted area and the chart is periodically refreshed to show new results.

Both charts display a path change icon () when the path a probe takes through the network differs from the path taken by the previous probe.

- Analysis Results—This section provides details about the overall trace results and about the selected probe:
 - The Latency table provides overall latency information for the trace: the highest and lowest latency experienced at a single hop and the average latency of all hops.
 - The Latency for Selected Path table shows the latency experienced by the selected probe at each hop.

You can perform the following actions in the Flow Analysis Details window.

General actions:

- For bidirectional traces, you can select the direction for which you want results by clicking one of the arrows at the top of the window (these arrows do not appear for unidirectional traces).
- To stop an active flow analysis, click **Stop Flow Analysis** at the bottom of the window. When you stop an active flow analysis, the results up to the time you stopped the flow analysis are retained and the previously active trace is marked as complete.

On the flow path diagram, you can:


- Reposition the topology diagram by dragging it or reposition devices by dragging them.
- Zoom in or out by clicking the plus or minus signs on the left.

- Mouse over the link connecting two devices to get the connecting port names. The names are displayed in green if the link is up and in red if the link is down.
- Mouse over a device to view details about the device, such as name, connection state, and IP address. The details shown depends on the device type.

If a device in the flow path does not support Cloud Analytics Engine, it is shown in the diagram in light grey color and minimal details, such as IP address, are available.

- Display the traffic statistics for switches by mousing over the device to display the device details and clicking the **Show Traffic Data** link. If you selected the Capture Bandwidth option when you started the flow analysis, the flow bandwidth is also displayed along with the traffic statistics.
- Display the active flows associated with a VM, BMS, or virtualized host by mousing over the device and clicking **Show Active Flows** in the details box.

On the Latency Trend charts, you can:

- Mouse over a bar group in the top bar chart. A pop-up box displays the latency figures for each hop taken by the probe.
- Click a bar group in the top bar chart. The flow path diagram and the Analysis Results change to reflect the information for the probe.
- Mouse over a path change icon in the top bar chart. Information about the old and new paths is displayed.
- Change the span and position of the focus indicator on the bottom area chart:
 - To increase or decrease the time span of the focus—in other words, to zoom in or zoom out on a portion of the trace—click one of the handle controls () and move it in either direction.
 - To change the focus to another time period, click on the arrows at either end of the slider bar.

Viewing Traffic on a Device

The Traffic on Device window displays the details of traffic that flows through the selected port on a device, such as the applications that are running on the client system, IP address of the client system and the destination, protocol used by the application, data usage, and the data usage percentage. You can choose to view the real-time traffic analysis data on an interface or to view data over a specified period of time.

The Traffic on Device window displays traffic details in two modes—top applications and top conversations. Network Director displays this data in graphical and tabular format, for each mode.

To view details of traffic that passes through a port on a device:

1. Log in to Network Director.

The Dashboard View is selected by default. All the devices that are managed by Network Director in a particular network are represented as cells in the Device & Port utilization widget.

2. Click a device cell to view the ports associated with that device.

All the ports in the selected device are represented as cells. Mouse over the cells to open up a pop-up displaying the port information.

3. Mouse over a cell (port) to view the port information in a pop-up.

4. Click **Analyse Traffic on this Port** in the pop-up.

The Traffic on Device :<port name> page is displayed.

5. Select **Top Applications** (default) or **Top Conversations** to view traffic details sorted based on applications or conversations respectively.

6. Select real-time or a time period for which you want to view the traffic analysis data.

Network Director displays the traffic details on the selected port for the time period you specified. If you chose to view the real-time data, the data in the graph refreshes after each sampling interval.

The graphical view displays traffic from each application or conversation plotted against time (x-axis) and data usage (y-axis). In the Top Applications mode, Network Director displays the names of well-known applications such as *http*, *ftp*, and *ssh*.

[Table 15](#) describes the fields that are displayed in the traffic details table.

7. If you are viewing traffic data in the Top Applications mode and if you know the application that uses a particular protocol port, then select that corresponding port number from the list and click **Associate Application with Port**.

8. Enter the name of the application and the protocol that the application uses. Click **Add**.

The name you entered replaces the name of the application in the list.

Table 15: Traffic on Device—Port Traffic Details Table Fields

Name	Mode	Description
Application	Top Applications	Name of the application.

Table 15: Traffic on Device—Port Traffic Details Table Fields (*continued*)

Name	Mode	Description
Protocol	Top Applications	Protocol that this application uses.
Ingress Bytes	Top Applications Top Conversations	Number of bytes that enter the device through the ingress interface for the given application or conversation for the selected duration.
Egress Bytes	Top Applications Top Conversations	Number of bytes that leave the device through the egress interface for the given application or conversation for the selected duration.
Total Bytes	Top Applications Top Conversations	Total number of bytes that traversed through the port for the given application or conversation for the selected duration.
Percentage of Total Traffic	Top Applications Top Conversations	Percentage of traffic that the application or conversation uses with respect to the total traffic that traverses the port for the selected duration.

RELATED DOCUMENTATION

[Virtual Machines & Bare Metal Servers Widget | 169](#)

[Recent Flow Analysis Widget | 160](#)

[Understanding Cloud Analytics Engine and Network Director | 83](#)

[Understanding the Dashboard | 144](#)

[Using Dashboard Widgets | 145](#)

Equipment By Type Widget

IN THIS SECTION

- [Equipment By Type Widget Summary | 158](#)
- [Equipment By Type Widget Details | 158](#)

The Equipment By Type widget provides summary and detailed information about the types of devices Network Director is managing.

This topic describes:

Equipment By Type Widget Summary

The Equipment By Type widget displays summary information about the types of devices Network Director is managing. The diagram represents the managed devices as a set of nested rings. The circle in the center of the diagram shows information about the ring segments when you mouse over them. The inner ring divides the devices into segments that represent wired and wireless device types. The outer ring divides each of those types into more specific device type segments. Mouse over any diagram segment to see the device type and number of those devices that it represents in the center circle.

Equipment By Type Widget Details

To open the Equipment By Type widget details page, click the **Maximize** button in the widget's title bar. The Equipment By Type widget details window has a table containing detailed information about the devices Network Director is managing. [Table 16](#) describes the columns in the table. To close the details page, click the **Minimize** button in the title bar.

Table 16: Equipment By Type Widget Details Table

Column	Description
Equipment Type	Device equipment type.
Platform	Device platform (model name).
Device Type	Device type.
Software Version	Software version running on the device.

Table 16: Equipment By Type Widget Details Table (*continued*)

Column	Description
Count	Number of devices of that platform in the inventory.

RELATED DOCUMENTATION

[Understanding the Dashboard | 144](#)

[Using Dashboard Widgets | 145](#)

[Network Director Documentation home page](#)

Port Status - Physical Widget

IN THIS SECTION

[Port Status - Physical Widget Summary | 159](#)

[Port Status - Physical Widget Details | 160](#)

The Port Status - Physical widget provides summary and detailed information about the status of physical ports on managed devices.

This topic describes:

Port Status - Physical Widget Summary

The Port Status - Physical widget displays summary information about the status of physical ports on managed devices. It has the following pie charts:

- Admin Status pie chart—Shows the distribution of ports that are administratively up or down and states the total number of ports. Mouse over a chart segment to see more information about it.
- Free vs. Used pie chart—Shows the distribution of ports that are free or used and states the total number of ports. Mouse over a chart segment to see more information about it.

Port Status - Physical Widget Details

The Port Status - Physical widget details window has a table containing detailed information about the status of physical ports on managed devices. See [“Port Status Monitor” on page 1397](#) for descriptions of the table columns.

RELATED DOCUMENTATION

[Understanding the Dashboard | 144](#)

[Using Dashboard Widgets | 145](#)

[Network Director Documentation home page](#)

Recent Flow Analysis Widget

IN THIS SECTION

- [Requirements for Flow Analysis | 161](#)
- [Recent Flow Analysis Main View | 161](#)
- [Flow Analysis Details Window | 163](#)
- [Simulate Flow Analysis Window | 165](#)

The Recent Flow Analysis dashboard widget enables you to view the results of all analyses of application flows that have been initiated by Network Director in your network. It also enables you simulate and analyze a flow between virtual machines (VMs) and between bare metal servers (BMSs) to determine the best placement for a new application in your data center.

Network Director uses Cloud Analytics Engine to perform flow analysis. The Compute Agent component of Cloud Analytics Engine creates a probe, or synthetic packet, that traces the path of the application flow through the network. When a device detects the probe, it collects various metrics that are sent to the Compute Agent, which then sends the metrics to the Data Learning Engine (DLE) component of the Cloud Analytics Engine. Network Director, in turn, obtains this information from DLE. For more information about Cloud Analytics Engine, see [“Understanding Cloud Analytics Engine and Network Director” on page 83](#).

Flow analysis provides you with the number of hops, latency per hop, and end-to-end latency. For each hop, you can view information about the device—for example, CPU utilization, traffic statistics, and ingress

and egress ports used. The information provided enables you to determine congestion points in the network that might be affecting application performance.

In addition to using the Recent Flow Analysis widget, you can perform flow analysis or view the results of flow analysis using the following widgets or tasks:

- Virtual Machine & Bare Metal Servers dashboard widget—This widget enables you to initiate flow analysis on existing application flows and to view the results on a per VM or BMS basis. It also enables you to put a VM or BMS on a watchlist, which automatically starts flow analysis on all flows.
- Device and Port Utilization dashboard widget—This widget provides a graphical view of the level of port utilization on a device. You can mouse over a representation of a port and click **View active flows through this link** to display the results of flow analysis on flows on that port.
- View Connectivity Task in Build mode in the Datacenter view—You can view the flows on a BMS or VM and start analysis on selected flows by mousing over the BMS or VM and clicking **Show active flows**.

This topic describes:

Requirements for Flow Analysis

To perform flow analysis with Network Director, you must:

- Ensure that the components of Cloud Analytics Engine are installed on your network devices and that the Compute Agent discovery file has been created and uploaded to the Data Learning Engine (DLE). See [“Understanding Cloud Analytics Engine and Network Director” on page 83](#) for more information.
- Specify the DLE server IP address under **Preferences > Monitoring > Data Learning Engine Settings**.
- Enable LLDP on the servers hosting the VMs, on the BMSs, and on the connecting switches. In addition, the switches must be discovered by using the SNMP option in Network Director.

Refer to the Cloud Analytics Engine documentation in the [QFX Series switch documentation](#) for more information about the Cloud Analytics Engine.

Network Director performs flow analysis between VMs and BMSs in the following topologies:

- BMSs connected by a Layer 3 IP fabric
- VMs connected by VMware NSX or OpenStack VXLAN tunnels overlaid on a Layer 3 IP fabric

Recent Flow Analysis Main View

The main view of the Recent Flow Analysis widget lists the 10 most recent flow analyses performed by Network Director, as determined by their start time. You can view all flow analyses by clicking the expand icon at the top of the widget or **View all** at the bottom of the widget.

Each flow analysis consists of one or more traces—a trace being a period during which flow analysis was active. Multiple traces can exist because a flow can become inactive and then active again during the duration of a flow analysis or because multiple analyses have been performed on the same flow over a period of time. By default, the summary and expanded views of the Recent Flow Analysis widget show the results of the most recent trace for a flow.

You can take the following actions:

- Click **View Details** to view in-depth analysis of the trace results.
- Click **All Traces** to view the results of all traces performed on a flow.

[Table 17](#) describes the fields in the Recent Flow Analysis expanded view.

Table 17: Recent Flow Analysis Expanded View

Field	Description
Source IP Source Port	Source IP address and source TCP/UDP port for the flow. If the port is associated with a well-known service, the service name is also shown.
Destination IP Destination Port	Destination IP address and destination TCP/UDP port for the flow. If the port is associated with a well-known service, the service name is also shown.
Status	Status of the most recent trace: <ul style="list-style-type: none"> • Active (rotating circle)—Trace has started and is ongoing. • Completed (check mark)—Trace completed at the end time shown. • Failed (red triangle)—Trace was unable to complete. • Scheduled (clock)—Trace is scheduled to start and finish at the start and end times shown.
Start Time End Time	Start and end time for the most recent trace.
Min Latency	Lowest end-to-end latency experienced by a probe during the trace, in milliseconds. NOTE: Time drift on the network devices between NTP synchronizations can result in negative values.
Max Latency	Highest end-to-end latency experienced by a probe during the trace, in milliseconds.
Avg Latency	Average end-to-end latency experienced by the probes in the trace, in milliseconds.

Table 17: Recent Flow Analysis Expanded View (continued)

Field	Description
Actions	<p>Click View Details to open the Flow Analysis Details window, which provides in-depth path analysis for this trace.</p> <p>Click All Traces to display all previously completed traces for this flow. For each trace, the information described in this table is provided.</p>


Flow Analysis Details Window

The Flow Analysis Details window provides detailed information about a flow trace.

The Flow Analysis Details window is divided into three sections:

- The flow path diagram—This diagram shows the path taken by a probe through the network. By default, the path shown is the path taken by the probe that experienced the highest per-hop latency in the trace. You can change this diagram to reflect the path taken by a different probe by selecting the probe from the top Latency Trend chart.
- Latency Trend charts—These charts show the change in latency experienced by the probes during the trace. The bars in the top chart are grouped by completed probes, with each bar in a probe group representing the latency experienced by the probe at a hop. By clicking a probe group, you can change the flow path diagram and the Analysis Results section to reflect the results of that probe. For traces of long duration, the bar chart shows only a portion of the trace results.

The bottom area chart graphs the highest latency experienced by each probe over the entire duration of the trace. You can use the provided controls to focus on a portion of the trace—the portion you choose is reflected in the top bar chart. By default, the focus is on the portion of the trace that had the highest latency. If the trace is ongoing, a rotating circle appears at the end of the plotted area and the chart is periodically refreshed to show new results.

Both charts display a path change icon () when the path a probe takes through the network differs from the path taken by the previous probe.

- Analysis Results—This section provides details about the overall trace results and about the selected probe:
 - The Latency table provides overall latency information for the trace: the highest and lowest latency experienced at a single hop and the average latency of all hops.
 - The Latency for Selected Path table shows the latency experienced by the selected probe at each hop.

You can perform the following actions in the Flow Analysis Details window.

General actions:

- For bidirectional traces, you can select the direction for which you want results by clicking one of the arrows at the top of the window (these arrows do not appear for unidirectional traces).
- To stop an active flow analysis, click **Stop Flow Analysis** at the bottom of the window. When you stop an active flow analysis, the results up to the time you stopped the flow analysis are retained and the previously active trace is marked as complete.


On the flow path diagram, you can:

- Reposition the topology diagram by dragging it or reposition devices by dragging them.
- Zoom in or out by clicking the plus or minus signs on the left.
- Mouse over the link connecting two devices to get the connecting port names. The names are displayed in green if the link is up and in red if the link is down.
- Mouse over a device to view details about the device, such as name, connection state, and IP address. The details shown depends on the device type.

If a device in the flow path does not support Cloud Analytics Engine, it is shown in the diagram in light grey color and minimal details, such as IP address, are available.

- Display the traffic statistics for switches by mousing over the device to display the device details and clicking the **Show Traffic Data** link. If you selected the Capture Bandwidth option when you started the flow analysis, the flow bandwidth is also displayed along with the traffic statistics.
- Display the active flows associated with a VM, BMS, or virtualized host by mousing over the device and clicking **Show Active Flows** in the details box.

On the Latency Trend charts, you can:

- Mouse over a bar group in the top bar chart. A pop-up box displays the latency figures for each hop taken by the probe.
- Click a bar group in the top bar chart. The flow path diagram and the Analysis Results change to reflect the information for the probe.
- Mouse over a path change icon in the top bar chart. Information about the old and new paths is displayed.
- Change the span and position of the focus indicator on the bottom area chart:
 - To increase or decrease the time span of the focus—in other words, to zoom in or zoom out on a portion of the trace—click one of the handle controls () and move it in either direction.
 - To change the focus to another time period, click on the arrows at either end of the slider bar.

Simulate Flow Analysis Window

The Simulate Flow Analysis window enables you to create a simulated application flow between VMs or between BMs in a data center and to analyze the flow results. By performing analysis of simulated flows, you can determine the best placement for new applications in the data center.

You can simulate flows between:

- Two existing BMSs in a non-overlay network.
- Between two VMs in an overlay (VXLAN) network. The VMs do not need to exist to simulate a flow between them.

To begin simulated flow analysis, you must specify which type of flow you are simulating:

- Select **Non Overlay Network** to simulate flows between two BMSs. [Table 18](#) describes the settings for this option.
- Select **Overlay Network** to simulate flows between two VMs. [Table 19](#) describes the settings for this options.

Table 18: Non Overlay Network Settings

Field	Action
Select Datacenter	Select a data center from the list.
IP Address	Select the source and destination IP addresses for the simulated flows. The IP addresses listed are existing IP addresses on the BMSs in the data center.
Port	Source and destination TCP/UDP ports.
Protocol	Select protocol—TCP or UDP.

Table 19: Overlay Network Settings

Field	Action
Select Datacenter	Select an existing data center from the list.
VNI ID	Select the VNI (VXLAN Network Identifier) for the overlay network.
VTEP IP Address	Select the source and destination VTEP (VXLAN Tunnel Endpoint) IP addresses.

Table 19: Overlay Network Settings (continued)

Field	Action
MAC Address	<p>If the source and destination VMs exist, enter their MAC addresses.</p> <p>If the source and destination VMs do not exist and the tunnel is from hypervisor to hypervisor without a VXLAN gateway between the hypervisors, enter the MAC addresses for possible virtual Ethernet ports on the servers where these VM might reside. If you do not know the VM location or the VM location is not decided, enter made-up MAC addresses.</p> <p>If the source and destination VMs do not exist and there is a VXLAN gateway on the tunnel between the hypervisors, enter MAC addresses from the MAC pools on the servers or hosts.</p>
Port	Source and destination TCP/UDP ports.
Protocol	Select protocol—TCP or UDP.

After you have defined the basic settings for the type of flow you are simulating, click **Simulate Flow & Analyze** to begin the simulated flow. You can also click **Advanced Settings** to configure the advanced settings described in [Table 20](#) before you start the flow analysis.

Table 20: Simulate Flow Analysis Advanced Settings

Field	Action
Schedule	<p>Run Now—Select to start the flow analysis as soon as you click Simulate Flow & Analyze.</p> <p>Schedule Later—Select to schedule the analysis to run at a later time and enter the date and time in the fields provided. When you click Simulate Flow & Analyze, the analysis is scheduled to run at the date and time you specified.</p> <p>Scheduled flow analyses are shown only on the Recent Flow Analysis widget and cannot be managed as scheduled jobs under System > Manage Jobs.</p>
Duration	Specify how long the flow analysis runs.
Frequency	Specify how often a probe is sent on the flow path during the duration in which flow analysis is active.
Timeout	Specify how long Network Director waits before timing out the flow analysis after a probe fails to respond.
Max. number of hops	Specify the maximum number of hops on which flow analysis is performed.

Table 20: Simulate Flow Analysis Advanced Settings (continued)

Field	Action
Select Direction	Click an arrow to select the direction of the flow you want to analyze: the right arrow for source to destination, the left arrow for destination to source, and the double-headed arrow for both directions.
Capture Flow Bandwidth	Select to have Network Director collect and report information about the flow bandwidth. If you select this option, you can view the captured flow bandwidth by: <ul style="list-style-type: none"> • Displaying the traffic statistics by mousing over a device in the Flow Analysis Details window. • Viewing active flows through a port in the Device & Port Utilization dashboard wizard.
Mirror Flows	Select to copy the packets in the flow to the IP address you specify.

RELATED DOCUMENTATION

[Virtual Machines & Bare Metal Servers Widget | 169](#)

[Understanding Cloud Analytics Engine and Network Director | 83](#)

[Understanding the Dashboard | 144](#)

[Using Dashboard Widgets | 145](#)

[Network Director Documentation home page](#)

Top Talker - Wired Devices Widget

IN THIS SECTION

[Top Talker - Wired Devices Widget Summary | 168](#)

[Top Talker - Wired Devices Widget Details | 168](#)

The Top Talker - Wired Devices widget provides summary and detailed information about the hosts that are using the most bandwidth. Hosts are endpoints that are directly connected to access ports of wired switches.

This topic describes:

Top Talker - Wired Devices Widget Summary

The Top Talker - Wired Devices widget has a bar chart that shows summary information about the hosts that are using the most bandwidth. Host names are listed on the vertical axis. Data usage in kilobytes is shown on the horizontal axis. Mouse over a bar to see more information about that host.

Top Talker - Wired Devices Widget Details

To open the Top Talker - Wired Devices widget details page, click the **Maximize** button in the widget's title bar. The Top Talker - Wired Devices widget details window has a table containing detailed information about the hosts that are using the most bandwidth. [Table 21](#) describes the columns in the table. To close the details page, click the **Minimize** button in the title bar.

Table 21: Top Talker - Wired Devices Widget Details

Column	Description
Host Name	Host's host name.
MAC Address	Host's MAC address
Data Usage (KBytes)	Data used by the host, in kilobytes.
Device Serial Number	Device's serial number.

RELATED DOCUMENTATION

[Understanding the Dashboard | 144](#)

[Using Dashboard Widgets | 145](#)

[Network Director Documentation home page](#)

Top Virtual Machines by Bandwidth Widget

The Top Virtual Machines by Bandwidth widget displays a bar chart of the virtual machines that are using the most bandwidth. Each horizontal bar represents a virtual machine. The horizontal axis shows the bandwidth utilization of the virtual machines in kilobits per second. You can mouse over a bar to see more information about that virtual machine.

RELATED DOCUMENTATION

[Understanding the Dashboard | 144](#)

[Using Dashboard Widgets | 145](#)

[Network Director Documentation home page](#)

Top vNetwork Hosts by Bandwidth Widget

The Top vNetwork Hosts by Bandwidth widget displays a bar chart of the virtual hosts that are using the most bandwidth. Each horizontal bar represents a virtual host. The horizontal axis shows the percentage of bandwidth utilization. Mouse over a bar to see more information about that host.

RELATED DOCUMENTATION

[Understanding the Dashboard | 144](#)

[Using Dashboard Widgets | 145](#)

[Network Director Documentation home page](#)

Virtual Machines & Bare Metal Servers Widget

IN THIS SECTION

- [Requirements for Flow Analysis | 170](#)
- [Virtual Machines & Bare Metal Servers Widget Main View | 171](#)
- [Initiating Analysis on Selected Flows from the All Tab | 171](#)
- [Initiating Flow Analysis on All Flows on a VM or BMS from the Watchlist Tab | 173](#)
- [Viewing the Results of Flow Analysis | 176](#)

The Virtual Machines & Bare Metal Servers widget provides information about the application flows on virtual machines (VMs) and bare metal servers (BMSs) in your data center. Use the widget to start flow analysis on selected active flows on a specific VM or a BMS and to view the analysis results. You can also use this widget to place a critical VM or BMS on a watchlist. Network Director will automatically initiate analysis on all flows on that VM or BMS.

Network Director uses Cloud Analytics Engine to perform flow analysis. The Compute Agent component of Cloud Analytics Engine creates a probe, or synthetic packet, that traces the path of the application flow through the network. When a device detects the probe, it collects various metrics that are sent to the Compute Agent, which then sends the metrics to the Data Learning Engine (DLE) component of the Cloud Analytics Engine. Network Director, in turn, obtains this information from DLE. For more information about Cloud Analytics Engine, see [“Understanding Cloud Analytics Engine and Network Director” on page 83](#).

Flow analysis provides you with the number of hops, latency per hop, and end-to-end latency. For each hop, you can view information about the device—for example, CPU utilization, ingress and egress ports used, and traffic statistics. The information provided enables you to identify congestion points in the network that might be affecting application performance.

In addition to using the Virtual Machines & Bare Metal Servers widget, you can perform flow analysis or view the results of flow analysis by using the following widgets or tasks:

- Recent Flow Analysis dashboard widget—This widget enables you to analyze simulated flows and to view the results of all flow analyses in the network, independent of the associated VM or BMS.
- Device and Port Utilization dashboard widget—This widget provides a graphical view of the level of port utilization on a device. You can mouse over a representation of a port and click **View active flows through this link** to display the results of analysis on flows on that port.
- View Connectivity Task in Build mode in the Datacenter view—You can view the flows on a BMS or VM and start analysis on selected flows by mousing over the BMS or VM and clicking **Show active flows**.

This topic describes:

Requirements for Flow Analysis

To perform flow analysis with Network Director, you must:

- Ensure that the components of Cloud Analytics Engine are installed on your network devices and that the Compute Agent discovery file has been created and uploaded to the Data Learning Engine (DLE).
See [“Understanding Cloud Analytics Engine and Network Director” on page 83](#) for more information.
- Specify the DLE IP address under **Preferences > Monitoring > Data Learning Engine Settings**.
- Enable LLDP on the servers hosting the VMs, on the BMSs, and on the connecting switches. In addition, the switches must be discovered by using the SNMP option in Network Director.

See the Cloud Analytics Engine documentation in the [QFX Series switch documentation](#) for more information about Cloud Analytics Engine.

Network Director performs flow analysis between VMs and BMSs in the following topologies:

- BMSs connected by a Layer 3 IP fabric
- VMs connected by VMware NSX or OpenStack VXLAN tunnels overlaid on a Layer 3 IP fabric

Virtual Machines & Bare Metal Servers Widget Main View

The main view of the Virtual Machines & Bare Metal Servers widget has two tabs that provide access to different flow analysis functions:

- **All tab**—Use this tab to analyze selected flows. You can choose which flows to analyze from all active flows on a VM or BMS. You can analyze flows in either direction (source to destination or destination to source) or in both directions, and you can schedule the start and duration of the flow analysis.
- **Watchlist tab**—Use this tab to analyze all flows on a VM or BMS. When you add a VM or BMS to the watchlist, Network Director starts analysis on all active flows on the device. While the device is on the watchlist, Network Director automatically starts flow analysis on any new flows that appear. To prevent the introduction of unintended overhead on devices that are not on the watchlist, flow analysis is always unidirectional—from source to destination.

Initiating Analysis on Selected Flows from the All Tab

The All tab for the Virtual Machines & Bare Metal Servers widget provides a list of VMs and BMSs known to Network Director. From this list, you can select the VMs or BMSs on which you want to perform flow analysis or for which you want to view flow analysis results.

By default, the All tab displays the 10 VMs or BMSs with the highest bandwidth utilization. To view all VMs and BMSs known to Network Director, click the expand icon at the top of the widget or **View all** at the bottom of the widget.

You can take the following actions:

- Click **Analyze Flows** for the VM or BMS whose flows you want to analyze. Doing so displays the Analyze Flows window. If the VM or BMS does not have Cloud Analytics Engine support, the **Analyze Flows** link is disabled.
- Click **All Traces** to view the results of previous or on-going flow analyses on the VM or BMS. For information about viewing results, see [“Viewing the Results of Flow Analysis” on page 176](#).

To support flow analytics, a device must have Cloud Analytics Engine Compute Agent support. For more information, see [“Understanding Cloud Analytics Engine and Network Director” on page 83](#)

NOTE: Each Compute Agent can analyze only up to 20 flows at a time.

[Table 22](#) describes the fields in the All tab.

Table 22: All Tab Fields

Field	Description	Default Display
Name	Name assigned to the VM or BMS.	Visible
Bandwidth Utilization (KBps)	Total bandwidth consumption in kilobytes per second.	Visible
Host	For VMs, the name of the host on which the virtual machine is running.	Hidden
IP Address	IP address of the VM or BMS.	Visible
Actions	<p>Click Analyze Flows for the device whose flows you want to analyze. Doing so displays the Current Flows window.</p> <p>Click View Results to view the results of previous or on-going flow analyses on the virtual machine.</p>	Visible

Current Flows Window

The Current Flows window shows the active flows on the selected VM or BMS. Each flow is uniquely defined by source IP address and TCP/UDP port, destination IP address and TCP/UDP port, and transport protocol—either TCP or UDP. When a port number is commonly associated with a well-known service, the service name is also shown.

You can perform the following actions:

- To start flow analysis immediately, select one or more flows and click **Run Flow Analysis**. By default, flow analysis runs for 2 hours, with a probe sent every 20 seconds.
- To schedule the flow analysis to start at a later time or to change the flow analysis settings, click **Scheduling & Advanced Settings** before you click **Run Flow Analysis**. The Scheduling and Advanced Settings window appears, which enables you to configure the optional settings described in [Table 23](#).

Table 23: Scheduling and Advanced Settings

Setting	Action
Schedule	<p>Run Now—Select to start the flow analysis as soon as you click Run Flow Analysis.</p> <p>Schedule Later—Select to schedule the analysis to run at a later time and enter the date and time in the fields provided. When you click Run Flow Analysis, the analysis is scheduled to run at the date and time you specified.</p> <p>Scheduled flow analyses are shown only on the dashboard widgets and cannot be managed as scheduled jobs under System > Manage Jobs.</p>

Table 23: Scheduling and Advanced Settings (*continued*)

Setting	Action
Duration	Specify how long the flow analysis runs.
Frequency	Specify how often a probe is sent on the flow path during the duration in which flow analysis is active.
Timeout	Specify how long the Compute Agent waits before timing out the flow analysis after a probe fails to respond.
Max. number of hops	Specify the maximum number of hops on which flow analysis is performed.
Select Direction	Click the arrow that indicates the direction of the flow you want to analyze: the right arrow for source to destination, the left arrow for destination to source, and the double-headed arrow for both directions.
Capture Flow Bandwidth	<p>Select to have Network Director collect and report information about the flow bandwidth.</p> <p>If you select this option, you can view the captured flow bandwidth by:</p> <ul style="list-style-type: none"> • Displaying the traffic statistics by mousing over a device in the Flow Analysis Details window, as described in “Viewing the Results of Flow Analysis” on page 176. • Viewing active flows through a port in the Device & Port Utilization dashboard wizard.
Mirror Flows	Select to copy the packets in the flow to the IP address you specify.
Apply above settings to all other selected flows	Select to apply these settings to all flows that are currently selected.

Initiating Flow Analysis on All Flows on a VM or BMS from the Watchlist Tab

The Watchlist tab of the Virtual Machines & Bare Metal Servers widget lists all VMs or BMSs currently on the watchlist. When you place a VM or BMS on the watchlist, flow analysis traces start immediately on all flows in the VM or BMS. Each trace lasts for an hour. If the VM or BMS remains on the watchlist after the hour is up, another hour-long trace is started. During the time the VM or BMS is on the watchlist, Network Director automatically starts flow analysis on any new flows that appear. To prevent the introduction of unintended overhead on devices that are not on the watchlist, flow analysis is always unidirectional—from source to destination.

You can take the following actions on this tab:

- Click **Add to Watchlist** to open the Add to Watchlist window, which lists all the VMs and BMSs you can add to the watchlist.
- Click **All Traces** to open the Flow Analysis Results window and view the flow analysis results for a device.
- Click **Remove** to remove a VM or BMS from the watchlist. All active traces are stopped and the VM or BMS is removed from the Watchlist tab. You can still access the flow analysis results from the All tab or from the Recent Flow Analysis widget.

TIP: If the Watchlist tab does not change immediately to reflect the VMs or BMSs you have added or removed from the watchlist, refresh the widget by clicking the Refresh icon on the widget title bar.

Table 24 describes the fields in the Watchlist tab.

Table 24: Watchlist Tab Fields

Field	Description	Default Display
Name	Name assigned to the VM or BMS.	Visible
Bandwidth Utilization (Kbps)	Total bandwidth consumption in kilobits per second.	Visible
Host	For VMs, the name of the host on which the virtual machine is running.	Hidden
IP Address	IP address of the VM or BMS.	Visible
Actions	Click Analyze Flows for the device whose flows you want to analyze. Doing so displays the Analyze Flows for VM window. Click All Traces to view the results of previous or on-going flow analyses on the virtual machine.	Visible

Add to Watchlist Window

To add VMs and BMSs to the watchlist, click **Add to Watchlist** on the Watchlist tab. The Add to Watchlist window opens, which lists the VMs and BMSs that can be added to the watchlist.

To add VMs and BMSs to the watchlist, select one or more VMs or BMSs and click **Add to Watchlist**.

NOTE: If a VM or BMS has more than 20 active flows or if adding a VM or BMS to the watchlist would result in more than 20 flows being traced by a Compute Agent, flow analysis on the VM or BMS is not started.

Analysis on all flows starts immediately on all the VMs and BMSs you add to the watchlist. Each trace runs for an hour. After an hour, another trace automatically starts unless you remove the VM or BMS from the watchlist.

By default, a probe is sent every 20 seconds. Because all flows on the VM or BMSs are traced, you might want to reduce the frequency with which probes are sent to one every 30 seconds or longer. To change probe frequency and other settings, click **Advanced Settings** before you click **Add to Watchlist**. The available advanced settings are described in [Table 25](#).

Table 25: Add to Watchlist Advanced Settings

Setting	Action
Frequency	Specify how often a probe is sent on the flow path during the duration in which flow analysis is active.
Timeout	Specify how long the Compute Agent waits before timing out the flow analysis after a probe fails to respond.
Max. number of hops	Specify the maximum number of hops on which flow analysis is performed.
Capture Flow Bandwidth	<p>Select to have Network Director collect and report information about the flow bandwidth.</p> <p>If you select this option, you can view the captured flow bandwidth by:</p> <ul style="list-style-type: none"> • Displaying the traffic statistics by mousing over a device in the Flow Analysis Details window, as described in "Viewing the Results of Flow Analysis" on page 176. • Viewing active flows through a port in the Device & Port Utilization dashboard wizard.
Mirror Flows	Select to copy the packets in the flow to the IP address you specify.
Apply above settings to all other selected flows	Select to apply these settings to all VMs and BMSs that are currently selected in the Add to Watchlist window.

Viewing the Results of Flow Analysis

IN THIS SECTION

- [Flow Analysis Results Window | 176](#)
- [Flow Analysis Details Window | 177](#)

You can view the flow analysis results for a VM or BMS by clicking **All Traces** on the All tab or Watchlist tab.

You can also view flow analysis results from the Recent Flow Analysis dashboard widget.

The following sections describe:

Flow Analysis Results Window

The Flow Analysis Results window shows summary results for the flow analyses run on the VM or BMS.

Each flow analysis consists of one or more traces—a trace being a period during which flow analysis was active. Multiple traces can exist because a flow can become inactive and then active again during the duration of a flow analysis or because multiple analyses have been performed on the same flow over a period of time. By default, the Flow Analysis Results window shows the results for the most recent trace for each flow.

You can:

- Click **View Details** to view in-depth analysis of the trace results.
- Click **All Traces** to view the results of all traces performed on a flow.

For each flow, the information described in [Table 26](#) is displayed.

Table 26: Flow Analysis Results Fields

Field	Description
Source IP Source Port	Source IP address and source TCP/UDP port for the flow. If the port is associated with a well-known service, the service name is also shown.
Destination IP Destination Port	Destination IP address and destination TCP/UDP port for the flow. If the port is associated with a well-known service, the service name is also shown.

Table 26: Flow Analysis Results Fields (continued)

Field	Description
Status	<p>Status of the most recent trace analysis:</p> <ul style="list-style-type: none"> • Active (rotating circle)—Trace has started and is ongoing. • Completed (check mark)—Trace completed at the end time shown. • Failed (red triangle)—Trace was unable to complete. • Scheduled (clock)—Trace is scheduled to start at the start time shown.
Start Time End Time	Start and end time for the most recent trace analysis.
Min Latency	<p>Lowest end-to-end latency experienced during the trace, in milliseconds.</p> <p>NOTE: Time drift on the network devices between NTP synchronizations can result in negative values.</p>
Max Latency	Highest end-to-end latency experienced during the trace, in milliseconds.
Avg Latency	Average end-to-end latency experienced by the trace, in milliseconds.
Action	<p>Click View Details to open the Flow Analysis Details window, which provides in-depth path analysis for this trace.</p> <p>Click All Traces to display all previously completed traces for this flow. For each trace, the information described in this table is provided.</p>

Flow Analysis Details Window


The Flow Analysis Details window provides detailed information about a flow trace.

The Flow Analysis Details window is divided into three sections:

- The flow path diagram—This diagram shows the path taken by a probe through the network. By default, the path shown is the path taken by the probe that experienced the highest per-hop latency in the trace. You can change this diagram to reflect the path taken by a different probe by selecting the probe from the top Latency Trend chart.
- Latency Trend charts—These charts show the change in latency experienced by the probes during the trace. The bars in the top chart are grouped by completed probes, with each bar in a probe group representing the latency experienced by the probe at a hop. By clicking a probe group, you can change the flow path diagram and the Analysis Results section to reflect the results of that probe. For traces of long duration, the bar chart shows only a portion of the trace results.

The bottom area chart graphs the highest latency experienced by each probe over the entire duration of the trace. You can use the provided controls to focus on a portion of the trace—the portion you choose

is reflected in the top bar chart. By default, the focus is on the portion of the trace that had the highest latency. If the trace is ongoing, a rotating circle appears at the end of the plotted area and the chart is periodically refreshed to show new results.

Both charts display a path change icon () when the path a probe takes through the network differs from the path taken by the previous probe.

- **Analysis Results**—This section provides details about the overall trace results and about the selected probe:
 - The Latency table provides overall latency information for the trace: the highest and lowest latency experienced at a single hop and the average latency of all hops.
 - The Latency for Selected Path table shows the latency experienced by the selected probe at each hop.

You can perform the following actions in the Flow Analysis Details window.

General actions:

- For bidirectional traces, you can select the direction for which you want results by clicking one of the arrows at the top of the window (these arrows do not appear for unidirectional traces).
- To stop an active flow analysis, click **Stop Flow Analysis** at the bottom of the window. When you stop an active flow analysis, the results up to the time you stopped the flow analysis are retained and the previously active trace is marked as complete.


On the flow path diagram, you can:

- Reposition the topology diagram by dragging it or reposition devices by dragging them.
- Zoom in or out by clicking the plus or minus signs on the left.
- Mouse over the link connecting two devices to get the connecting port names. The names are displayed in green if the link is up and in red if the link is down.
- Mouse over a device to view details about the device, such as name, connection state, and IP address. The details shown depends on the device type.

If a device in the flow path does not support Cloud Analytics Engine, it is shown in the diagram in light grey color and minimal details, such as IP address, are available.

- Display the traffic statistics for switches by mousing over the device to display the device details and clicking the **Show Traffic Data** link. If you selected the Capture Bandwidth option when you started the flow analysis, the flow bandwidth is also displayed along with the traffic statistics.
- Display the active flows associated with a VM, BMS, or virtualized host by mousing over the device and clicking **Show Active Flows** in the details box.

On the Latency Trend charts, you can:

- Mouse over a bar group in the top bar chart. A pop-up box displays the latency figures for each hop taken by the probe.
- Click a bar group in the top bar chart. The flow path diagram and the Analysis Results change to reflect the information for the probe.
- Mouse over a path change icon in the top bar chart. Information about the old and new paths is displayed.
- Change the span and position of the focus indicator on the bottom area chart:
 - To increase or decrease the time span of the focus—in other words, to zoom in or zoom out on a portion of the trace—click one of the handle controls () and move it in either direction.
 - To change the focus to another time period, click on the arrows at either end of the slider bar.

RELATED DOCUMENTATION

[Recent Flow Analysis Widget | 160](#)

[Understanding Cloud Analytics Engine and Network Director | 83](#)

[Understanding the Dashboard | 144](#)

[Using Dashboard Widgets | 145](#)

[Network Director Documentation home page](#)

Top Overlay Networks Widget

IN THIS SECTION

[Top Overlay Networks Widget Summary | 180](#)

[Top Overlay Networks Widget Details | 180](#)

Virtual Extensible Local Area Network (VXLAN) represents a technology that enables you to segment your networks (as VLANs do), but that also solves the scaling limitation of VLANs and provides benefits that VLANs cannot. VXLAN is often described as an overlay technology because it enables you to stretch Layer 2 connections over an intervening Layer 3 network by encapsulating (tunneling) Ethernet frames in a VXLAN packet that includes IP addresses.

This topic describes:

Top Overlay Networks Widget Summary

The Top Overlay Networks widget displays a summary of the VXLANs in your network in a table. [Table 27](#) displays the details that are displayed in this table.

Table 27: Top Overlay Networks Widget Summary page Field Descriptions

Column	Description
VXLAN	Unique ID of the VXLAN.
Tenant	Name of the tenant that uses the given VXLAN.
Network	IP address and the subnet mask of the network that is assigned to the tenant.
VMs	Number of virtual machines (VMs) that are active for the tenant.
Aggregate Bandwidth	Aggregated bandwidth used by the overlay network for the last 10 minutes.
Datacenter	Name of the data center to which the tenant is connected.

Click **Show Details** corresponding to a VXLAN entry to view detailed information about that VXLAN. The Top Overlay Networks Widget Details page opens.

Top Overlay Networks Widget Details

To open the Top Overlay Networks Widget Details page, click **Show Details** in the Top Overlay Networks Widget table in the Summary View. Top Overlay Networks Widget Details page has two tabs—List View and the Topology View.

The list view displays detailed information about the VMs that are part of the selected overlay network or VXLAN. [Table 28](#) describes the fields in this page.

Table 28: Top Overlay Details page Field Descriptions

Column	Description
VM Name	Name of the VM.
BW Utilization	Average bandwidth utilized by the VM for the last 10 minutes.
Host Name	The ESXi hostname of the VM.

Table 28: Top Overlay Details page Field Descriptions (*continued*)

Column	Description
Guest Operating System	Operating system running on the VM.
IP Address	IP address of the VM.
MAC Address	MAC address of the VM.

The Topology view highlights the VMs and the bare metal servers that are part of the selected VXLAN. You can mouse over each entity to view more details.

RELATED DOCUMENTATION

[Understanding the Dashboard | 144](#)

[Using Dashboard Widgets | 145](#)

[Network Director Documentation home page](#)

3

PART

Working in Build Mode

About Build Mode | **184**

Discovering Devices | **204**

Setting Up Sites and Locations Using the Location View | **222**

Building a Topology View of the Network | **242**

Creating Custom Device Groups | **273**

Configuring Quick Templates | **282**

Configuring Device Settings | **290**

Configuring Authentication, Authorization, and Access for Your Network | **335**

Configuring Interfaces and VLANs | **408**

Configuring Firewall Filters (ACLs) | **540**

Configuring Class of Service (CoS) | **609**

Configuring Media Access Control Security (MACsec) | **637**

Configuring Link Aggregation Groups (LAGs) | **646**

Configuring Fibre Channel Gateways | **683**

Creating Configurations for Fabrics | **693**

Creating and Managing Datacenter Fabrics | **707**

Configuring Cloud-Based Datacenter Networks | **783**

Configuring Overlay Networks and Tenants | **828**

Configuring VRRP Profiles | **845**

Configuring Wireless Access Points and Radios | **851**

Configuring Wireless Controllers | **1037**

Configuring Wireless Mobility and Network Domains | **1051**

Configuring WLAN Service (SSIDs) | **1061**

Managing Network Devices | **1135**

About Build Mode

IN THIS CHAPTER

- [Understanding Build Mode in Network Director | 184](#)
- [Understanding the Build Mode Tasks Pane | 189](#)
- [Understanding Network Configuration Profiles | 197](#)
- [Assigning Profiles to an Interface, Device, or a Group of Devices | 201](#)

Understanding Build Mode in Network Director

IN THIS SECTION

- [Discovering Devices | 184](#)
- [Building the Logical, Location, and Custom Views | 185](#)
- [Configuring Devices | 186](#)
- [Managing Devices | 188](#)

In Build mode, you build the network managed by Junos Space Network Director. It provides you with the ability to use device discovery to bring devices under Network Director management, to customize your view of the devices, to configure devices, and to perform some common device management tasks.

This topic describes:

Discovering Devices

Device discovery finds your network devices and brings them under Network Director management. You provide Network Director with identifying information about the devices you want Network Director to manage—an IP address or hostname, an IP address range, an IP subnetwork, or a CSV file that contains this information. Network Director uses the information to probe the devices by using either ping or SNMP get requests. If a device probe is successful, Network Director then attempts to make an SSH connection

to the device using the login credentials you supply. If the connection is successful and the device is a supported device, Network Director adds the device to its database of managed devices. Network Director uses Juniper Network's Device Management Interface (DMI), which is an extension to the NETCONF network configuration protocol, to connect to and configure its managed devices.

You can also discover devices using the device discovery feature provided by the Junos Space Network Management Platform. Devices you discover using Junos Space device discovery are brought under Network Director management if they are supported by Network Director.

Besides bringing your devices under Network Director management, device discovery:

- Reads the device configuration and saves it in the Junos Space configuration database. Network Director uses this record of the device configuration to determine what configuration commands it needs to send to a device when you deploy the configuration on the device. For this reason, it is important for the Junos Space configuration record to match, or be in sync with, the device configuration. For more information about how the Junos Space configuration record and device configuration are kept in sync, see [“Understanding Resynchronization of Device Configuration” on page 1214](#).
- Imports the device configuration into the Build mode configuration. For more information about importing device configurations, see [“Importing Device Configurations” on page 187](#).

Building the Logical, Location, and Custom Views

When a device is discovered in the physical network mode, it is added to the network tree in the View pane. In Logical View, all switches are added to the Unassigned node under the switching network. You can then assign them to the Access, Aggregation, or Core nodes to complete the Logical View of the switching network. The Logical View of the wireless network is built for you.

Similarly, in Location View, all discovered devices are added to the Unassigned node. You can use Build mode to create the Location View—that is, create the sites, buildings, floors, closets, and outdoor areas that reflect the physical location of your network devices—and to assign the discovered devices to these locations.

NOTE: Network Director displays the Virtual Chassis, Virtual Chassis Fabric, and QFabric systems in the Location view network tree only if at least one of their member devices are *not* assigned to any location entity. If all the member devices are assigned to location entities, then the Virtual Chassis, Virtual Chassis Fabric, and QFabric systems are not displayed in the network tree.

The Custom Group View displays only the top level—My Network—until you create one or more custom groups. Custom group is another way of grouping your devices based on your business needs. You can create custom groups and add devices to each custom group. You can manually add devices to a custom group or you can define rules to add devices, that match the rule condition, to the custom group once they

are discovered by Network Director. You can view the custom groups and devices that are assigned to each group in the Custom Group view.

NOTE: This section does not apply to virtual devices that Network Director manages.

Configuring Devices

In Build mode, you can define the configuration of network devices in your Physical network. To support rapid, large-scale deployment of devices, you can define much of your Build mode configuration in a set of profiles. You can reference profiles in other profiles or apply them to multiple objects in your network—devices, ports, radios, logical entities. For example, you can create a Port profile that sets up class-of-service (CoS), authentication, firewall filters, and Ethernet switching settings that are appropriate for access ports that connect to employee desktop VoIP phones and then assign that profile to access ports across multiple switches.

NOTE: This section does not apply to virtual devices that Network Director manages.

[Figure 8](#) shows an example landing page for profile configuration, in this case VLAN profiles. This page lists all existing VLAN profiles. From this page you can create new profiles, modify or delete existing profiles, assign profiles to objects, and change or view assignments. For more information about profiles, see [“Understanding Network Configuration Profiles” on page 197](#).

Figure 8: Manage VLAN Profiles Page

Profile Name	VLAN Name	Family Type	VLAN ID	VLAN Range	VLAN ID List	Description	Assignment State
v300_152	v300	Campus Switching ELS	300	-	-	-	Deployed
v200_150	v200	Campus Switching ELS	200	-	-	-	Deployed
v123_148	v123	Campus Switching ELS	123	-	-	-	Deployed
v100_146	v100	Campus Switching ELS	100	-	-	-	Deployed
v1_144	v1	Campus Switching ELS	1	-	-	-	Deployed
v100_5	v100	Campus Switching ELS	1000	-	-	-	Deployed

In addition to creating configuration profiles, in Build mode you can configure wireless mobility and network domains, enable Smart Mobile Virtual Controller Clustering, configure access points on controllers, configure Link Aggregation Groups (LAGs) on switches, and so on.

Deploying Device Configurations

After you build your device configurations in Build mode, you need to deploy the configurations on the devices. None of the configurations you create in Build mode affect your devices until the configurations are actually deployed on the devices.

To deploy the configuration on devices, use Deploy mode. When you change a device's configuration in Build mode, the device becomes available in Deploy mode for configuration deployment.

For more information about deploying configuration changes, see ["Deploying Configuration Changes" on page 1173](#).

Importing Device Configurations

As part of device discovery, Network Director analyzes the configuration of a newly discovered device and automatically imports the configuration into the Build mode configuration for that device. For example, as part of the discovery of a wireless LAN controller, Network Director imports the configurations of wireless access points from the controller and makes them available for viewing and modification under the Manage Access Point task for that controller.

As it imports the device configuration, Network Director automatically creates profiles to match the configuration. It first determines whether any existing profiles match the configuration, and if so, assigns those profiles to the device. It then creates and assigns new profiles as needed. For example, if an access switch has some ports that match the configuration of an existing Port profile, Network Director assigns

the existing Port profile to those ports. For the other ports, Network Director creates as many Port profiles as needed to match the port configurations and assigns them to the ports.

You can manage the profiles that Network Director creates as part of device discovery in the same way that you manage user-created profiles—that is, you can modify, delete, or assign them to other devices.

Out-of-Band Configuration Changes

Out-of-band configuration changes are configuration changes made to a device outside of Network Director. Examples include changes made by:

- Using the device CLI.
- Using the device Web-based management interface (the J-Web interface or Web View).
- Using the Junos Space Network Management Platform configuration editor.

NOTE: You cannot use the Junos Space configuration editor to configure wireless LAN controllers.

- Using RingMaster software.
- Restoring or replacing device configuration files.

When an out-of-band change is made, the device configuration no longer matches the Build mode configuration, and the device configuration state changes to out of sync. You cannot deploy configuration on a device that is out of sync. Use the Resynchronize Device Configuration task in Deploy mode to resynchronize the device configuration. For more information about how Network Director resolves out-of-band configuration changes and synchronizes the Build mode configuration with the device configuration, see [“Understanding Resynchronization of Device Configuration” on page 1214](#).

TIP: Before you make configuration changes in Build mode, make sure that devices that will be affected are in sync. Resynchronizing the device configuration can result in losing pending Build mode configuration changes for that device.

Managing Devices

In addition to the tasks that allow you to build your network, Build mode provides a number of tasks for day-to-day device management. For example, you can:

- View a device’s hardware component inventory or its installed licenses
- Reboot a device or groups of devices

- Connect to a device's CLI through SSH or to its web-based management interface
- View the profiles assigned to a device
- Set up QFabric devices.
- View Aruba wireless devices inventory that are connected to Juniper switches.

RELATED DOCUMENTATION

[Understanding the Build Mode Tasks Pane | 189](#)

[Understanding the Network Director User Interface | 85](#)

[Understanding Network Configuration Profiles | 197](#)

[Deploying Configuration Changes | 1173](#)

[Understanding Resynchronization of Device Configuration | 1214](#)

[Network Director Documentation home page](#)

Understanding the Build Mode Tasks Pane

The Tasks pane in Build mode contains all the tasks you can do in Build mode. Click a specific task to begin that task.

The tasks listed in the Tasks pane depend on the scope you select in the View pane—that is, what view (Logical, Location, Device, Virtual, or Custom Group) you have selected and what object you have selected. Not all tasks are available in all scopes. As you change your selections in the View pane, the contents of the Tasks pane also change.

Build mode tasks are divided into the following categories in the Tasks pane.

Network Director enables you to perform the following tasks for devices in your physical network:

- **Device Discovery**—Before your devices can be managed by Network Director, you must use device discovery to discover them. As Network Director discovers devices, it adds them to your network view in the View pane. [Table 29](#) describes the device discovery tasks.
- **Device Management**—After devices have been discovered, you can perform administrative tasks on them, such as viewing a list of the device's physical components, connecting to a device using SSH, rebooting a device, or assigning a switch to its logical role in the network. You can also view the inventory of the Aruba wireless devices connected to the Juniper Network switches and launch the Aruba Airwave application to manage these devices. [Table 30](#) describes the device management tasks.

- **Wired**—You can create configuration profiles and quick templates for the different wired devices—EX Series Ethernet Switches, EX Series switches with ELS, QFX Series switches, MX Series routers, MX Series routers with ELS, QFabric systems.
- **Wireless**—You can create configuration profiles for controllers and access points.
- **Domain Management**—You can create and modify mobility and network domains in your wireless network. [Table 31](#) describes the domain management tasks.
- **Location Management**—You can build your Location view of the network by creating sites, buildings, floors, closets, and outdoor areas and assigning devices to these locations. [Table 32](#) describes the location management tasks.
- **Connectivity**—For switches in your network that are connected to your virtual network, you can view the connectivity between a given switch and the corresponding virtual switch and between the virtual switch and the virtual machine. [Table 33](#) describes the connectivity tasks.
- **Profile and Configuration Management**—Network Director provides a set of configuration profiles that you can create to provision multiple devices in your network. [Table 34](#) describes the profile and configuration management tasks.
- **Key Tasks**—Network Director enables you to group the tasks that you perform frequently and create a list of key tasks. You can add any task from the Tasks pane to the Key Tasks list by selecting a task and clicking the plus (+) sign that appears adjacent to the task. For some modes, you can see that Network Director has predefined some key tasks for you. You can modify this set of tasks to suit your requirements. This feature is available in Task pane irrespective of your current mode, scope, or view.

For more information about Build mode features, see [“Understanding Build Mode in Network Director” on page 184](#).

[Table 29](#) through [Table 34](#) describe the tasks that you can perform in the physical network category, including the scope in the View pane that you must select to access the task.

Table 29: Device Discovery Tasks

Task	Description	Scope
Discover Devices	Discovers supported switches and wireless LAN controllers in the network and brings them under Network Director management. Access points are discovered when their associated controllers are discovered.	Any
View Discovery Status	Displays the status of device discovery jobs.	Any

Table 30: Device Management Tasks

Task	Description	Scope
Assign Device to Logical Category	Assigns switches to one of the following logical categories within the switching network: Core, Aggregation, Access.	View: Logical Object: Switching Network, Core, Aggregation, Access, Unassigned, or an individual switch
Change Location of Device	Changes where a device is located in Location view.	View: All Object: Individual switch, controller, or access point
Create MC-LAG	Configures multichassis link aggregation groups (MC-LAG).	View: All Object: All
Create ESI-LAG	Configures Ethernet Segment Identifier link aggregation groups (ESI-LAG) in a campus environment.	View: All Object: All
Delete Devices	Deletes a switch or a wireless LAN controller as a managed device from Network Director. If you select a scope that contains more than one switch or controller, you can choose which devices are deleted.	View: All Object: All, except access points
Launch RingMaster	Launches Juniper Networks RingMaster, using the URL set under Preferences.	View: All Object: Wireless Network, wireless domains and clusters, wireless devices
Launch Web View	Launches the Web-based management interface for the selected device in a separate window: the J-Web interface for EX Series switches and Web View for wireless LAN controllers.	View: All Object: Individual switch or controller
Manage Access Point	Adds an access point to the selected wireless LAN controller or cluster. Also, deletes an access point from the selected controller or cluster or changes the access point configuration.	View: Logical Object: Controller or cluster
Manage LAG	Creates and manages Link Aggregation Groups (LAGs).	View: All Object: Individual switch

Table 30: Device Management Tasks (continued)

Task	Description	Scope
Manage Port Groups	Enables you to create and manage port groups. Port groups enable you to configure a group of ports with a single configuration.	View: Logical Object: Switching Network, Core, Aggregation, Access, Unassigned, an individual switch (both EX Series and QFX Series)
Manage Virtual Chassis Fabric (VCF)	Manages a Virtual Chassis Fabric (VCF)	View: All Object: VCF
Reboot Devices	Reboots devices. If you select a scope that contains more than one switch or controller, you can choose which devices get deleted.	View: All Object: All
Setup QFabric	Sets up a QFabric device.	View: All Object: QFabric device
Setup Virtual Chassis Fabric (VCF)	Sets up a Setup Virtual Chassis Fabric (VCF).	View: All Object: All
Show Current Configuration	Shows the running configuration on a switch or a wireless LAN controller.	View: All Object: Individual switch or controller
SSH to Device	Launches an SSH connection to the selected device.	View: All Object: Individual switch or controller
Validate Pending Configuration	Validates configuration changes that have not yet been deployed on devices.	View: All Object: All
View Assigned Profiles	Displays the profiles assigned to the selected device.	View: All Object: Individual switch, controller, or radio
View Inventory	Displays information about all the devices in the currently selected object and all its child objects.	View: All Object: All
View License Information	View the licenses installed on the device and their status.	View: All Object: Individual switch or controller
View Physical Inventory	Displays information about the selected device's hardware components.	View: All Object: Individual switch or controller

Table 30: Device Management Tasks (continued)

Task	Description	Scope
View Aruba Wireless Device Inventory	Displays the list of Aruba devices connected to the Juniper Network switches.	View: Logical, Location, Device, and Custom Group Object: Aruba devices
Launch Aruba Airwave	Launches the Home > Overview page of the Aruba Airwave application to manage Aruba devices.	View: Logical, Location, Device, and Custom Group Object: Aruba devices

Table 31: Domain Management Tasks

Task	Description	Scope
Create Mobility Domain	Creates a mobility domain within the wireless network and optionally enables clustering.	View: Logical Object: My Network or Wireless Network only
Create Network Domain	Creates a network domain within the wireless network.	View: Logical Object: My Network or Wireless Network only
Delete Mobility Domain Edit Mobility Domain	Modifies or deletes a mobility domain.	View: Logical Object: A mobility domain
Delete Network Domain Edit Network Domain	Modifies or deletes a network domain.	View: Logical Object: A network domain node

Table 32: Location Management Tasks

Task	Description	Scope
Add Building	Creates a new building in the selected site. NOTE: Use this task only to create the building. Floors and closets in the building must be created separately.	View: Location Object: A site
Add Closet	Creates a new closet in the selected floor.	View: Location Object: A floor
Add Floor	Creates a new floor in the selected building. NOTE: Use this task only to create the floor. Closets in the building must be created separately.	View: Location Object: A building

Table 32: Location Management Tasks (continued)

Task	Description	Scope
Add Outdoor Area	Creates a new outdoor area in the selected site.	View: Location Object: A site
Add Site	Creates a new site in Location view. NOTE: Use this task only to create the site object. Buildings, floors, closets, and outdoor areas in the site must be created separately.	View: Location Object: My Network only
Delete Building/Edit Building	Deletes or modifies the selected building.	View: Location Object: A building
Delete Closet/Edit Closet	Deletes or modifies the selected closet.	View: Location Object: A closet
Delete Floor/Edit Floor	Deletes or modifies the selected floor.	View: Location Object: A floor
Delete Outdoor Area/Edit Outdoor Area	Deletes or modifies the selected outdoor area.	View: Location Object: An outdoor area
Delete Site/Edit Site	Deletes or modifies the selected site.	View: Location Object: A site
Assign Devices to Building	Assigns switches or wireless LAN controllers to a building. You cannot assign access points to a building.	View: Location Object: A building
Assign Devices to Closet	Assigns switches or wireless LAN controllers to a closet. You cannot assign access points to a closet.	View: Location Object: A closet
Assign Devices to Floor	Assigns access points, switches, or wireless LAN controllers to a floor.	View: Location Object: A floor
Assign Devices to Outdoor Area	Assigns access points to an outdoor area.	View: Location Object: An outdoor area
Setup Locations	Opens the page by using which you can create an entire site—that is, define buildings, floors, closets, outdoor areas and to assign devices to these locations. NOTE: Use this task only to create a site. Do not use it to modify an existing site.	View: Location Object: My Network and any location node within an existing site.

Table 33: Connectivity Tasks

Task	Description	Scope
View Virtual Network Connectivity	<p>Pictorially displays the network connectivity between the selected switch and the virtual switch and between the virtual switch and the virtual machine.</p> <p>If the selected switch is not connected to a virtual network, Network Director displays the standalone switch.</p>	View: Logical, Location, Device Object: Individual switch
View Virtual Machines	Displays the virtual machines that are connected to the selected switch.	View: Logical, Location, Device Object: Individual switch

Table 34: Profile and Configuration Management Tasks

Task	Description	Device Family	Scope
Manage Quick Templates	Enables you to create and manage quick templates. Quick templates enable you to define your network configuration in the form of templates that you can apply to multiple devices in your network.	EX Series QFX Series MX Series	All, except wireless devices
View Deployed Templates	Enables you to view the list of quick templates that are deployed.	EX Series QFX Series MX Series	All, except wireless devices
Access	Creates and manages Access profiles. Use Access profiles to configure authentication methods (RADIUS, LDAP, and local), accounting methods (RADIUS and LDAP), and authentication/accounting servers.	EX Series QFX Series WLC Series	Any
Authentication	Creates and manages Authentication profiles. Use Authentication profiles to specify authentication method and authentication parameters for authenticating clients and users who connect to a WLAN or to an access port on a switch.	EX Series QFX Series WLC Series	Any
Authorization	Creates and manages authorization profiles. Use Authorization profiles to configure authorization attributes for users, such as the VLAN to place the users in, firewall filters to apply to the user traffic, time-of-day access restrictions, and so forth.	WLC Series	Any

Table 34: Profile and Configuration Management Tasks (continued)

Task	Description	Device Family	Scope
Auto AP	Creates and manages Auto AP profiles. Use Auto AP profiles to define the configuration of distributed access points that are not explicitly configured on the controller.	WLC Series	Any
CoS	Creates and manages CoS profiles. Use CoS profiles to configure class-of-service (CoS) attributes to be applied to interfaces or to user traffic.	EX Series QFX Series WLC Series	Any
Device Common Settings	Creates and manages Device Common Settings profiles. Use Device Common Settings profiles to configure basic system settings, such as users, time and time servers, SNMP, system logging, and so on.	EX Series QFX Series WLC Series	Any
Fabric	Creates and manages Fabric profiles. Use Fabric profiles to configure gateway FC fabrics on QFX Series devices that act as a FCoE-FC gateway.	QFX Series	Any
FC Gateway Service	Provides a quick way to configure Fibre Channel (FC) gateways on Data Center Switching devices.	QFX Series	Any
Filter	Creates and manages Filter profiles. Use Filter profiles to define Layer 2 and Layer 3 firewall filters (ACLs).	EX Series QFX Series WLC Series	Any
LDAP	Creates and manages LDAP profiles. Use LDAP profiles to specify details about your LDAP authentication and accounting server. LDAP profiles can then be linked to an access profile.	EX Series WLC Series	Any
Port	Creates and manages Port profiles for EX Series switches. Use Port profiles to configure interface settings, such as PoE settings, protocol family, port mode, physical link settings, firewall filters, and port security settings for interfaces.	EX Series QFX Series	Any
Radio	Creates and manages Radio profiles. Use Radio profiles to specify radio behavior that is in common across multiple radios, including the SSIDs advertised by the radio.	WLC Series	Any

Table 34: Profile and Configuration Management Tasks (*continued*)

Task	Description	Device Family	Scope
Radius	Creates and manages Radius profiles. Use Radius profiles to specify details about your RADIUS authentication and accounting server. Radius profiles can then be linked to an access profile.	EX Series WLC Series	Any
VLANs	Creates and manages VLAN profiles. Use VLAN profiles to define VLANs, including the firewall filters to be applied to the VLANs and other settings.	EX Series QFX Series WLC Series	Any
WLAN Service	Creates and manages WLAN Service profiles. Use WLAN Service profiles to create SSIDs and their attributes, such as beaconing, encryption, and default authorization for users accessing the SSID.	WLC Series	Any

RELATED DOCUMENTATION

[Understanding Build Mode in Network Director | 184](#)

[Understanding Network Configuration Profiles | 197](#)

[Understanding the Network Director User Interface | 85](#)

[Understanding the Build Mode Tasks Pane for Datacenter View | 788](#)

[Understanding Quick Templates | 282](#)

[Network Director Documentation home page](#)

Understanding Network Configuration Profiles

To support rapid network deployment, Junos Space Network Director enables you to define your network configuration in a set of profiles that you can apply to multiple objects in your network. For example, you can define a Port profile to set up class-of-service (CoS), authentication, firewall filters, and Ethernet switching settings that are appropriate for all access ports in your network that connect to employee desktop VoIP phones.

After you have defined a profile, you can associate it with devices in one of two ways:

- By directly assigning it to a device (or to ports or radios on the device). When you assign a profile to a device, you can configure certain device-specific parameters. For example, when you assign a VLAN profile to a device, you can configure the IP address for that VLAN on that device. Or when you assign a Port profile for a Layer 3 interface to the interface, you can configure the IP address for that interface.
- By referencing the profile in another profile. Some profiles are not assigned directly to network devices—instead they are referenced from other profiles that are, in turn, assigned to network devices. For example, the settings in the CoS, Filter, and Authentication profiles are assigned indirectly to a port by the profiles being included in the Port profile.

Because a child profile might be a required setting in its parent profile, you must create the child profiles before you create the parent profiles. For example, to create a Radio profile, create the profiles in this order:

1. Access, VLAN, CoS, and Filter profiles
2. Authentication and Authorization profiles
3. WLAN Service profile
4. Radio profile

Network Director also includes six predefined Port profiles and one predefined CoS profile for EX Series switches. You can choose to apply the Port profiles to one or more ports of a single device or a group of devices, and the CoS profile to a Port profile or a Radio profile (using a WLAN profile and an Authorization profile).

After you have created and included the child profiles in to parent profiles, you can assign these parent profiles at various levels in your wired and wireless networks. [Table 35](#) shows the levels at which you can assign each of these parent profiles.

Table 35: Profile Associations at Various Levels

Name of the Profile	Wireless Network (WLC)	EX Series Ethernet Switches (With or Without ELS)	QFX Series Switches (With or Without ELS) and QFabric System
Device Common Settings profile	Device	Device	Device
Radio profile	Radio	Not applicable	Not applicable
VLAN profile	Device, Port	Device	Device
Authorization profile	Not applicable	Not applicable	Not applicable

Table 35: Profile Associations at Various Levels (continued)

Name of the Profile	Wireless Network (WLC)	EX Series Ethernet Switches (With or Without ELS)	QFX Series Switches (With or Without ELS) and QFabric System
Port profile	Not applicable	Port	Port
Fabric profile	Not applicable	Not applicable	Port
Auto AP profile	Device, Cluster	Not applicable	Not applicable
FC Gateway Service	Not applicable	Not applicable	Port
Wireless Filter profile	Device, Cluster	Not applicable	Not applicable
Wireless CoS profile	Device, Cluster	Not applicable	Not applicable
Local Switching VLAN profile	Access Point	Not applicable	Not applicable
mDNS profile	Access Point	Not applicable	Not applicable
Remote Sites profile	Access Point	Not applicable	Not applicable
RF Snoop profile	Radio	Not applicable	Not applicable
RF Detection profile	Device, Cluster	Not applicable	Not applicable

Once you have assigned profiles to devices, radios, ports, or access points, you can view the profile associations in the **Profiles Assigned to the Device** page. For more information, see [“Viewing Profiles Assigned to a Device” on page 1144](#).

In addition to the profiles you create yourself, Network Director creates profiles for you from existing device configuration. Typically, you create profiles and associations manually when you are setting up a new network from scratch, adding a new device to your existing network, or when you want to make certain customized changes to the way your network is currently operating. Network Director creates profiles for you when:

- You discover existing devices in your network. As part of device discovery, Network Director examines the configurations present in the discovered device or devices. If configurations match existing profiles, Network Director assigns the matching profiles to the appropriate levels on the devices. If configurations do not match existing profiles, Network Director creates the required profiles and associates them at the appropriate levels. For more information about device discovery, see [“Discovering Devices in a Physical Network” on page 204](#).

NOTE: If Network Director fails to read the configuration of one or more devices after the device discovery, such devices are not displayed in the Assign Profile page. You will not be able to assign profiles to such devices. The Manage Jobs page in System mode displays details of the device discovery jobs. Use the information displayed on this page to take appropriate corrective steps to enable Network Director to reread the configuration of the failed device. For more information, see [“Troubleshooting Device Discovery Error Messages” on page 217](#).

- You first install Network Director and supported devices are already being managed by Junos Space. In this case, Network Director imports the device configurations into profiles the same way it does when you discover devices with Network Director.
- You resynchronize the Network Director configuration with the device configuration in order to resolve out-of-band configuration changes—that is, configuration changes that are not made with Network Director. Out-of-band configuration changes result in the device configuration not matching or being in sync with the Network Director configuration for the device. When you resynchronize the Network Director configuration with the device configuration, Network Director creates and associates new profiles if none of the existing profiles match the changed configuration. For more information about resynchronization of device configuration, see [“Understanding Resynchronization of Device Configuration” on page 1214](#).

After a profile is created, you can edit it from the Manage Profile page by selecting the profile and clicking Edit. The only exception is when the profile that you want to edit is part of a job that is scheduled for deployment. When you schedule a deployment job, that job and any profiles assigned to that job are locked. You cannot edit the job or any of its assigned profiles until the job is completed or gets cancelled. For more information, see [“Deploying Configuration Changes” on page 1173](#).

When you delete a device in Network Director, the system deletes only the device and the profile associations. The profiles are retained in the system. If you rediscover the deleted devices into the system at a later stage, without making any configuration changes on the device, Network Director identifies this and reinstates the previous profile associations.

RELATED DOCUMENTATION

[Understanding Access Profiles | 351](#)

[Understanding Authentication Profiles | 381](#)

[Understanding Wireless Authorization Profiles | 395](#)

[Understanding Auto AP Profiles | 883](#)

[Understanding Class of Service \(CoS\) Profiles | 609](#)


[Understanding Device Common Settings Profiles | 290](#)

Understanding Fabric Profiles 693
Understanding Filter Profiles 540
Understanding Port Profiles 408
Understanding Radio Profiles 879
Understanding VLAN Profiles 499
Understanding WLAN Service Profiles 885
Creating and Managing FC Gateway Service Profiles 684
Network Director Documentation home page

Assigning Profiles to an Interface, Device, or a Group of Devices

After you create an authorization profile, CoS profile, device common settings profile, fabric profile, FC Gateway services profile, filter profile, port profile, radio profile, VLAN profile, Auto AP profile, Local Switching profile, and RF Snooping Filter profile, you can assign each of these profile to an interface, device, or a group of devices.

To assign a profile:

1. Click  **Build** in the Network Director banner.
2. Under Select View, select either **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View** or **Topology View**.

3. From the Tasks pane, select the type of network (Wired or Wireless), the appropriate functional area (System, AAA, or Wireless), and select the name of the profile that you want to create. For example, to create a radius profile for a wireless device, click **Wireless > AAA > Radius**. The Manage Profile page opens.
4. Select the profile that you want to assign and click **Assign**.

The Assign Authorization Profile page appears displaying a hierarchal list of network objects, including the network, mobility domain, controllers and clusters that are already defined or discovered for your network.

5. Select a level and click **Next** to view the objects available at that level. If you select a network or a mobility domain, all the controllers or clusters that are part of that network or mobility domain are also selected.
6. Select one or more devices, groups, or clusters from the list.

NOTE: If Network Director fails to read the configuration of one or more devices after the device discovery, such devices are not displayed in the Assign Profile page. You will not be able to assign profiles to such devices. The Manage Jobs page in System mode displays details of the device discovery jobs. Use the information displayed on this page to take appropriate corrective steps to enable Network Director to reread the configuration of the failed device. For more information, see [“Discovering Devices in a Physical Network” on page 204](#).

7. If you want to view the existing assignment of a device or a cluster, select it and click **View Assignments**.
The Profile Details window opens displaying the device’s current profile assignment.
Click **Close** to close the window.
8. If you want to remove an existing assignment from a device or a cluster, select it and click **Remove**.
The system removes the assignment from the selected device or cluster.
9. Do one of the following to assign the Authorization profile to a device or a cluster:
 - Click **Assign > Assign to Device** to assign the Authorization profile to the selected devices.
 - Click **Assign > Assign to Cluster** to assign the Authorization profile to the selected clusters.
10. Click **Next** or **Review**.
The system displays the associations that you created. To modify any of these assignments, click **Edit** or **Profile Association**.
11. Click **Finish** to save the profile associations.
After you click Finish, the Create Profile Assignments Job Details window opens with a report on the status of the profile assignment job. If you have assigned the profile to a large number of objects, the profile assignment job can take some time to complete. Instead of waiting for the Job Details window to report job completion status, you can close it and check the details of the profile assignment job at a later time using the Manage Job task in System mode.

NOTE: If any assignment fails, the profile assignment job fails and none of the assignments are created. Check the details for the profile assignment job for information about why the assignment failed.

RELATED DOCUMENTATION

[Understanding Filter Profiles | 540](#)

[Understanding Class of Service \(CoS\) Profiles | 609](#)

[Network Director Documentation home page](#)

Discovering Devices

IN THIS CHAPTER

- Discovering Devices in a Physical Network | 204
- Discovering Devices in a Datacenter Network | 212
- Understanding the Device Discovery Process | 216
- Troubleshooting Device Discovery Error Messages | 217

Discovering Devices in a Physical Network

IN THIS SECTION

- Specifying Target Devices | 205
- Specifying Discovery Options | 208
- Specifying Schedule Options | 209
- Reviewing Device Discovery Options | 210
- Viewing the Discovery Status | 210

You can discover and synchronize physical devices such as EX Series switches, QFX Series devices, MX Series routers, and wireless LAN controllers in your network that are managed by Network Director.

You can also discover and manage virtual devices in your virtual network from Network Director. To discover devices in your virtual network, follow the steps described in [“Discovering Devices in a Datacenter Network”](#) on page 212.

NOTE: To discover wireless LAN controllers and access points, Network Director connects to port 8889 of the wireless LAN controller for sending an *HTTPS* request to the controller. Therefore, port 8889 must be open on the wireless LAN controller for Network Director to discover wireless LAN devices. From the Junos Space JA2500 Appliance or the Junos Space Virtual Appliance, you can verify the availability of port 8889 using the CLI command `root@space# nmap <IP address of controller> -p 8889`.

On EX Series switches, Network Director connects to port 22 (the default port) on the Junos Space JA2500 Appliance or the Junos Space Virtual Appliance by using SSH. You can configure port 22 on the Junos Space appliances through **Administration > Applications** on the Junos Space Platform page. Select **Network Application Platform** and click **Actions > Modify Application Settings**. Change the SSH port for device connection to 22.

Device discovery is a three-step process in which you specify the target devices, the discovery options, and the schedule options.

While in Build mode, from the Tasks pane, click **Discover Devices** from the Device Discovery menu. The Discover Devices page is displayed.

This topic describes:

Specifying Target Devices

You can add devices to Network Director for device discovery by clicking either **Import from CSV** or **Add**, or both together. Click **Import from CSV** to add devices in bulk. You can add a large number of devices to Network Director by using a CSV file that contains information extracted from an LDAP repository. During device discovery, you can associate the devices with logical, location, and custom groups. You can list all devices to be discovered in the CSV file along with their logical, location, and custom groups. This eliminates the need to make an explicit association later. If you do not assign groups to the devices, the devices are added to the Unassigned folder by default. You can also change the assignment later. Associating new devices with groups makes the network simpler to manage and maintain.

To specify the target devices that you want Network Director to discover:

1. Enter a name for the device discovery job.

The default name is ND Discovery.

2. To add devices in bulk, click **Import from CSV** from the Device Targets window.

The Upload CSV File dialog box is displayed.

3. Click **Browse**.

The File Upload dialog box is displayed.

4. Navigate to the target CSV file on your computer, select the file, and click **Open**.

The CSV File Upload dialog box reappears, this time displaying the name of the selected file.

NOTE: The selected CSV file must follow the same file format as that of the sample CSV file.

5. Click **Upload** to upload the selected CSV file.
6. To add individual devices by specifying the IP address credentials, click **Add** in the Device Targets table.
The Add Device Target dialog box appears.

7. In the Add Device Target dialog box, perform the following steps:

- a. Choose one of the following options to specify target devices:

- Select the **IP** option and enter the IP address of the device.
- Select the **IP-Range** option and enter a range of IP addresses for the devices.

The maximum number of IP addresses for an IP range target is 1024.

- Select the **IP-Subnet** option and enter an IP subnet for the devices.
- Select the **HostName** option and enter the hostname of the device.

- b. In the Assign To section, specify the following groups to which the newly discovered devices can be assigned:

- From the Logical Group drop-down menu, select **Core**, **Aggregation**, **Access**, or **Layer 3 Fabric** to specify the logical grouping of the device.

Select the Layer 3 Fabric option to discover a Layer 3 fabric that is not created using Network Director and OpenClos. You can discover devices in Network Director that belong to the same IP subnet. To discover a Layer 3 Fabric, specify the IP subnet range, as all the devices that belong to the same Layer 3 Fabric resides in the same subnet. Network Director expands the IP subnet range and reaches every single IP address that you have specified in the IP range.

Network Director initially discovers the Layer 3 Fabric based on the IP subnet and range. However, you can manually discover Layer 3 Fabric at a later stage by entering the host name of the Layer 3 Fabric.

- For the Location Group field, click **Select** to choose the location group for the device or input the location path for the association. To clear the selection, click **Clear**.

Use the following format for the location path for the respective associations:

- *site-name#S/building-name#B*
- *site-name#S/building-name#B/floor-name#F#floor-level*
- *site-name#S/building-name#B/floor-name#F#1 - 1st level*
- *site-name#S/building-name#B/floor-name#F#floor-level/closet-name#C*
- *site-name#S/building-name#B/floor-name#F#floor-level/aisle-name#A/rack-name#R*
- *site-name#S/outdoorarea-name#O*

If the location paths do not point to existing locations in Network Director, new location groups are created before the devices are assigned to them.

- For the Custom Group field, click **Select** to choose the custom group for the device or input the custom group path. To clear the selection, click **Clear**.

Use the following format for the custom group path:

- *customgroup1-name/customgroup2-name*

If the custom group paths do not point to existing custom groups, new groups are created before the devices are assigned to them.

- c. Click **Add** to save the target devices that you specified, or click **Add More** to add more target devices. When you have added all target devices that you want Network Director to discover, click **Add**.

The Discover Targets table displays the addresses of the configured target devices.

8. Following device discovery management options are available:
 - To edit a target device, select a row from the Device Targets table and click **Edit**. Make the required changes and click **Add** to display the IP addresses in the Device Targets table
 - To delete a target device, select a row from the Device Targets table and click **Delete**.
 - To view and download a sample CSV file, click **CSV Sample**. The Opening Device_Discovery_CSV.csv file dialog box is displayed. You can open the sample CSV file or save the sample CSV file.
9. Click **Next** or click **Discovery Options** from the top wizard workflow to go to the Discovery Options page. Specify the options as described in [“Specifying Discovery Options” on page 208](#).

Specifying Discovery Options

To add the device credentials and specify the probes:

1. Add the device credentials. To add the credentials, click **Add** from the Device Credentials table.

The Add Device Credentials dialog box is displayed.

NOTE: If the credentials were specified in the CSV file, the Credentials table displays those values. If the credentials were not specified in the CSV file, then enter the values in the Add Device Credentials dialog box.

- Specify the administrator username and password, and confirm the password. The username and password must match the name and password configured on the device. The username is a mandatory field.
- Click **Add** to save the username and password that you specified or click **Add More** to add another username and password.

Click **Add** after you have finished adding all login credentials. The Device Credentials table displays the usernames that you configured.

2. Specify the probes from the Specify Probes table. Select a probe method to discover the target devices.

- Select **Use Ping** if SNMP is not configured for the device and clear the **Use SNMP** check box.

Network Director uses the Juniper Networks Device Management Interface (DMI) to directly connect to and discover devices. DMI is an extension to the NETCONF network management protocol.

- Select **Use SNMP** if SNMP is configured for the device, and clear the **Use Ping** check box.

Network Director uses the **SNMP GET** command to discover target devices.

- Select both the **Use Ping** and the **Use SNMP** check boxes, to enable Network Director for faster discovery of the target devices, provided the device is pingable and also SNMP is enabled on the device.

NOTE: Network Director uses the Juniper Networks Device Management Interface (DMI) adapter to manage devices that do not run Junos OS. Wireless LAN controllers are not DMI-complaint and use the MSS software. However, if you enable Use SNMP, Network Director detects whether the device is running a DMI-complaint software or not. The controllers always use a DMI adapter and hence can be detected.

If you have not enabled Use SNMP, then Network Director assumes that all devices that failed during device discovery are controllers and retry the process if the port 8889 is open.

3. Click **Add** if you have selected the Use SNMP check box.

The Add SNMP Settings dialog box is displayed.

Select either **SNMP V1/V2C** or **SNMP V3**. Based on the selection, you need to enter the details as follows:

- If you selected SNMP V1/V2C, specify a community string, which can be *public*, *private*, or a predefined string.

Click **Add** in the Add SNMP Settings dialog box or click **Add More** to add more strings to the community. If you click **Add More**, when you are done adding all the strings, click **Add** to save the SNMP settings for V1/V2C.

- If you selected SNMP V3:
 - Enter a username
 - Select the privacy type (AES 128, DES, or None).
 - Enter the privacy password (if AES 128 or DES). If you specify none for the privacy type, the privacy function is disabled.
 - Select the authentication type (MD5, SHA, or none).
 - Enter the authentication password (if MD5 or SHA). If you specify none for the authentication type, the authentication function is disabled.
 - Click **Add** to save the SNMP settings and close the dialog box, or click **Add More** to add additional configurations. If you clicked Add More, click **Add** to save the settings and close the dialog box.

The Specify Probes table displays the configured SNMP settings.

4. Click **Next** or click **Schedule Options** from the top wizard workflow to go to the Discovery Schedule Options page. Specify the options as described in the [“Specifying Schedule Options” on page 209](#).

Specifying Schedule Options

To specify the scheduler details:

1. Click **Run Now** if you want to discover the devices immediately or Click **Schedule at a later time** if you want to schedule the device discovery for a future time.

If you select **Schedule at a later time**, specify the date and time to run the device discovery.

2. Click **Next** or click **Review** from the top wizard workflow to view the configuration. See [“Reviewing Device Discovery Options” on page 210](#).

Reviewing Device Discovery Options

From this page, you can save or make changes to the device discovery options.

- To make changes to the device discovery options, click the **Edit** button associated with the configuration you want to change.

Alternatively, you can click the appropriate buttons in the profile workflow at the top of the page that corresponds to the configuration you want to change.

When you are finished with your modifications, click **Review** to return to this page.

- Click **Finish** when you are done with the configurations.

A message window opens, displaying the status of the device discovery job name and job ID. Click **OK**.

The Device Discovery Jobs page is displayed with the list of jobs scheduled.

Viewing the Discovery Status

After you have configured the device discovery options, you can view the device discovery status from the **View Discovery Status** option from the **Device Discovery** menu.

The **Device Discovery Jobs** page displays all the scheduled device discovery jobs. You can view the following details from the Device Discovery Jobs page as described in [Table 36](#).

Table 36: Viewing Device Discover Jobs

Field	Description
Job ID	An identifier assigned to the job.
Job Name	The name of the job (user-created).
Percent	Percentage of the job that is complete.
Status	<p>Job status. The possible states are:</p> <ul style="list-style-type: none"> • CANCELLED—The job was cancelled by a user. • FAILURE—The job failed. This state is displayed if any of the devices in the job failed. But some of the devices might have completed successfully. View the job details for the status of each device. • INPROGRESS—The job is running. • SCHEDULED—The job is scheduled but has not run yet. • SUCCESS—The job completed successfully. This state is displayed if all of the devices in the job completed successfully.
Summary	Summary of the job scheduled and executed with status.

Table 36: Viewing Device Discover Jobs (continued)

Field	Description
Scheduled Start Time	The UTC time on the client computer when the job is scheduled to start.
Actual Start Time	The actual time when the job started.
End Time	The time when the job was completed.
User	The login ID of the user that initiated the job.
Recurrence	The recurrent time when the job will be restarted.

To view the details of a job, select the check box against Job ID or Job Name and click **Show Details**. The Discover Network Elements window displays details of the device discovery job.

NOTE: During device discovery, if Network Director is unable to read the device configurations, then the status displays Failed state. For such failures, you can check the reason for failure from the Manage Jobs page in System mode. You must make the required changes to the device configuration using the CLI so that Network Director can read the configuration. Network Director automatically resynchronizes once you enable a device discovery job. If Network Director cannot discover the device even after resynchronization, then you must rediscover the device after making the appropriate changes in the device configurations by using the CLI.

RELATED DOCUMENTATION

[Viewing the Device Inventory Page | 1136](#)

[Troubleshooting Device Discovery Error Messages | 217](#)

[Network Director Documentation home page](#)

Discovering Devices in a Datacenter Network

IN THIS SECTION

- [Specifying the vCenter Targets | 213](#)
- [Specifying the Virtual Network Credentials | 213](#)
- [Reviewing the Virtual Network Discovery Options | 214](#)
- [Viewing the Discovery Status | 214](#)

You use virtual network discovery to add virtual networks to Network Director. Discovery is the process of finding and adding a virtual network (VMWare vCenter server) and then synchronizing its inventory and configuration with Network Director. To use virtual network discovery, Network Director must be able to connect to the vCenter server.

When discovery succeeds, Network Director creates an object in its database to represent the vCenter server and maintains a connection between the object and the vCenter server to maintain the information flow.

When configuration changes are made on the vCenter server, Network Director automatically resynchronizes with the vCenter server so that the inventory information in the Junos Space database matches the current vCenter server inventory and the configuration information. You can also manually synchronize the vCenter server configuration changes to the configuration stored in the corresponding Network Director object.

At a high level, the following vCenter server inventory and configuration data is captured and stored in relational tables in the Junos Space database:

- Virtual Switches: Port groups, interfaces, VLANs, and other configuration details such as QoS parameters, teaming, and failover details
- Hosts: Physical NICs and virtual machines
- Virtual Machines: Virtual adapters and their associations

Virtual network discovery is a three-step process in which you specify the target devices, the credentials of devices to be discovered, and review the discovery options.

While in Build mode with Datacenter View selected, from the Tasks pane, click **Discover Virtual Network** from Virtual Network Management menu. The Discover Virtual Network page is displayed.

This topic describes:

Specifying the vCenter Targets

You can add VMware vCenter servers to Network director by specifying the hostname or the IP address of one or more vCenters.

To specify the target devices that you want Network Director to discover:

1. In the vNetwork Targets and Credential tab, enter a name for the virtual network discovery job. The default name is vNetwork Discovery.
2. Click **Add** in the vNetwork Targets table.
The Add vNetwork Target dialog box appears.
3. Choose one of the following options to specify the target vCenter server that you want to discover:
 - Select the **IP** option and enter the IP address of the vCenter server.
 - Select the **HostName** option and enter the hostname of the vCenter server.
4. Specify the port that Network Director uses to connect to the vCenter server. The default port is 443.

NOTE: You can modify this and specify a port of your choice. If you do so, make sure to manually change the Junos Space firewall settings and apply to this port.

5. Click **Add** to save the target vCenter server that you specified, or click **Add More** to add more target vCenter servers. When you have added all target devices that you want Network Director to discover, click **Add**.

The vNetwork Targets table displays the addresses of the configured target devices.

6.
 - To edit a target vCenter server, select a row from the vNetwork Targets table and click **Edit**. Make the required changes and click **Add** to display the changes in the vNetwork Targets table.
 - To delete a target vCenter server, select a row from the vNetwork Targets table and click **Delete**.
7. Click **Next** or **vNetwork Credentials** from the top wizard workflow to go to the vNetwork Credentials page. Specify the options as described in the [“Specifying the Virtual Network Credentials” on page 213](#).

Specifying the Virtual Network Credentials

1. Add the vCenter server credentials. To add the credentials, click **Add** from the Device Credentials table.

The Add vNetwork Credentials dialog box is displayed.

- Specify the administrator username and password, and confirm the password. The username and password must match the name and password configured on the device. The username is a mandatory field.
- Click **Add** to save the username and password that you specified or click **Add More** to add another username and password.

Click **Add** after you have finished adding all login credentials. The Device Credentials table displays the usernames that you configured.

2. Click **Next** or click **Review** from the top wizard workflow to go to the Review page.

Reviewing the Virtual Network Discovery Options

From the review page, you can save or make changes to the virtual network discovery options.

- To make changes to the virtual network discovery options, click the **Edit** button associated with the configuration you want to change.

Alternatively, click the appropriate buttons in the profile workflow at the top of the page that corresponds to the configuration you want to change.

When you are finished with your modifications, click **Review** to return to this page.

- Click **Finish** when you are done with the configurations.

A message window opens, displaying the status of the virtual network discovery job name and job ID. Click **OK**.

You can view the status of the discovery job in the Virtual Network Discovery Jobs page.

Viewing the Discovery Status

After you have configured the virtual network discovery options, you can view the discovery status from the **View Discovery Status** option from the **virtual network discovery** menu.

The **Virtual Network Discovery Jobs** page displays all the discovery jobs. You can view the following details from the Virtual Network Discovery Jobs page as described in [Table 37](#).

Table 37: Viewing Device Discover Jobs

Field	Description
Job ID	An identifier assigned to the job.
Job Name	The name of the job (user-created).

Table 37: Viewing Device Discover Jobs (continued)

Field	Description
Percent	Percentage of the job that is complete.
Status	<p>Job status. The possible states are:</p> <ul style="list-style-type: none"> ● CANCELLED—The job was cancelled by a user. ● FAILURE—The job failed. This state is displayed if any of the devices in the job failed. But some of the devices might have completed successfully. View the job details for the status of each device. ● INPROGRESS—The job is running. ● SCHEDULED—The job is scheduled but has not run yet. ● SUCCESS—The job completed successfully. This state is displayed if all of the devices in the job completed successfully.
Summary	Summary of the job scheduled and executed with status.
Scheduled Start Time	The UTC time on the client computer when the job is scheduled to start.
Actual Start Time	The actual time when the job started.
End Time	The time when the job was completed.
User	The login ID of the user that initiated the job.
Recurrence	The recurrent time when the job will be restarted.

To view the details of a job, select the check box against Job ID or Job Name and click **Show Details**. The Discover Network Elements window displays details of the virtual network discovery job.

NOTE: During virtual network discovery, if Network Director is unable to read the device configurations, then the status displays Failed state. For such failures, you can check the reason for failure from the Manage Jobs page in System mode. You must make the required changes to the device configuration using the CLI so that Network Director can read the configuration. Network Director automatically resynchronizes once you enable a virtual network discovery job. If Network Director cannot discover the device even after resynchronization, then you must rediscover the device after making the appropriate changes in the device configurations using the CLI.

RELATED DOCUMENTATION

| [Troubleshooting Device Discovery Error Messages](#) | 217

Understanding the Device Discovery Process


When a new device with network configurations is added to the network, Network Director runs a job to discover the device. Two minutes after device discovery, Network Director initiates another job called the brownfield process. The brownfield process ensures that the new device is ready to be used in the network by deploying the required configurations to the device.

To support rapid network deployment, Junos Space Network Director enables you to define your network configuration in a set of profiles that you can apply to multiple objects in your network. For example, you can define a Port profile to set up class-of-service (CoS), authentication, firewall filters, and Ethernet switching settings that are appropriate for all access ports in your network that connect to employee desktop VoIP phones.

You can manually create Profiles from the Network Director user interface or the profiles may be created automatically by Network Director when you discover a device. Once a device, that has network configurations, is discovered, Network Director initiates a Brownfield process to read the configuration and create the necessary profiles for all the supported configuration from the discovered device.

Figure 9 displays how the brownfield process works in Network Director.

Figure 9: Brownfield Process



Using the brownfield process, Network Director completes the following actions:

- Fetching the complete configuration file of the newly discovered device and looking for matching profiles in the Network Director database.
- Using basic configurations from the already existing matching profiles in the database (such as VLAN IDs, ports, authentication protocols, class of service, firewalls, and so on) to deploy on the newly

discovered device. If a matching profile does not exist, Network Director uses the configuration in the newly discovered device to create a new profile that can then be associated with other devices added to the network in the future.

Benefits of the Device Discovery Process

- A newly discovered device becomes functional in a matter of minutes after it is brought into the network because Network Director automatically assigns an existing profile to the device or creates a new profile without manual intervention.
- Device configurations are reused so you do not need to configure basic features (for example, the VLAN ID) for every newly discovered device added to the network.
- Bulk provisioning of profiles on devices means you can change any parameter (for example, VLAN) on the profile to effect the changes on multiple devices simultaneously.

RELATED DOCUMENTATION

| [Understanding Network Configuration Profiles](#) | 197

Troubleshooting Device Discovery Error Messages

While you are discovering devices by using Network Director, you might encounter some issues. Network Director enables you to detect the errors and provide solutions to the potential errors that you encounter.

Error Message	Solution
---------------	----------

Error Messages Displayed During Discovery of Wireless LAN Controllers

Error Message	Solution
<p>User authentication failed.</p>	<p>This message is displayed when SNMP option is not enabled during the device discovery process.</p> <p>Check whether the firewall is disabled on Junos Space Virtual Machine or the JA 1500 Appliance on port 8889 (outbound), if not, disable the firewall. Port 8889 is opened during Network Director installation, but in case of a failure, you must check whether this port is open or not.</p>
<p>Adapter mssos cannot manage the device - Either the Device is down or the port 8889 is not reachable.</p>	<p>This message is displayed when SNMP is enabled during the device discovery process.</p> <p>Check whether port 8889 is open or not. You can check using the CLI command root@space# nmap <IP address of controller> -p 8889 from the Junos Space JA2500 Appliance or Junos Space Virtual Appliance.</p>
<p>Adapter mssos cannot manage the device-Can't manage Device, Wrong Credentials.</p>	<p>Check the credentials entered during the device discovery process. Always select the <i>Enable Password</i> credential for device discovery in Network Director.</p>
<p>Check if adapter is running or Platform can't connect to it.</p>	<p>Check whether the MSSOS adapter process is running on Junos Space JA2500 Appliance or the Junos Space Virtual Appliance. Verify using the CLI command root@space# ps -ef grep mssos.</p> <p>If no process is running, restart the adapter service by using the CLI command root@space# service mssosadapter start.</p>

Error Message	Solution
<p>Device is not reachable.</p>	<p>If ping is enabled during device discovery, then check whether the controller is reachable using the CLI command ping.</p>
<p>Junos Space is unable to query the device information through SNMP. Check the SNMP settings on the device to verify SNMP is not blocked and the SNMP settings specified in Junos Space match the device SNMP settings.</p>	<p>If the SNMP option is enabled in Network Director during device discovery, then check and ensure that SNMP is enabled on the controller. Also, ensure that the SNMP settings on Network Director and Junos Space match with the SNMP settings on the controller.</p>
<p>Error Messages Displayed During Discovery of EX Series Switches</p>	
<p>SSH connection failed. Device might not be reachable.</p>	<p>For EX Series switches, Network Director connects to port 22 (default port) on the JA2500 Appliance or the Junos Space Virtual Appliance by using SSH. Ensure that you have configured port 22 on the Space appliance through Administration > Applications in the Junos Space Platform page. To do this, select Network Application Platform and click Actions > Modify Application Settings. Change SSH port for device connection field to 22.</p> <p>If port 22 is open on the Junos Space Appliance, and you still get the error, then check if port 22 is open on the switch and if the switch is accepting SSH connections on port 22.</p>
<p>User Authentication failed.</p>	<p>Check the read and write credentials used during device discovery.</p>

Error Message	Solution
Device is not reachable.	If ping is enabled during device discovery, then check whether the switch is reachable using the CLI command ping .
Junos Space is unable to query the device information through SNMP. Check the SNMP settings on the device to verify SNMP is not blocked and the SNMP settings specified in Junos Space match the device SNMP settings.	If the SNMP option is enabled in Network Director during device discovery, check and ensure that SNMP is enabled on the switch. Also, check and ensure that the SNMP settings on Network Director and Junos Space match with the SNMP settings on the switch.
General Error Messages	
Device Failed to return System information.	This message is displayed if the switch is too busy to respond to operational commands. Try discovering the device again.
Failed to configure device, Check Device state.	Check whether the Edit lock is open on the switch and close it if it is open. The configuration commit fails if the Edit lock is open.
Device has been added, but failed to synchronize. Please try manual re-synchronization. Error while reading config from device: device_name, Detail - Fail while executing following RPC: <get-configuration database=committed><configuration></configuration></get-configuration>	Try to resynchronize the devices manually. For details, see "Resynchronizing Device Configuration" on page 1220 .
Error while reading config from device: device-name Failed while executing the following RPC: <get-hardware-inventory/>	<p>Check the hardware details of the switch using the CLI command show chassis hardware detail.</p> <p>If the output displays a message error: command is not valid, then the Junos OS image on the specified switch is corrupted and you need to upgrade to the latest version of Junos OS.</p>

RELATED DOCUMENTATION

- [Discovering Devices in a Physical Network | 204](#)
- [Resynchronizing Device Configuration | 1220](#)
- [Network Director Documentation home page](#)

Setting Up Sites and Locations Using the Location View

IN THIS CHAPTER

- Understanding the Location View | 222
- Setting Up the Location View | 224
- Creating a Site | 228
- Configuring Buildings | 229
- Configuring Floors | 231
- Setting Up Closets | 233
- Assigning and Unassigning Devices to a Location | 234
- Changing the Location of a Device | 236
- Deleting Sites, Buildings, Floors, Wiring Closets, and Devices | 238
- Configuring Outdoor Areas | 240

Understanding the Location View

The Location View is one of the perspectives that Network Director enables you to view and analyze your network. Using this view, you can view devices and data based on their physical location and proximity in the network. By physical location, we mean the buildings, floors, aisles, racks, wiring closets, and outdoor areas where the devices reside. After these locations are defined and devices assigned, the Location View gives you a visual representation of your devices based on where they reside.

You can define the physical location where the devices in the network are deployed in a hierarchical way, and define location entities from a site down to the wiring closet. When in the Location View, the network tree shows the network in terms of buildings, floors, aisles, racks, wiring closets, and outdoor areas nested beneath the building. The hierarchy of the locations is:

- Site—Your campus or data center; the highest node in your location.
- Building—One entry for every building at your site. Buildings are listed in alphabetical order, not by address or the order in which you identified them to the system.

- Floors—One entry for each floor within the building; Floors are nested within the building.
- Aisles—One entry for each aisle in a floor. Aisles are nested within the floor.
- Racks—One entry for each rack in an aisle. Racks are nested within the aisle.
- Outdoor Area—One entry for each named area; Outdoor areas are associated with buildings.
- Devices—Most are assigned to buildings, floors, outdoor areas, or racks. Access points can be assigned only to outdoor areas and floors. Devices are not assigned at the site level; those devices are considered unassigned.

The hierarchical model enables you to define a location by using either of these methods:

- Using the Location Setup wizard to set up a location in a single process, starting at the site level and progressing to the racks and outdoor areas. The wizard also provides an option to create part of the location, such as defining the site and building, then to use the individual procedures to create floors and wiring closets for the building you created.
- Using separate tasks to create location entities in sequence in a top-to-bottom order. You can create the higher level entities such as a site or building first and save them. Later, you can add floors and wiring closets when information about them becomes available.

RELATED DOCUMENTATION

[Setting Up the Location View | 224](#)

[Creating a Site | 228](#)

[Network Director Documentation home page](#)

Setting Up the Location View

You can build a new location site by the individual nodes, or you can use the Location Setup page. The wizard guides you through the top-down process from the site node down to the assignment of devices.

NOTE: Use the Location Setup page only to design new sites; it is not meant for editing existing sites. If you enter data for an existing site, it is rejected when you attempt to commit the data.

A site is the cornerstone of the location-based view of your network. Until you define a site, the default view of your network tree only shows you a list of your unassigned devices. After you define a site, you can build a tree structure of buildings, floors, wiring closets, aisles, and outdoor areas. As you define your network, you can assign devices to the various components of your network. [Table 38](#) describes the devices that you can assign to each of the location component.

Table 38: Devices that can be Assigned to each Location Component

Component	Devices that can be assigned
Site	None
Building	EX Series switches, QFX Series switches, QFabric components (Interconnects, Directors, Node devices, Control Plane devices), Wireless Controllers
Floor	EX Series switches, QFX Series switches, QFabric components (Interconnects, Directors, Node devices, Control Plane devices), Wireless Controllers, Access Points
Closet	EX Series switches, QFX Series switches, QFabric components (Interconnects, Directors, Node devices, Control Plane devices), Wireless Controllers
Aisle	None
Rack	EX Series switches, QFX Series switches, QFabric components (Interconnects, Directors, Node devices, Control Plane devices), Wireless Controllers
Outdoor Area	EX Series switches, QFX Series switches, QFabric components (Interconnects, Directors, Node devices, Control Plane devices), Wireless Controllers, Access Points

The Location Setup page displays the network tree as you add components to your network. Use the buttons on this page to add various components—such as, buildings, outdoor areas, floors, aisles, racks—to your network. These buttons change depending on the component that you select in the network tree.

After the location is set up, you can view the devices in the network by expanding and collapsing these location nodes in the Location view.

To set up your Location view:

1. Ensure you are in the Build mode and Location or Topology view. Click **Build** in the Network Director banner to enter Build mode; select **Location** view or **Topology** view from the View selector.
2. If you are accessing the Location Setup page from the Location view, select the root node (for example, My Network) in the View pane.
3. Do one of the following depending on the view you are in:
 - From the Tasks pane in the Location view, select **Location Management > Setup Locations**.
 - From the Tasks pane in the Topology view, select **Location > Setup Locations**.

The Location Setup page opens.

4. Click **Add Site** to add a new site.

Network Director adds a new site under the root node and names it as **Site-1**.

5. Select the new site and perform any of the following actions:

- Click **Edit Site** to modify the name of the site and specify the site address. The Edit Site window opens.

Topology view uses this address to place the devices assigned to this site on the topology map. For more details on editing a site, see [“Creating a Site” on page 228](#).

- Click **Add Building** to add a building to your site.

Network Director adds a new building under the site and names it as **Building-1**.

- Click **Outdoor Area** to add an outdoor area to your site. Network Director adds a new outdoor area under the site and names it as **Outdoor Area-1**. You can associate an outdoor area to a site or a building for wireless coverage and upload an image or map of that area. After you designate an outdoor area, you can edit or view the map using the Edit Outdoor Area task.

- Click **Delete** to delete the site.

6. If you added a building, select the building and perform any of the following actions to continue building your network:

- Click **Add Floor** to add floors to the building.
- Click **Assign Device** to assign devices to the selected building. The Associate Devices to Building window opens displaying all the unassigned devices in your network. Select the devices that you want to add to the building and click **Add**.

Network Director adds the selected devices to the network tree.

- Click **Edit Building** to edit the name and address of the building. For more details on editing a building, see [“Configuring Buildings” on page 229](#).
 - Click **Delete** to delete the building.
7. If you added an outdoor area, select the outdoor area and perform any of the following actions to continue building your network:
- Click **Assign Device** to assign devices to the selected outdoor area. The Associate Devices to Outdoor window opens displaying all the unassigned devices in your network. Select the devices that you want to add to the building and click **Add**.
 - Click **Edit Outdoor Area** to edit the name of the outdoor area and to upload the image of the outdoor area. For more details on editing an outdoor area, see [“Configuring Outdoor Areas” on page 240](#).
 - Click **Delete** to delete the building.
8. If you added a floor to the building, select the floor and perform any of the following actions to continue building your network:

NOTE: You can add aisles and racks to a floor only from the Location view. However, you can view aisles, racks, and devices that you have assigned to these components from the Topology view.

- Click **Add Closet** to add a wiring closet to the floor.
 - Click **Add Aisle** to add an aisle to the floor.
 - Click **Assign Device** to assign devices to the selected floor. The Associate Devices to Floor window opens displaying all the unassigned devices in your network. Select the devices that you want to add to the floor and click **Add**.
 - Click **Edit Floor** to modify the name of the floor, floor level and upload the floor plan. For more details on editing a floor, see [“Configuring Floors” on page 231](#).
 - Click **Delete** to delete the floor.
9. If you added a wiring closet, select the wiring closet and perform any of the following actions:
- Click **Assign Device** to assign devices to the selected closet. The Associate Devices to Closet window opens displaying all the unassigned devices in your network. Select the devices that you want to add to the closet and click **Add**.
 - Click **Edit Closet** to modify the name of the wiring closet. In the Edit Closet window, modify the wiring closet name and click **Done**.
 - Click **Delete** to delete the wiring closet.
10. If you added an aisle, select the aisle and perform any of the following actions:

NOTE: You can add aisles and racks to a floor only from the Location view. However, you can view aisles, racks, and devices that you have assigned to these components from the Topology view.

- Click **Add Rack** to add a rack to the aisle.
- Click **Edit Aisle** to modify the name of the aisle. In the Edit Aisle window, modify the name and click **Done**.
- Click **Delete** to delete the aisle.

11. If you added a rack, select the rack and perform any of the following actions:

NOTE: You can add aisles and racks to a floor only from the Location view. However, you can view aisles, racks, and devices that you have assigned to these components from the Topology view.

- Click **Assign Device** to assign devices to the selected rack. The Associate Devices to Rack window opens displaying all the unassigned devices in your network. Select the devices that you want to add to the rack and click **Add**.
- Click **Edit Closet** to modify the name of the rack. In the Edit Rack window, modify the name and click **Done**.
- Click **Delete** to delete the rack.

12. Click **Done** to save the location details.

Network Director displays the location details along with the assigned devices in Location view.

RELATED DOCUMENTATION

[Understanding the Location View | 222](#)

[Deleting Sites, Buildings, Floors, Wiring Closets, and Devices | 238](#)

[Changing the Location of a Device | 236](#)

[Configuring Buildings | 229](#)

[Configuring Floors | 231](#)

[Setting Up Closets | 233](#)

[Network Director Documentation home page](#)

Creating a Site


IN THIS SECTION

- [How to Add or Edit a Location Site | 228](#)
- [Creating or Editing a Site | 228](#)

A site is the cornerstone of the location-based view of your network. Until you define a site, the default view of your network tree merely shows you a list of your unassigned devices. After you define a location site, you can build a tree structure of buildings, floors, wiring closets, and outdoor areas that can each be assigned devices. You are able to view the devices in the network by expanding and collapsing these location nodes. To setup a location in Network Director, the first step is to create a site.

This topic describes:

How to Add or Edit a Location Site

1. Click the Build Mode icon  **Build** in the Network Director banner.
2. Select **Location View** from the list in the View pane.
3. Click **Add Site** to add a new site or click **Edit Site** in the Tasks pane.
4. Fill in or change the fields on the page that opens.
5. Click **Done** to define the site and to save the configuration.

Creating or Editing a Site

Only a few fields are required to establish a site as shown in [Table 39](#).

Table 39: Site Creation Fields

Site Name	A descriptive name for the site. This field is mandatory.
City	The city where the site is located.
State	The state where the site is located.

Table 39: Site Creation Fields (*continued*)

Country	<p>The country where the site is located. Select the country from the list.</p> <p>This field is mandatory because it sets the regulatory country code for wireless devices. Network Director validates the country code against the country codes in the network's controllers and access points. If the codes do not match, a warning message is sent.</p>
---------	--

RELATED DOCUMENTATION

[Understanding the Location View | 222](#)

[Configuring Buildings | 229](#)

[Network Director Documentation home page](#)

Configuring Buildings

IN THIS SECTION

- [How to Add or Edit a Building | 229](#)
- [Adding or Editing a Building for a Location | 230](#)

At any time after you create a site, you can grow your location by adding buildings. You add a building to a site either from within the Location wizard or independently from the Add Building page.

This topic describes:

How to Add or Edit a Building

To add or change a building definition:

1. Ensure you are in the Build mode and Location view. Click **Build** in the Network Director banner to enter Build mode; select **Location View** from the list in the View pane.
2. If you want to add a building to a site:
 - a. Select the site in the Tasks pane , for example, Main Campus.

The Tasks pane refreshes to show your selected site and the tasks available at the site node.

- b. Click **Add Building** in the Tasks pane to open the **Add Building** page.
3. If you want to edit an existing building definition:
 - a. Select the building within the site, for example, Headquarters Building.
The Tasks pane refreshes to show your selected building and the available tasks that you can perform at the building node.
 - b. Click **Edit Building** in the Tasks pane to open the **Edit Building** page.
 4. Fill in the fields and click **Done** to submit the information and to refresh the network tree.

Adding or Editing a Building for a Location

[Table 40](#) describes the fields needed to establish a building.

Table 40: Add or Edit Building Fields

Field	Description
Building Name	Type a representative name for the building. The Building Name is a required field.
Address	Type an address. The address can be the street address, building number, or any other identification that helps distinguish it from other buildings.
Done	Click to submit the information. Your view updates to reflect the building change under the site name in the network tree.
Cancel	Click to close the window without changes.

RELATED DOCUMENTATION

[Understanding the Location View | 222](#)

[Configuring Floors | 231](#)

[Assigning and Unassigning Devices to a Location | 234](#)

[Network Director Documentation home page](#)

Configuring Floors

IN THIS SECTION

- [How to Add or Edit a Floor | 231](#)
- [Adding or Editing a Building Floor for a Location | 232](#)


You can refine the a building location and designate floors within the building. Use the Add Floor page to:

- Name a floor
- Note the floor level
- Upload a floor plan for viewing
- View an uploaded floor plan

This topic describes:

How to Add or Edit a Floor

Within each building you can define the number of floors and attach the floor plan for online viewing.

1. Click the Build Mode icon  in the Network Director banner.
2. Select **Location View** from the list in the View pane.
3. If you want to add a floor to a building:
 - a. Select the building in the network tree to which you want to add floors, for example, Headquarters.
The Tasks pane refreshes to show your selected building and the available tasks for the building.
 - b. Click **Add Floor** in the Tasks pane to add a new floor to the building.
4. If you want to edit an existing floor definition:
 - a. Select the floor within the building, for example, Lobby-Floor 1.
The Tasks pane refreshes to display the selected building floor and the available tasks that you can perform at the floor node.

- b. Click **Edit Floor** in the Tasks pane to open the Edit Floor page.
5. Fill in the fields for the floor name and level.
6. (Optional) Upload an image of the floor plan.
7. (Optional) View the floor plan, if available.
8. Click **Done** to submit the information and to refresh the network tree.

Adding or Editing a Building Floor for a Location

To add or change information about a building floor, use the fields in [Table 41](#).

Table 41: Floor Field Descriptions

Field	Description
Floor Name	Type the name of the floor. This field is required.
Floor Level	Use the arrow keys to set the floor number.
Add/Update	Upload a image of the floor plan.
View	View an existing floor plan.
Done	Saves the floor configuration information, and returns you to Device Inventory page in the default view.
Cancel	Discards any configuration changes.

RELATED DOCUMENTATION

[Understanding the Location View | 222](#)

[Setting Up Closets | 233](#)

[Configuring Buildings | 229](#)

[Assigning and Unassigning Devices to a Location | 234](#)

[Network Director Documentation home page](#)

Setting Up Closets

IN THIS SECTION


- [How to Add or Edit a Closet | 233](#)
- [Adding or Editing a Wiring Closet | 234](#)

Use the Add Closet or Edit Closet tasks to create or change the name of a wiring closet. These tasks are visible only from a floor node in a building.

This topic describes:

How to Add or Edit a Closet

To add or change a wiring closet:

1. Click the Build Mode icon  **Build** in the Network Director banner.
2. Select **Location View** from the list in the View pane.
3. Navigate to the building and floor where you are adding or changing the closet.
4. If you are adding a wiring closet:
 - a. Select a building floor in the network tree to which you want to add a wiring closet.
The Tasks pane refreshes to show your selected floor and the available tasks for the floor.
 - b. Click **Add Closet** in the Tasks pane.
5. If you are changing a wiring closet, click **Edit Closet** in the Tasks pane.
6. Type the closet name and click **Done** to save the configuration.
The closet appears with the change in the network tree.

Adding or Editing a Wiring Closet

The Add Wiring Closet or Edit Wiring Closet pages allow you to name a wiring closet. Simply type the name of the new or changed wiring closet and click **Done** to submit the information to the system. Your network tree refreshes to show the wiring closet.

RELATED DOCUMENTATION

[Understanding the Location View | 222](#)

[Configuring Floors | 231](#)

[Assigning and Unassigning Devices to a Location | 234](#)

[Network Director Documentation home page](#)

Assigning and Unassigning Devices to a Location

IN THIS SECTION

- [How to Assign or Unassign Devices | 235](#)
- [Assigning Devices | 235](#)

You can assign devices or remove assignments from devices by their location. Your choices for device assignment are dependent upon the type of device and your position in the site. For example, you cannot assign access points to buildings or closets. However, you can assign access points to floors or outdoor areas. For details on which devices can be assigned to a location node, see the Devices that can be Assigned to each Location Component table from the [“Setting Up the Location View” on page 224](#).

This topic describes:

How to Assign or Unassign Devices

To assign devices to a specific location:

While in Build mode,

1. Select **Location View** from the list in the View pane.

The network tree displays discovered devices under the physical locations already defined in Network Director. The root node (for example, My Network) is selected by default. The devices that are assigned to the locations are displayed under the nodes for respective locations, such as buildings and floors. All devices that are not assigned to any location are displayed under the Unassigned node.

2. Navigate the network tree to the location where you want to add a device.

Both the Tasks pane and Device Inventory page update to reflect the location's current configuration.

3. Select one of the following tasks in the pane to open Add/Remove Devices for Selected Location.

- Assign Devices to Building
- Assign Device to a Floor
- Assign Devices to a Wiring Closet
- Assign Devices to an Outdoor Location

4. Navigate the tree to find an available device under Unassigned in the left portion of the page.

5. Select the device and click the double right arrows to assign it to the target location on the right. To unassign a device, select the device in the Assigned Devices to Selected Location column and click the double left arrows. Repeat this step until you have finished assigning and unassigning devices.

6. Click **OK** at the bottom of the page to save the assignment. The network tree refreshes to display the device in the new location.

Assigning Devices

Use the Add/Remove Devices for Selected Location to find a device and assign it to a location within a site. Locate the device in the Available Devices column and assign it by clicking the double right arrows. Use the same method to unassign a device by selecting it in the Assigned Devices to Selected Location column and double clicking the double left arrows.

You can assign switches, access points, controllers, QFabric devices and members, Virtual Chassis devices and members, Virtual Chassis Fabrics and corresponding member devices to buildings, floors, aisles, and closets.

While assigning QFabric devices, Virtual Chassis devices, or Virtual Chassis Fabric devices to a location within a site, you can either assign the logical device—QFabric system, Virtual Chassis, or the Virtual Chassis Fabric—as a single device *or* one or more member devices that belong to these logical devices, but not both.

NOTE: Network Director displays the Virtual Chassis, Virtual Chassis Fabric, and QFabric systems in the Location view network tree only if the following conditions are met:

- QFabric system, Virtual Chassis, or the Virtual Chassis Fabric is assigned to a location.
- At least one of their member devices are *not* assigned to any location entity.

If all the member devices are assigned to location entities, then the Virtual Chassis, Virtual Chassis Fabric, and QFabric systems are not displayed in the network tree.

RELATED DOCUMENTATION

[Understanding the Location View | 222](#)

[Configuring Buildings | 229](#)

[Configuring Floors | 231](#)

[Setting Up Closets | 233](#)

[Configuring Outdoor Areas | 240](#)

[Network Director Documentation home page](#)

Changing the Location of a Device

IN THIS SECTION

- [How to Move a Device to a New Location | 237](#)
- [Changing the Location of a Device | 237](#)

The Change Location of Device task is an easy way to move a device address to another building, floor, or wiring closet location within the site. You can move an access point to another floor or to another outdoor area. However, you cannot move an access point to a building or wiring closet. The Change

Location of Device task is available whenever you select an assigned device in the Location or Logical views.

This topic describes:

How to Move a Device to a New Location

To move a device address to another location:

1. Select a device in the network tree that is currently assigned to a building, floor, or closet.
2. Click **Change Location of Device** to open the Change Location of Device page.
3. Using the Location View, navigate the tree and select the new location for the device. You can move an access point, only to another floor or outdoor area.
4. Click **OK** to move the device assignment and to save the new configuration.

Changing the Location of a Device

The Change Location of Device page consists of two components: Selected Device Details and Location View. Use the Selected Device Details portion of the page to review information about the device and its current location. The fields in Selected Device Details page are described in [Table 42](#).

Table 42: Contents of Selected Device Details

Field	Description
Device Name	Hostname
Device IP	Device Address
Device Family	Hardware family of products, for example, Junos-QFX.
Location	Gives the current location of the device in the format of site/building/floor/cabinet

Location View is a copy of the network tree for you to navigate and designate the new location for the device.

RELATED DOCUMENTATION

[Understanding the Location View | 222](#)

[Assigning and Unassigning Devices to a Location | 234](#)

[Network Director Documentation home page](#)

Deleting Sites, Buildings, Floors, Wiring Closets, and Devices

IN THIS SECTION

- [How to Delete a Location Object | 238](#)
- [Deleting Sites | 239](#)
- [Deleting Buildings | 239](#)
- [Deleting Floors | 239](#)
- [Deleting Closets | 239](#)
- [Deleting Devices | 239](#)

From the Build mode Tasks pane, you can delete any sites, buildings, floors, wiring closets and their associated devices. When you delete one of these objects, it removes not only that item but all child objects within the node. All associations related to the node and below are also removed. Devices are moved to the Unassigned node in the network tree. Be sure you understand what is being deleted on the node when you choose to delete a node.

For example, if you delete a building, it deletes the building, all floors, all wiring closets in that building. All of the devices in the building are moved to Unassigned in the network tree. When you delete a building, the site and any other buildings and their associations remain.

How to Delete a Location Object

1. Ensure you are in the Build mode and Location view. Click **Build** in the Network Director banner to enter Build mode; select **Location View** from the list in the View pane.
2. Select any object within the site. The option to delete the object appears in the Tasks pane.
3. Confirm the deletion of the object.

Deleting Sites

There is only one method of deleting a site: select the site in the Tasks pane and click **Delete Site**. Use caution with this selection. When you click **Delete Site** you are given the opportunity to confirm or cancel the deletion. If you confirm the deletion, you remove the site and everything in the site. All devices become unassigned and are not associated with any buildings, floors, or wiring closets.

Deleting Buildings

When you delete a building, it removes the building, all floors, and all wiring closets within that building. All devices become unassigned and are not associated with the building, its floors, or its wiring closets. Only one building can be deleted at a time. To delete a building, select the building in the network tree and click **Delete Building**. Confirm the deletion to remove the objects and to disassociate the devices. If a site is deleted, all of the buildings within the site are also deleted.

Deleting Floors

When you delete a floor, it removes the selected floor and all wiring closets on that floor. All devices assigned to the floor or to the closets on that floor become unassigned and become available for reassignment. To delete a floor, select the floor in the network tree and click **Delete Floor**. Confirm the deletion to remove the objects and to disassociate the devices. If a site or building is deleted, the floors are also deleted.

Deleting Closets

When you delete a closet, it removes the selected closet and unassigns the devices in the closet. Those devices then become available for reassignment. To delete a closet, select the closet in the network tree and click **Delete Closet**. Confirm the deletion to remove the objects and to disassociate the devices. If a site, building, or floor is deleted, the associated closets are also deleted.

Deleting Devices

At every node in the network tree, you can choose to delete devices directly.

BEST PRACTICE: However, it is usually best to select the node directly above the device so that you do not accidentally unassign more devices than desired.

- Select the node (site, building, floor, or closet) directly above the device.
- Click **Delete Devices** to open the Delete Devices page.
- Click the plus signs to expand the node until you locate the device.

- Click one or more boxes to select the devices. If you do not navigate down to the device level and select a node at a higher level (such a closet or floor), the system selects all devices at and below the node.
- Click **OK** and confirm your selection to remove the assignment. The devices are moved to the Unassigned node of the network tree.

RELATED DOCUMENTATION

[Understanding the Location View | 222](#)

[Network Director Documentation home page](#)

Configuring Outdoor Areas

IN THIS SECTION

- [How to Configure an Outdoor Area | 240](#)
- [Configuring an Outdoor Area | 241](#)

You can associate an outdoor area to a site or a building for wireless coverage and upload an image or map of that area. After you designate an outdoor area, you can edit or view the map using the Edit Outdoor Area task.

This topic describes:

How to Configure an Outdoor Area

To create an outdoor area without using the wizard:

- Ensure you are in Build mode and Location view. Click **Build** in the Network Director banner to enter Build mode; select **Location View** from the list in the View pane.
- Click **Add Outdoor Area** in the Tasks pane. The Add Outdoor Area page opens.
- Fill in the name and upload the optional map.
- Click **Done** to save the data and to return to the default view.

Configuring an Outdoor Area

Table 43 describes the fields and buttons necessary to create or change an outdoor area.

Table 43: Outdoor Area Fields

Field	Description
Outdoor Area Name	Type the name of the outdoor area. Network Director associates the outdoor area with the building.
Upload	Optional step to upload an image of the outdoor area. Use the Upload Map window to navigate to the image file location.
Done	Click to save the configuration. The network tree is updated to reflect the change.
Add/Update	Click to add a map or overlay an existing map of the area.

RELATED DOCUMENTATION

[Understanding the Location View | 222](#)

[Setting Up the Location View | 224](#)

[Assigning and Unassigning Devices to a Location | 234](#)

[Network Director Documentation home page](#)

Building a Topology View of the Network

IN THIS CHAPTER

- Understanding the Network Topology in Network Director | 243
- Understanding the Topology View Tasks pane | 247
- Setting Up the Topology View | 250
- Managing the Topology View | 251
- Adding and Managing OUI Data in Network Director | 271

Understanding the Network Topology in Network Director

Junos Space Network Director provides features for monitoring and managing Juniper Networks EX Series Ethernet Switches, QFX Series devices, and Juniper Networks WLC Series Wireless LAN Controllers (WLCs) besides enabling connectivity visualization between discovered and managed devices such as routers, switches, controllers, and access points. Connectivity between devices and their association with their location provide the foundation for rendering topology in a complete manner.

As a network administrator, you must have a clear understanding of the various networking devices in your network, their physical locations, and how these devices are interconnected in your network. The network topology represents the interconnection between various devices in your network, which are managed by Network Director, based on their connectivity and association to their physical surroundings. The network topology provides a visual insight into the network, which is useful for debugging, troubleshooting, planning, and executing administrative actions.

Before you access the topological view of your network, you must:

- Connect your Network Director system to the Internet before accessing Topology View in Network Director as this feature works only while the system is connected to the Internet.

NOTE: Ensure that Internet connection is available for both the Network Director and Network Director client systems.

- Discover the devices managed by Network Director in your network. For details about discovering devices, see [“Discovering Devices in a Physical Network” on page 204](#).

NOTE: You must specify the SNMP parameters during device discovery to have all the devices discovered and managed by Network Director available in Topology View. However, you can specify the SNMP parameters in the **Refresh Topology** task from the Topology View also.

NOTE: Ensure that you have enabled the LLDP, STP, or RSTP protocols on the devices as Network Director uses these protocols to determine the connectivity of devices with their neighbors in the network. LLDP and RSTP protocols are enabled by default on all EX Series switches and QFX Series devices.

- Set up the physical location of the devices in your network based on the geographical location of the devices. Some examples of location nodes are: sites, buildings, floors, closets, and aisles. You can set up one location node within another location node as in buildings within a site or floors within a building. You can then assign the network devices based on their location in buildings, floors, outdoor areas, closets, and racks. Apart from using the Location Management tasks in the Location View, you can set up the location details of the network devices managed by Network Director by using the *Setup Locations*

task from the Task menu in the Topology View. For details, see [“Setting Up the Location View” on page 224](#). The device-to-topology-group location relationship is established when a device is placed at a specific location on a topology map.

Network topology enables you to view all the discovered devices in your network, overlaid on a map where the devices are located across sites, buildings, floors, outdoor area, closets, and racks along with their physical interconnection with other devices in your network. Topology also provides visualization around physical connectivity between various discovered interconnected devices.

You can use the Topology View to zoom in or zoom out of a site to a building and a building to a site. In the Topology View, you can also double-click a node such site, buildings, and so on to navigate to the next node. You can also see the connectivity between a device and its immediate neighbors, alarms details, and so on. Network Director also enables you to assign devices to buildings, floors, closets, and outdoor areas on the map.

Network topology also provides visualization around physical connectivity between various discovered interconnected devices. You can upload a floor plan at the floor level or a map at the outdoor area if you already have the floor plan or map available. You can then move the nodes in the Topology View page based on the floor plan.

Network Director supports QFabric internal topology.

The topology display is created by layering the device images on top of the imported floor plan images as shown in [Figure 10](#) and [Figure 11](#).

Figure 10: Typical Floor Plan Displaying the Closets and Devices

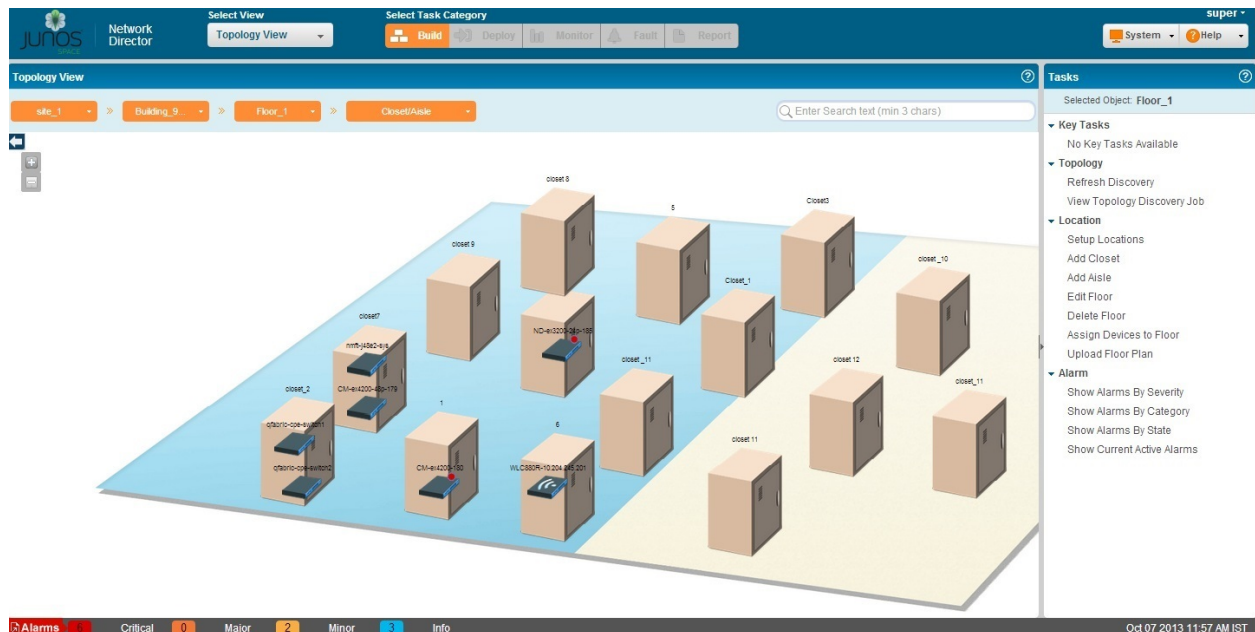
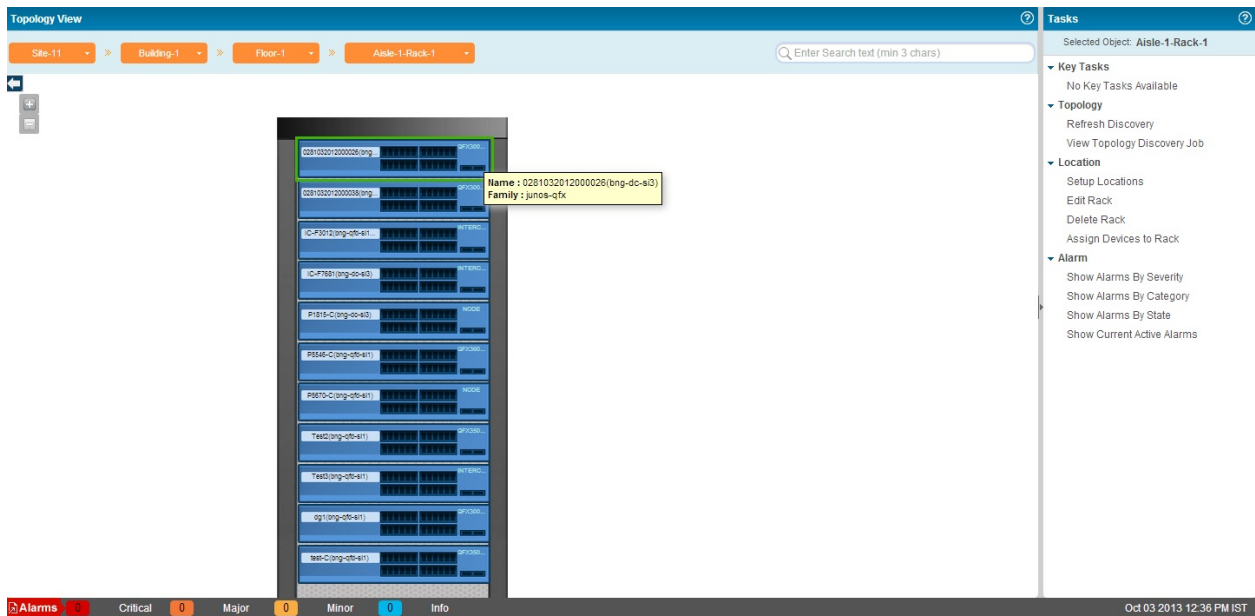
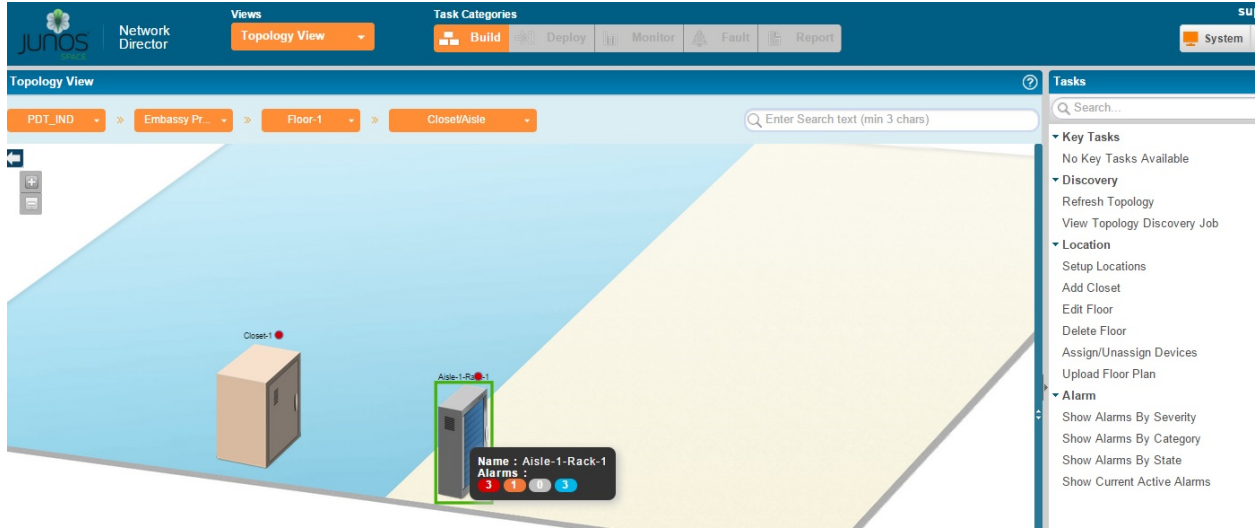


Figure 11: Typical Floor Plan Displaying the Racks and QFX Devices



RELATED DOCUMENTATION

[Discovering Devices in a Physical Network | 204](#)

[Setting Up the Location View | 224](#)

[Managing the Topology View | 251](#)

[Network Director Documentation home page](#)

Understanding the Topology View Tasks pane

The Tasks pane in Topology View contains all the tasks you can do in the Topology View. Click a specific task to begin that task.

Not all tasks are available by default in the Topology View. As you change your selections in the Topology map pane, the tasks in the Tasks pane also change—for example, tasks such as View Virtual Chassis Fabric connectivity are visible only after you select a device that is part of a Virtual Chassis fabric (VCF). Similarly, Alarm tasks are available only after you select a site, building, floor, closet, or device.

Topology View tasks are divided into the following categories in the Tasks pane.

- **Key Tasks**—The most preferred tasks that you want to do while you are using the Topology View. You can add important tasks from the topology task menu and hence the Key Tasks are a duplicate of select tasks from the Topology and Location tasks menu.
- **Discovery**—The tasks you do to create the topology of your network.
- **Location**—The tasks you do to create a location such as a site, building, floor, closet, aisle, rack, and outdoor area.
- **Alarm**—The tasks that enable you to monitor the fault alarm details of the devices.

[Table 44](#) describes the Topology View tasks.

Table 44: Topology View Tasks

Task	Description
Key Tasks	A duplicate of the most important tasks from other tasks menu. You can add your frequently used tasks to the key tasks menu.
Connectivity	
Displays the device-level connectivity	
View Virtual Network Connectivity	Displays the connectivity of the selected device with the virtual network. NOTE: This task is available only after you select a device that is connected to a virtual network.
View Virtual Machines	Displays virtual machines in the grid format that are connected through the switch.
View Device Connectivity	Displays the connection details of a device with its neighbors in graphical and grid views. If the selected device is connected to more than 60 devices, then the connection details are displayed only in grid view.

Table 44: Topology View Tasks (continued)

Task	Description
View QFabric Connectivity	<p>Displays the control plane and data plane connectivity details between the interconnects and nodes.</p> <p>NOTE: This task is applicable for QFabric devices and its members only.</p>
View VC/VCF Connectivity	<p>Displays the Virtual Chassis (VC) or Virtual Chassis Fabric (VCF) connectivity from the selected device.</p> <p>NOTE: This task is available only after you select a device that is part of the VC or the VCF.</p>
View Layer 3 Fabric Connectivity	Displays the devices and their physical connectivity in the spine and leaf topology.
Discovery	
Displays the device discovery tasks	
Refresh Topology	Refreshes the devices discovered and managed by Network Director earlier from Build mode. This task also refreshes the device connectivity.
View Topology Discovery Job	Displays the discovery jobs in the Topology View pane.
Location	
Displays the location related tasks	
Setup Locations	<p>Creates a new location. This task has sub tasks to add sites, buildings, floors, outdoor area, and closets.</p> <p>NOTE: You cannot create aisles and racks from Topology View. You must create aisles and racks from the Location View work flow.</p>
Add Site	<p>Creates a new site in Location View.</p> <p>NOTE: Use this task only to create the site object. Buildings, floors, closets, and outdoor areas in the site must be created separately.</p>
Add Building	<p>Creates a new building in the selected site.</p> <p>NOTE: Use this task only to create the building. Floors and closets in the building must be created separately.</p>
Add Outdoor Area	Creates a new outdoor area in the selected site.

Table 44: Topology View Tasks (continued)

Task	Description
Upload Floor Plans	Enables you to upload a floor plan for a floor in a building, if you already have a floor plan.
Upload map	Enables you to upload a map to an outdoor area.
Delete Site/Edit Site	Deletes or modifies the selected site.

Alarm

Displays the alarm related tasks

Show Alarms by Severity	Displays the fault alarm details sorted based on the severity; that is from critical, major, minor, and info.
Show Alarms by Category	Displays the fault alarm details sorted based on the category; that is from active, acknowledged, and cleared.
Show Alarms by State	Displays the fault alarm details sorted based on the state; that is from active, acknowledged, and cleared.
Show Current Active Alarms	Displays any active alarm that has not yet been cleared.

Device Management

Displays the various device management tasks

SSH To Device	Launches the SSH connection to the device. You can launch the SSH connection for a device from a location such as a building, floor, rack, or closet. This task is available when you click the device in a particular location. For more details, see “Accessing a Device’s CLI from Network Director” on page 1151 .
Launch Web View	Launches the Web user interface connection to the device. You can launch the Web user interface connection for a device from a location such as a building, floor, rack, or closet. This task is available when you click the device in a particular location. For more details, see “Accessing a Device’s Web-Based Interface from Network Director” on page 1153 . NOTE: This task is applicable to only to those devices that support Web user interface access.
Manage Port Admin State	Enables or disables one or more ports of the selected device. For more details, see “Enabling or Disabling Network Ports on Switches” on page 1250 .

RELATED DOCUMENTATION

[Understanding the Network Topology in Network Director | 243](#)

[Managing the Topology View | 251](#)

[Network Director Documentation home page](#)

Setting Up the Topology View

Topology View enables you to view all the discovered devices in your network, overlaid on a map. You can create sites, buildings, floors, outdoor areas, closets, and racks by using the Location wizard in the Topology View. You can use the Topology View to zoom in or zoom out of a site or a building. You can also see the connectivity between a device and its immediate neighbors, alarms details, port details, and so on.

Before you start, ensure that:

- The devices are discovered using the Device Discovery task. For detailed steps, see [“Discovering Devices in a Physical Network” on page 204](#) or [“Discovering Devices in a Datacenter Network” on page 212](#).
- You have specified SNMP parameters during device discovery.
- Ensure that you have enabled the LLDP, STP, or RSTP protocols on the devices as Network Director uses these protocols to determine the connectivity of devices with their neighbors in the network. LLDP and RSTP protocols are enabled by default on all EX Series and QFX Series devices.

Perform the following steps to set up your network in the Topology View:

1. After you have discovered the devices using the Device Discovery task, open Location View or the Topology View. You can set up sites, buildings, floors, outdoor areas, closets, aisles, and racks using Location View. The Topology View allows you to create most of these components except racks and aisles. For detailed steps on building your location view, see [“Setting Up the Location View” on page 224](#).
2. Assign devices to the various entities in the Location View. Network Director enables you to assign devices to the various components of your network. [Table 38](#) describes the devices that you can assign to each of the location component.
3. Save the location settings.
4. Select **Topology** from the Network View Selector.

Network Director lays out the sites that you created in Location View, in the topology map. Network Director uses the location that you specified for each site and building to place them on the map. A site that does not have the location specified, will be placed in the default location—United States.

5. Click **Discovery > Refresh Discovery** from the Tasks menu to refresh the topology of devices. Network Director refreshes the Topology View for all the devices that you have added to Topology View. After a successful discovery, you can select and view the connectivity of a particular device with their immediate neighbors using the tasks in the Connectivity section of the Tasks pane.
6. Follow the steps outlined in [“Managing the Topology View” on page 251](#) to perform various tasks from Topology View.

RELATED DOCUMENTATION

[Understanding the Network Topology in Network Director | 243](#)

[Managing the Topology View | 251](#)

[Network Director Documentation home page](#)

Managing the Topology View

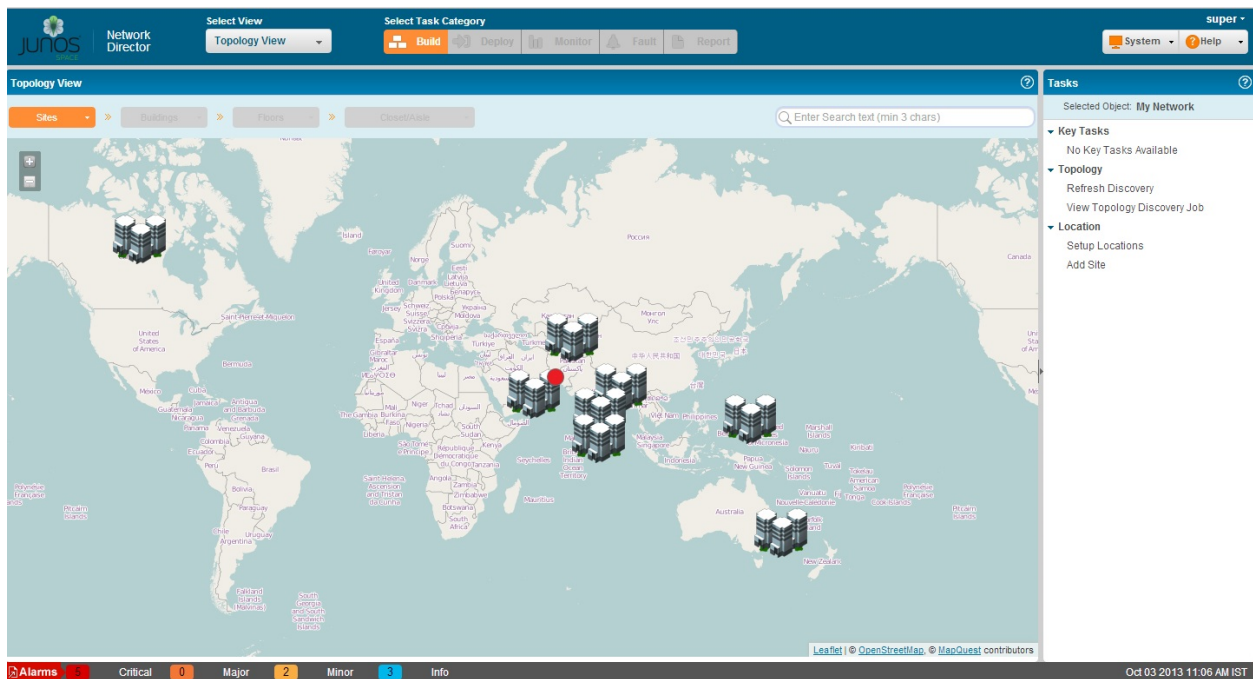
IN THIS SECTION

- [Viewing the Network Topology | 252](#)
- [Refreshing the Topology | 254](#)
- [Viewing Topology | 255](#)
- [Viewing Topology Discovery Job | 256](#)
- [Setting Up Locations | 257](#)
- [Viewing the Alarm Details | 257](#)
- [Discovering the Linux Hosts | 257](#)
- [Displaying Device Connectivity | 257](#)
- [Displaying QFabric Connectivity | 262](#)
- [Displaying Virtual Chassis and Virtual Chassis Fabric Connectivity | 265](#)
- [Displaying Virtual Network Connectivity | 268](#)
- [Displaying Third-Party Device Details | 268](#)
- [Uploading Floor Plans | 269](#)
- [Uploading Topology Map | 270](#)

Viewing the Network Topology



Topology View enables you to view all the discovered devices in your network, overlaid on a map. You can create sites, buildings, floors, outdoor areas, closets, and racks by using the Location wizard in Topology View. You can use the Topology View to zoom in or zoom out of a site or a building. You can also see the connectivity between a device and its immediate neighbors, alarms details, port details, and so on. An example of how the topology map looks like after you have added the location details is shown in [Figure 12](#).

Figure 12: Main Topology Window



The topology display is created by layering the device images on top of the imported floor plan images.

In addition to the tasks outlined in the [Table 44](#), you can perform the following tasks from the Topology View pane:





- **Zoom In and Zoom Out**—You can use the zoom in () or zoom out () buttons to get a detailed or high-level view. Network Director enables you to zoom in and view details up to the rack level, if you have defined racks and assigned devices to the rack.
- **Pan**—Network Director enables you to pan the topolog. You can move the devices to different parts on the topology by holding the mouse button down and dragging to a specific part of the map.


Network Director displays the sites, buildings, and geographical coordinates on the topology map based on the address specified while setting up the locations. However, you can move the devices around to another location by selecting, dragging, and dropping the device to the correct location. For example, you can move a site from a US site to a Bangalore site on the topology by using the Pan feature. This

helps to determine the correct location of a site because while setting up the location details, it is not mandatory to provide the address information.

- **Search and Locate**—You can search for all nodes such as sites, building, floors, specific devices, all devices in a floor, and so on by entering a search keyword or the complete name of the device in the Search field. On entering the search criteria, you might see the list of nodes or just one node based on the search criteria. When you select the node, Network Director locates the node and Topology View is panned to ensure the selected node is centered on the page. In cases where a device is at the edge of the map the corner is aligned accordingly to bring device in view.
- **View details**—If you mouse over a entity (site, building, device), the entity gets highlighted and Network Director displays details such as the building name or IP address and the number of active alarms on that entity. A colored dot that appears on the upper right corner of the entity identifies whether there are alarms on the entity and the color of the dot indicates the severity level of the alarm. See [Table 45](#) to know more about the alarm severity indicator for each alarm severity.


Table 45: Alarm Severity Indicator

Alarm severity	Indicator
Critical	
Major	
Minor	
Informational	

NOTE: Each of these entities might have alarms of different severities. Network Director displays the indicator based on the most severe alarm on an entity. For example, if a device has 3 informational alarms and one major alarm, Network Director displays the indicator for the major alarm ().

- **Highlight and select View Device Connectivity**—Select a device and click **View Device Connectivity**, **View QFabric Connectivity**, **View VC/VCF Connectivity**, or **View Virtual Network Connectivity** to view the selected device's connectivity with other devices. Each of these tasks, except the **View Device Connectivity**, are visible only when you select a device that is part of a QFabric system, Virtual Chassis, Virtual Chassis Fabric, or a Virtual network. The device images are displayed along with details such as name, IP address, and the connectivity link between the devices.
- **Navigation**—Use the navigation breadcrumbs at the top of the page to navigate through sites, buildings, floors, closets, or aisles. For there are more than one entities at any give level, you can use the Down

arrow in the breadcrumb to navigate to that entity. For example, if you want to navigate from floor-1 in building-1 to floor-3 in building-2, you can use the down arrow in the building breadcrumb to select building-2 and the down arrow in the floor breadcrumb to select floor-3.

- **Host Information**—You can use the  button to expand the members to view the host details. You can also view the virtual machines if the host is a hypervisor and managed by Network Director.

Refreshing the Topology

You can refresh the devices discovered and managed by Network Director from the Tasks pane in the Topology View and the Datacenter View. You can add the Refresh Topology task to the Key tasks in both the views if you will use this task frequently.

To refresh the device discovery process:

1. Click **Refresh Topology** from Key Tasks or the Topology View, or from the Datacenter Discovery panes.

The Refresh Discovery window is displayed along with the SNMP details that you specified in the discovery options in the Device Discovery task page.

If you have not specified the SNMP details while discovering the devices, proceed to Step 3.

2. Select the **SNMP version** from the SNMP version table and click **Discover**. The Refresh Topology window appears displaying the progress of the topology discovery.
3. If SNMP version details are not displayed in the SNMP version table, click **Add** on the Refresh Discovery window.

The Add SNMP Settings dialog box is displayed.

Select either **SNMP V1/V2C** or **SNMP V3**. Based on the selection, you need to enter the details as follows:

- If you selected SNMP V1/V2C, specify a community string, which can be *public*, *private*, or a predefined string.

Click **Add** in the Add SNMP Settings dialog box or click **Add More** to add more strings to the community. If you click **Add More**, when you are done adding all the strings, click **Add** to save the SNMP settings for V1/V2C.

- If you selected SNMP V3, specify the following:
 - Enter a username.
 - Select the privacy type (AES 128, DES, or None).
 - Enter the privacy password (if AES 128 or DES). If you specify none for the privacy type, the privacy function is disabled.
 - Select the authentication type (MD5, SHA, or none).

- Enter the authentication password (if MD5 or SHA). If you specify none for the authentication type, the authentication function is disabled.
- Click **Add** to save the SNMP settings and close the dialog box, or click **Add More** to add additional configurations. If you clicked Add More, click **Add** to save the settings and close the dialog box.

The specified details are displayed in the SNMP version table.

4. Click **Discover**.

The Refresh Topology window is displayed, showing details of the device discovery.

Viewing Topology

The Refresh Topology window displays the refresh device discovery job details as described in [Table 46](#).

Table 46: Refresh Topology Job Details

Field	Description
Job Name	The Refresh topology job name along with the Job ID.
Start Time	The time at which the refresh discovery job is initiated.
End Time	The time at which the refresh discovery job is completed.
Percentage Progress	Displays the progress of the job in percentage. When the job is completed, displays 100 percentage.
Status	The status of the job. The status is <i>In Progress</i> until the job is completed.
Target Devices Count	The total number of targets for the discovery of devices.
Discovered Devices Count	The total number of discovered devices for which the SNMP parameters are specified. NOTE: In the Topology View, you can view the network connections of only those discovered devices for which LLDP, SNMP, and STP parameters are set.
Discovered Subnets Count	The total number of subnets discovered.
Targets	The Targets table displays the Management IP addresses of the devices discovered along with the status of the discovery process.

Use the right and left arrows to navigate through the discovered pages. You can specify the details to be displayed in a page by selecting the **show items** list box and specifying the number of items to be displayed in one page.

Click **Close** to close the Refresh Topology page and return to the main Topology page.

Viewing Topology Discovery Job

To view the discovery jobs in the Topology View and the Datacenter View, click **View Topology Discovery Job**. The Topology Discovery Jobs window opens displaying details of the topology related jobs as described in [Table 47](#). To view any hidden column, mouse over column heading, select the down arrow, and then click **Columns**. Select the check box to display the hidden columns.

Table 47: Job Details for Topology Discovery

Field	Description
Job ID	For each job-based task, the audit log includes a job ID.
Job Name	The name of the job.
Percent	The percentage of completion of the job.
State	The status of the job: <ul style="list-style-type: none"> • Success—Job completed successfully • Failure—Job failed and was terminated • Job Scheduled—Job is scheduled but has not yet started • In progress—Job is has started, but not completed • Cancelled—Job is cancelled
Summary	Summary of the job scheduled and executed with status.
Scheduled Start Time	The UTC time on the client computer when the job is scheduled to start.
Actual Start Time	The actual time when the job started.
End Time	The time when the job was completed.
User	The login ID of the user that initiated the task.
Recurrence	The recurrent time when the job will be restarted.

To view the details of a topology discovery job, select a row and click **Show Details**. To cancel a scheduled job, select a job that is scheduled for a later time or a job that is in progress and click **Cancel**.

Setting Up Locations

You can set up locations and assign devices to these locations by setting up the Location View. To set up Location View, see [“Setting Up the Location View” on page 224](#).

Viewing the Alarm Details

You can view the alarm details at site, building, floor, closet, aisle and rack levels. Based on the location node, the alarms are aggregated and displayed. That is, all the alarms for a particular building is aggregated and displayed at the building level. The alarms display as red, orange, yellow, and blue dots indicating critical, major, minor, and info alarms. Network Director updates and displays the alarm status changes in real time in the Topology view. To view the Alarm details for a device, navigate to the device level, select a device, and click one of the following options from Alarm in the Tasks pane:

- Show Alarms By Severity, see [“Alarms by Severity Monitor” on page 1466](#)
- Show Alarms By Category, see [“Alarms by Category Monitor” on page 1466](#)
- Show Alarms By State, see [“Alarms by State Monitor” on page 1467](#)
- Show Current Active Alarms, see [“Current Active Alarms Monitor” on page 1464](#)

Discovering the Linux Hosts

Based on the LLDP discovery method, you can discover the hosts for various Linux platforms such as Ubuntu, CentOS, and Red Hat by clicking **Refresh Topology**. You can also view the additional information in the tool tip that appears when you mouse over a host. These hosts also show icons that identify if a server is based on its respective platform. For example, if a server is a generic Linux server, a Linux icon is shown in the topology.

NOTE: You must enable LLDP both on the switch and on the host.

Displaying Device Connectivity

At the device level, you can view the connectivity details of a device and the details of all the devices that are connected to the specified device by using Topology View in Network Director. The Device Connectivity View also displays various details about the selected device and the immediate neighbors. The level of detail that Network Director displays in the Device Connectivity View differs based on the type of device that you select.

To view the connectivity details of a device:

1. Do one of the following:

- While in the Logical, Location, Device, or Custom View, select the device for which you want to view connectivity from the View pane and click **Connectivity > View Device Connectivity** from the Tasks pane.
- In the Topology View, navigate to a device in a building, floor, outdoor area, closet or rack and select the device for which you want to view the connectivity details and click **Connectivity > View Device Connectivity** from the Topology Tasks pane.

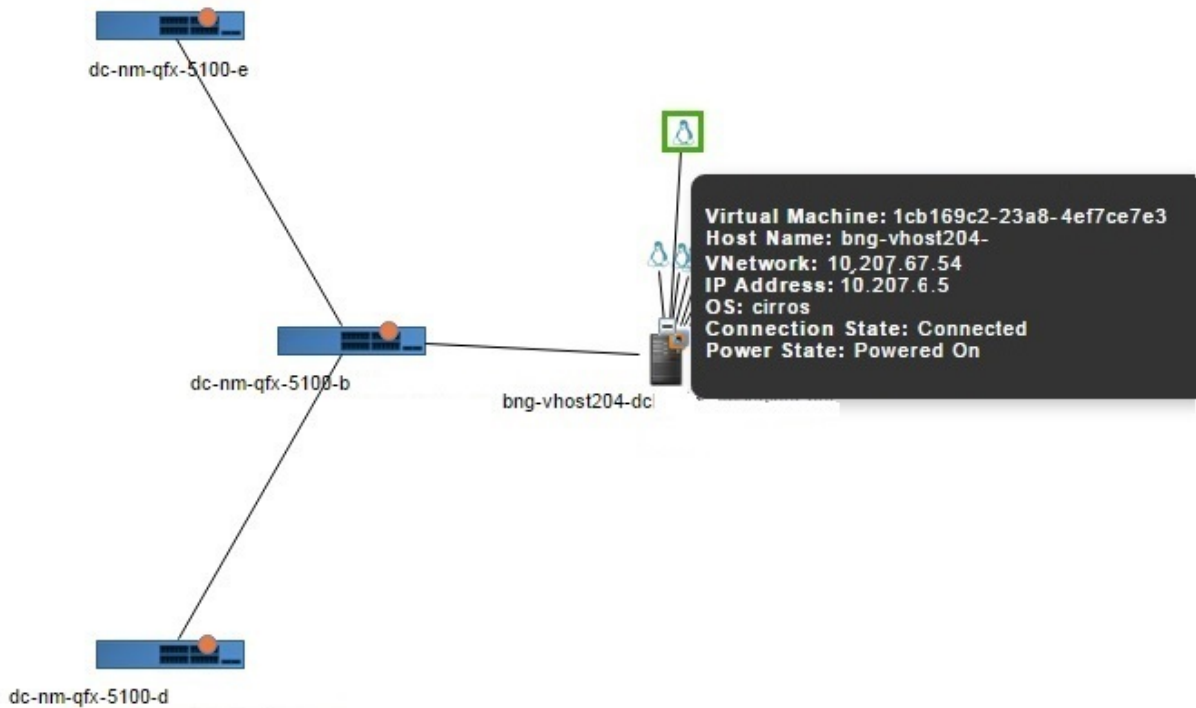
NOTE: The Connectivity task container is available only after you select a device.

The Device Connectivity page opens. You can view the device connectivity details either in graphical view or in grid view. The default view is the graphical view.

In the graphical view, the device is displayed in the center and its network connectivity to all the connected devices are displayed as in [Figure 13](#). Mouse over a device to view details of the highlighted device.

NOTE: If the selected device is connected to a device that is not a Juniper Networks device, the latter appears dimmed in the Device Connectivity page indicating that the device is not managed by Network Director.

Figure 13: Displaying the Connection Details in Graphical View



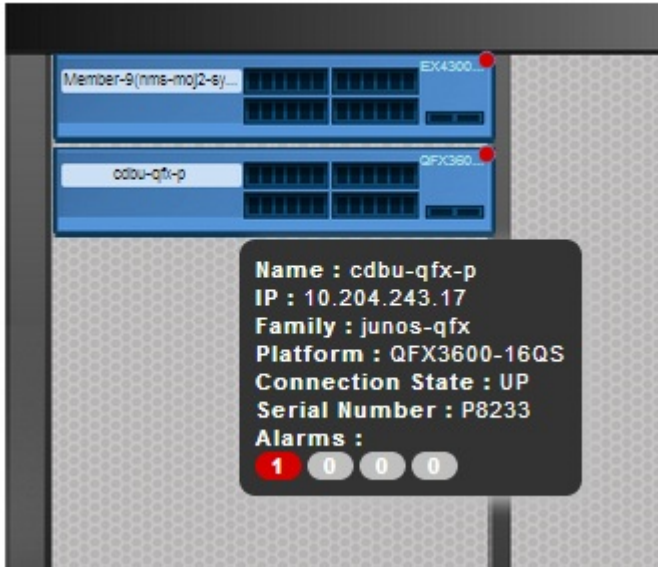
If the selected device is connected to more than sixty devices, then all the connected devices are highlighted in a circular form or a grid form. If the selected device is connected to less than 60 devices, then the links between the interconnected devices are displayed.

The device images are displayed along with details such as name, IP address, and alarm state information in colored labels that provide health and reachability information. You can also view the details of the hosts or virtual machines that are connected to the switches.

You can view the following details in the Device Connectivity - Graphical view:


- Name—The name of the device provided while configuring the device. The device name is displayed as a label.
- IP—The IP address of the device.
- Family—The family or platform to which the device belongs to. Family can be an JUNOS-EX, JUNOS-QFX, JUNOS-QF, MSSOS, and WLC-AP.
- Serial Number—Serial number of the device.
- Alarms—The alarm details displaying the number of critical, major, minor alarms, or info for the device. Alarms details are color coded to indicate their severity level as shown in [Figure 14](#). Network Director updates and displays the alarm status changes in real time in the Topology view.

Figure 14: Alarm Indicator



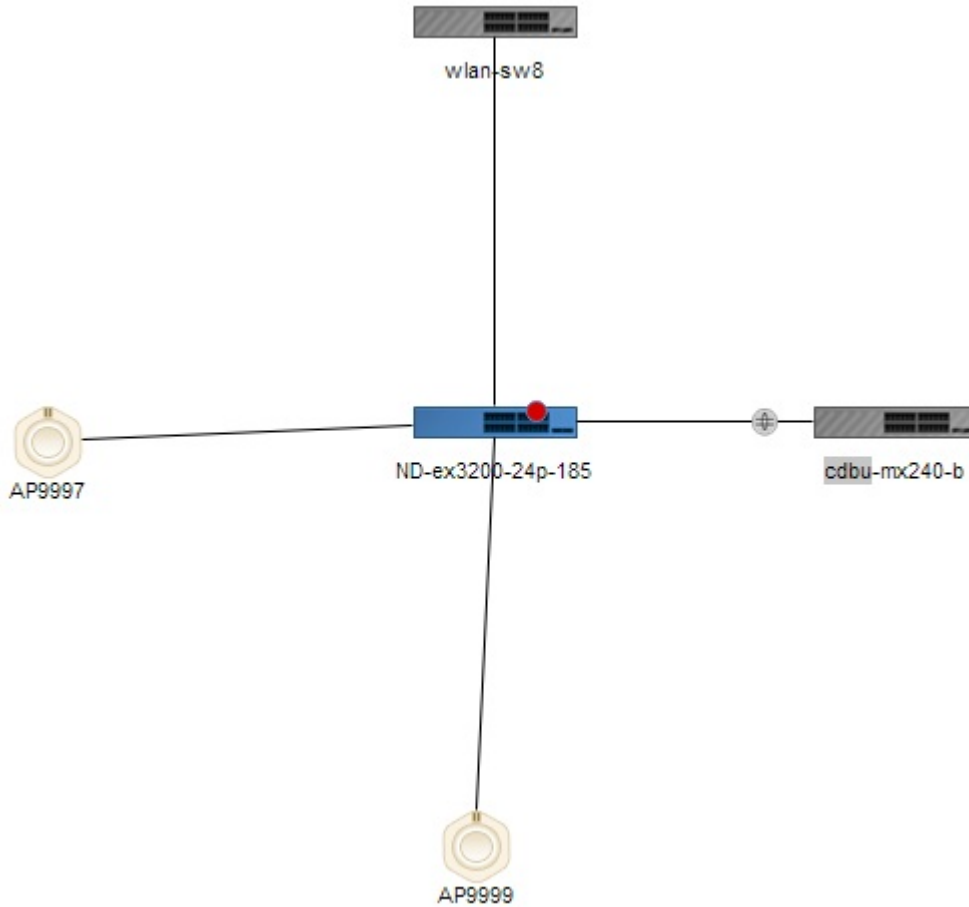
- Connection State—Connection status of the device. Connection state can be UP, DOWN or N/A. Network Director updates and displays the connection state changes in real time in the Topology view.
- Link status—Indicates whether the link between two devices is UP or DOWN as shown in [Figure 15](#). Network Director updates and displays the link status changes in real time in the Topology view.

Figure 15: Link Status Indicator



- LAG—Identifies connections that are configured as LAGs as shown in [Figure 16](#).

Figure 16: LAGs in the Device Connectivity view



You can view the following details of the virtual machine that are connected to hosts:

- Virtual Machine—Name of the virtual machine.
- Host Name—Name of the host to which the virtual machine is connected to.
- VNetwork—Name of the virtual network.
- OS—Name of the operating system on which the virtual machine is running.
- Connection State—Connection status of the virtual machine. Connection state can be UP, DOWN or N/A.
- Power State—State of the power supply: Powered On or Powered Off.

2. Click **Show Grid View** to view the device connectivity details in a tabular format as displayed in [Figure 17](#).

Figure 17: Displaying the Connection Details in Grid View

Device Connectivity : nd-72q1-ellit						
External Links Fabric Links						
Show Graph View						
Source Device	Source Port	Source Port Bandw...	Destination Device	Destination Port	Destination Port Bandw...	Link Status
nd-72q1-ellit	[LAG] ae0	NA	nd-36q1-ellit	[LAG] ae0	NA	UP
nd-72q1-ellit	[LAG] ae1	NA	nd-36q1-ellit	[LAG] ae1	NA	DOWN
nd-72q1-ellit	et-0/0/0 (ae0)	0	nd-36q1-ellit	et-0/0/0 (ae0)	0	UP
nd-72q1-ellit	et-0/0/1 (ae1)	0	nd-36q1-ellit	et-0/0/1 (ae1)	0	DOWN
nd-72q1-ellit	et-0/0/2 (ae2)	0	nd-opus-48s4	et-0/0/48	0	UP

The following details are displayed in the table:

- Source Node—Name of the device specified while configuring the device.
- Source Port—Source port of the device.
- Source Port Bandwidth %—Realtime percentage of bandwidth utilized at the source port.
- Destination Node—Name of the device or devices the device is connected to.
- Destination port—The port number on the destination device to which the source device is connected to.
- Destination Port Bandwidth %—Realtime percentage of bandwidth utilized at the destination port.
- Link Status—Indicates if the link to the device is UP or DOWN.

You can sort the details in the table in the ascending order or descending order for each column. You can also use filters to display only the desired device connectivity details.

Displaying QFabric Connectivity

You can view the QFabric devices available in your network by navigating to the rack, closet, or floors from the Topology View. Alternatively, you can search for the QFabric devices by entering a search criteria in the Search field. You can also assign the QFabric devices to a rack, closet, or floor from Topology View.

Network Director enables you to view the data plane topology and the control plane topology from the Fabric Connectivity page. Data plane is a redundant, high-performance, and scalable data plane that carries QFabric system data traffic. Control plane constitutes of the internal network connection that carries control traffic between the various QFabric system components such as node devices, interconnect devices, director group processes, and the control plane Ethernet switches.

You can view the connection details of the QFabric Interconnects with the nodes by using all views except the Dashboard View and Datacenter View.

To view the connection details of the QFabric devices:

Do one of the following:

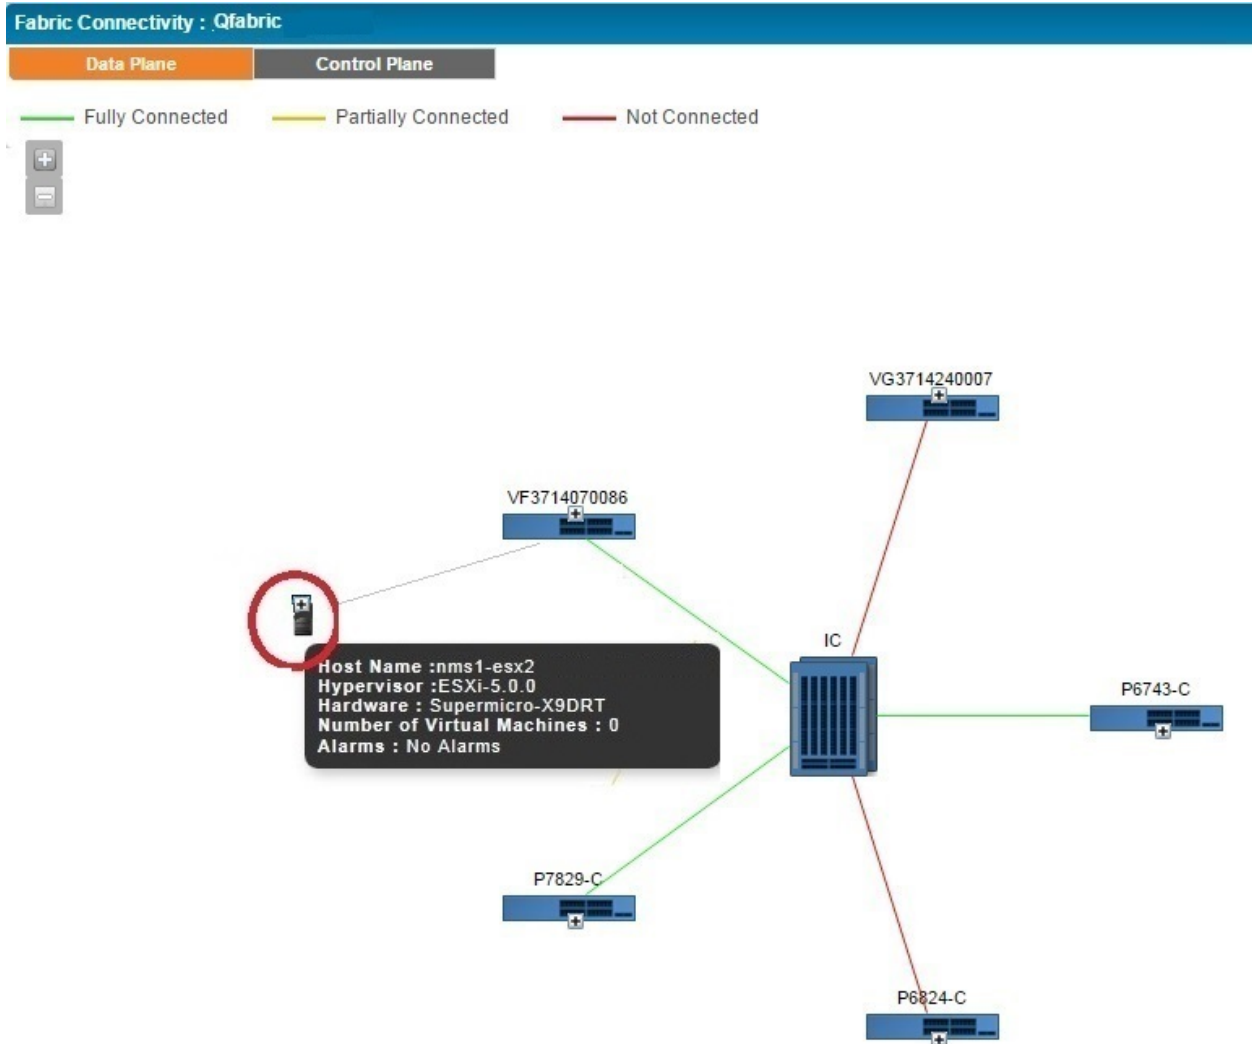
- While in the Logical, Location, Device, or Custom View, select the device for which you want to view the QFabric connectivity from the View pane and click **Connectivity > View QFabric Connectivity** from the Tasks pane.
- In the Topology View, navigate to a device in a building, floor, outdoor area, closet or rack and select the device for which you want to view the connectivity details and click **Connectivity > View QFabric Connectivity** from the Topology Tasks pane.

NOTE: The Connectivity task container is available only after you select a device.

Click **Connectivity > View QFabric Connectivity** from the Tasks pane. Network Director displays the QFabric connectivity in the Data plane—with the interconnect in the center surrounded by the nodes—as shown in [Figure 18](#). Mouse over a device to view the port details and the connection state of that device.

By default, Network Director displays the Data plane topology for the selected QFabric. To view the control plane topology, click **Control Plane**.

Figure 18: Displaying the Data Plane Connectivity for a QFabric Switch



In the [Figure 18](#), the connection details are represented by green, yellow, and red lines.

- The green line indicates that the node is connected to all the Interconnects properly and all functions are normal.
- The yellow line indicates that the node is only partially connected to the Interconnect. That is, the node might be connected to some of the interconnects, but not all.
- The red line indicates that the node is not connected to any of the interconnects.

The following details are displayed for each device, if the details are configured on the device:

- Name of the device provided while configuring the device. The device name is displayed as a label.
- Platform of the device. Platform can be QFX3500, QFX3600, or QFX5100 switches.
- Node group name. The name of the node group to which this device belongs to.

- Node group type. The server name of the node group to which this device belongs to.
- Source Port—source port of the device.
- Destination port—The port number on the destination device to which the source device is connected to.

The following details are displayed for each host that is connected to the member:

- Host Name—Name of the host machine connected to the member.
- Hypervisor—Name of the hypervisor, which is used to create multiple virtual machines on a hardware device.
- Hardware—Type of hardware on which the hypervisor is deployed.
- Number of Virtual Machines—Number of the virtual machines that are connected to the hosts.
- Alarms—Number of active alarms on the device.

Displaying Virtual Chassis and Virtual Chassis Fabric Connectivity

You can view the connectivity between the components of Virtual Chassis and Virtual Chassis Fabric using the View VC/VCF Connectivity task. You can access this tasks from all views except Dashboard View and Datacenter View.

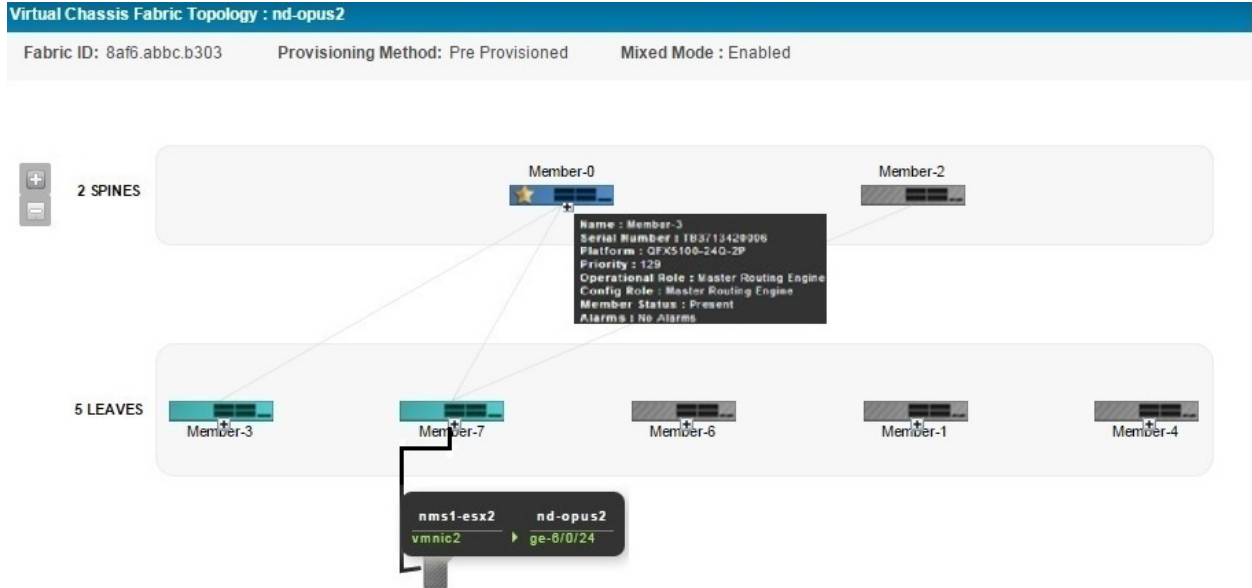
To view the connectivity details for a Virtual Chassis or a Virtual Chassis Fabric:

1. Do one of the following:
 - While in the Logical, Location, Device, or Custom View, select the device for which you want to view the Virtual Chassis or Virtual Chassis Fabric connectivity from the View pane and click **Connectivity > View VC/VCF Connectivity** from the Tasks pane.
 - In the Topology View, navigate to a device in a building, floor, outdoor area, closet or rack and select the device for which you want to view the Virtual Chassis or Virtual Chassis Fabric connectivity details and click **Connectivity > View VC/VCF Connectivity** from the Topology Tasks pane.

NOTE: The Connectivity task container is available only after you select a device.

2. Click **Connectivity > View VC/VCF Connectivity** from the Tasks pane. Network Director displays the connectivity between the members of the selected Virtual Chassis as shown in [Figure 19](#). The inactive members and any members having alarms in down state are shown as grey icons in the topology identifying the state of the member. Mouse over a member to view details of that member.

Figure 19: Displaying the Connectivity for a Virtual Chassis Fabric



The following details are displayed in the top panel of the Virtual Chassis Topology View as shown in Table 48.

Table 48: Common Details for the Virtual Chassis and Virtual Chassis Fabric

Details	Description
Fabric ID VC ID	All members of a Virtual Chassis configuration share one Virtual Chassis identifier (VCID). This identifier is derived from internal parameters.
Provisioning Method	<p>The provisioning mode of the member. Provision mode can be <i>autoprovisioned</i>, <i>preprovisioned</i> or <i>not preprovisioned</i>.</p> <p>Autoprovisioning a Virtual Chassis Fabric (VCF) enables you to “plug and play” devices into your VCF after minimal initial configuration.</p> <p>In a VCF, you can have two to four members configured in the Routing Engine role. Of this, one member acts as the master Routing Engine and another member acts as the backup Routing Engine. In a preprovisioned configuration, the selection of which member functions as the master Routing Engine and which as the backup Routing Engine is determined by the software based on the master-election algorithm.</p> <p>In a configuration that is not preprovisioned, the selection of the master and backup is determined by the mastership priority value and secondary factors in the master-election algorithm.</p>
VC Mode	Indicates whether the Virtual Chassis is mixed or not.

The following details are displayed in the Virtual Chassis Topology view for each member depending on the role of the member as shown in [Table 49](#).

Table 49: Details of VC/VCF Members

Details	Description	Role
Name	Name of the member switch provided while configuring the device. The device name is displayed as a label.	Master Backup Line Card
Serial Number	Serial number of the member switch.	Master Backup Line Card
Platform	Platform of the device. Platform can be QFX3500, QFX3600, QFX5100, or QFX5110.	Master Backup Line Card
Priority	The mastership priority value. This is the most important factor in determining the role of the member switch within the Virtual Chassis configuration.	Master Backup Line Card
Operational Role	Operational role of the device. A device might be configured for a particular role, but can operate in the same or a different role. For example, a spine device configured with a Routing Engine role might operate as a line card. Therefore, the operational role of this device is Line Card. Operational role can be Routing Engine or Line Card.	Master Backup Line Card
Config Role	The configured role of the device. This can be Routing Engine or Line card.	Master Backup Line Card
Member Status	Displays the status of each member device: <ul style="list-style-type: none"> ● Present—The device is connected and working fine. ● Not Present—The device is not connected to the VC or VCF. ● Inactive—The device is connected, but not running. ● Non Provisioned—A configuration in which the roles of the members are assigned automatically; not configured statically (preprovisioned). ● Pre Provisioned—A configuration that allows you to deterministically control the member ID and role assigned to a member by associating the member with its serial number. 	Master Backup Line Card

In addition to the above details, when you expand the host details, you can also view the details of the member and the link connected to the virtual machine.

Displaying Virtual Network Connectivity

To view the virtual network connectivity from Topology View, see [“Viewing the Virtual Machine Inventory in a Cloud Infrastructure” on page 803](#).

Displaying Third-Party Device Details

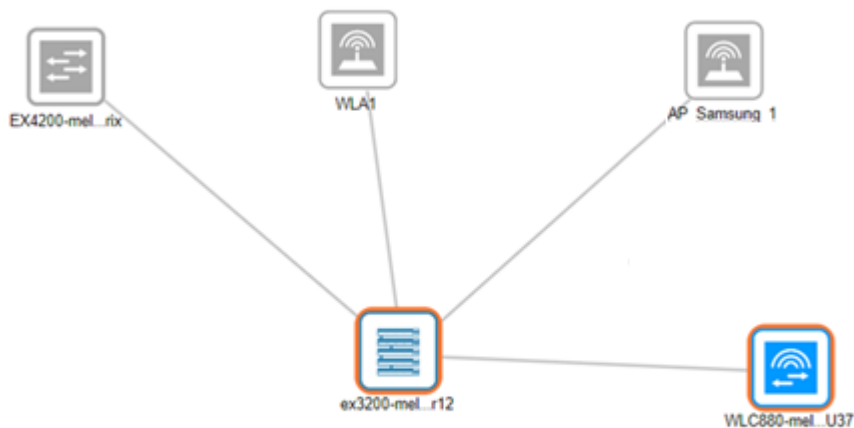
Network Director now uses a unique icon to depict Samsung Access Points (APs) in the Device Connectivity page. You can mouse over the Samsung AP icon to view details of the device.

To view the Samsung AP details:

1. From the View pane in the **Logical View**, **Location View**, **Device View**, **Custom View**, or **Topology View**, select the device for which you want to view the device connectivity.
2. From the **Tasks** pane, click **Connectivity > View Device Connectivity**.

The Device Connectivity page opens showing the connectivity between various devices as shown in [Figure 20](#).

Figure 20: Samsung AP Device Connectivity



3. To view the Samsung AP details, mouse over the Samsung AP icon.

A tool-tip, as shown in [Figure 21](#), is displayed showing the Samsung AP details listed in [Table 50](#).

Figure 21: Displaying Samsung AP Details

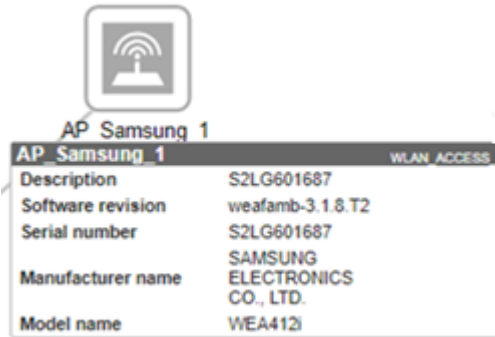


Table 50: Samsung AP Details

Field	Description
Name	The name of the device. The device name is displayed as a label.
Description	The description provided for the access point.
Software version	Software version that is running on the access point.
Serial number	Serial number of the access point.
Manufacturer name	Name of the access point manufacturer (Samsung Electronics CO., LTD.)
Model name	Model name of the access point.

NOTE: You can view the Samsung AP details that are connected to Juniper devices that runs Junos OS 12.1 version.

Uploading Floor Plans

You can upload the floor plan from the Topology View if you already have a floor plan for a specified building in a site.

To upload the floor plan:

1. Select a **Site > Building > Floor**. Alternatively, create a site, building, and floor. Click the **Upload Floor Plan** task from the Location task in the Tasks pane.

The Upload Floor Plan dialog box is displayed.

2. Click **Browse** next to Image File to choose a floor plan file.
3. Navigate to the folder where you have saved the floor plan on your system and click **Open**.
4. Click **Upload** to upload the floor plan image.
5. Click **Cancel** if you do not want to upload the floor plan and quit the Upload Floor Plan dialog box.

Uploading Topology Map

You can upload a topology map for an outdoor area from the Topology View.

To upload the map:

1. Select a **Site > Outdoor area**. Alternatively, create an outdoor area within a site. Click the **Upload map** task from the Location task in the Tasks pane.

The Upload Map dialog box is displayed.

2. Click **Browse** next to Image File to choose a map file.
3. Navigate to the folder where you have saved the topology map for an outdoor area on your system and click **Open**.
4. Click **Upload** to upload the topology map image.
5. Click **Cancel** if you do not want to upload the map and quit the Upload Map dialog box.

RELATED DOCUMENTATION

[Discovering Devices in a Physical Network | 204](#)

[Setting Up the Location View | 224](#)

[Understanding the Network Topology in Network Director | 243](#)

[Network Director Documentation home page](#)

Adding and Managing OUI Data in Network Director

Network Director uses Link Layer Discovery Protocol (LLDP) to detect the type of network device (such as desktop computer, VoIP phone, or network servers) in a campus network. However, network printers are an exception as most printers do not use LLDP. As a result, Network Director might not be able to identify the device as a printer. This is where the organizationally unique identifier (OUI) of printers come into play. OUI is formed using the first three octets of the device MAC address. OUI is unique to a vendor or manufacturer and can be identified globally.

You can build a database of OUIs that Network Director can use to identify the device type. During the topology discovery, if the LLDP-based discovery does not identify the device type, Network Director looks up in the OUI database to see whether there is a match. If Network Director finds a matching entry, the device is notated as a printer in the network topology. If there is no match, the device is marked as an unknown device.

To add and manage OUI data:

1. While in the Build mode with Logical View selected, click **Connectivity > Manage OUI** from the Tasks pane.

The Manage OUI page opens.

NOTE: Network Director displays a list of well-known printer manufacturers and their OUI details in the Manage OUI page. Make sure that the details that you want to add are not already listed before you proceed with adding OUI data for your device.

2. Click **Add** to add new OUI details to Network Director.

The Add MAC OUI window opens.

3. Enter the MAC OUI for the device that you want to add. The first three octets of a MAC address of the device forms the OUI. OUI is unique to a vendor or manufacturer and can be identified globally.
4. Enter the name of the vendor and select the type of device.
5. Click **Save** to save the OUI details and return to the Manage OUI page.
6. To delete one or more OUI details, select the rows that you want to delete and click **Delete**.

RELATED DOCUMENTATION

Creating Custom Device Groups

IN THIS CHAPTER

- [Understanding Custom Device Groups | 273](#)
- [Creating Custom Device Groups | 276](#)

Understanding Custom Device Groups

IN THIS SECTION

- [Where Is the Custom Group Function Located in Network Director? | 274](#)
- [How Do Custom Group Rules Work? | 274](#)
- [What Happens When I Edit a Custom Group Rule? | 276](#)
- [When Are Rules Executed? | 276](#)

Custom group is way of grouping your devices based on your business needs. You can create custom groups and add devices to each custom group. You can manually add devices to a custom group or you can define rules to add devices, that match the rule condition, to the custom group once they are discovered by Network Director. You can view the custom groups and devices that are assigned to each group in the Custom Group view.

A custom group can include devices such as switches, wireless controllers, and access points. Creating custom device groups enables the configuration of multiple devices simultaneously—you can create multiple custom groups and directly associate devices at any level. Up to this point, Custom Groups are the same as selecting related items in the location view tree. What makes Custom Groups unique is that you can also configure a custom group to automatically add devices after discovery. You indicate the criteria for additional devices by editing rules. Custom groups can then be created in a hierarchy up to eight levels deep. Each layer can contain up to 32 peer containers under a single parent container.

Where Is the Custom Group Function Located in Network Director?

Network Director has different views that you select to see different aspects of your data. You select one of these views at a time from the Select View option in the Network Director banner. The options are Logical View, Location View, Device View, Custom Group View, Datacenter View, and Topology View. To create a Custom Group, Network Director must be in Custom Group View. Custom Groups are created at the top level of the network—My Network.

Once Custom Groups are created, they appear in all views as options for profile assignment—assigning a profile to a Custom Group assigns that profile to all members of the group.

How Do Custom Group Rules Work?

Adding rules to a Custom group consists of creating a three part rule statement, with a rule basis, an operator, and matching criteria. Possible combinations are shown in [Table 51](#).

Table 51: Three Options of a Rule Statement

Rule Basis	Operator	Matching Criteria
Device Role	Equals	Access
		Aggregation
		Core
		Unassigned (available only for EX Series switches for which logical category can be defined)

Table 51: Three Options of a Rule Statement (continued)

Rule Basis	Operator	Matching Criteria
Device Type	Equals or Not Equals	Access Point
		QFabric Member
		QFabric
		Switch
		Virtual Chassis
		Wireless Controller
Serial Number	Equals or Contains	<i>You provide serial numbers or letters</i>
SKU or Model	Equals or Contains	<i>You provide model numbers or letters</i>
Management IP Address	Equals or Regex	<i>You provide IP address</i>
Location	<p>Select a previously configured location:</p> <p>NOTE: For location directions, see “Setting Up the Location View” on page 224.</p> <ol style="list-style-type: none"> 1. Click Please select. 2. From the Select Location window, select a location. 3. Click OK. 	
Device Role	Equals	<i>You provide preconfigured device role</i>
Device Type	Equals or Not Equals	Access Point QFabric Member QFabric Switch Virtual Chassis Wireless Controller
Firmware Version	Equals or Contains	<i>You provide a full or partial firmware version for devices</i>

What Happens When I Edit a Custom Group Rule?

When you edit a rule, devices that were added to the group but no longer qualify because of the rule edit are not automatically removed from the group. You must remove those devices manually. If more devices are now qualified to join the group because of your rule edit, the devices are added to the group on the next device notification change to the network.

When Are Rules Executed?

The option **Associate devices based on the rules for the custom groups while saving group information** is enabled by default. If the option is disabled, the rule engine will be activated only when there is some change in the device property. When a device property change occurs, rules are processed and devices are added to the group, if the group has a rule for those actions.

RELATED DOCUMENTATION

[Creating Custom Device Groups | 276](#)

[Network Director Documentation home page](#)

Creating Custom Device Groups

IN THIS SECTION

- [Creating Custom Groups | 277](#)
- [Creating a Custom Group | 277](#)

From Network Director, you can create a custom group, then add devices such as switches, wireless controllers, and access points to the group. Creating custom device groups enables the configuration of multiple devices simultaneously—you can also create multiple custom groups and directly associate devices at any level. Up to this point, Custom Groups behave the same way as selecting related items in the location view tree. What makes Custom Groups unique is that you can also configure a custom group to automatically add devices after discovery. You indicate the criteria for additional devices with rules. Custom groups can then be created in a hierarchy up to eight levels deep. Each layer can contain up to 32 peer containers under a single parent container.

NOTE: A device can be part of a group at only one level in a hierarchy.

This topic describes:

Creating Custom Groups

To create custom groups:

1. In the top banner, under **Views**, select **Custom Group View**.

2. Click  in the Network Director banner.

3. Click **Set Up Custom Group** under Key Tasks in the Tasks pane.

The Set Up Custom Group page opens, displaying a list of currently configured Custom Groups.

4. Configure the custom group, following the directions [“Creating a Custom Group” on page 277](#).

5. Click **Done**.

The new custom group appears in the Groups List.

Creating a Custom Group

Use the Set Up Custom Group page to define a group of devices that you can configure simultaneously.

To add a new custom group:

1. Type a Custom Group Name for the new group and then click **Add**.

The Custom Group tree is displayed with your new group added.

2. Click **Done** now to create the group with no child groups, devices, or rules. The Message *Data Saved Successfully* is displayed. Click **OK**.

For additional configuration, select your new group.

The options **Add Child Group**, **Assign Devices**, and **Add/Edit Rule** appear.

3. To add a child group under the new custom group:
 - a. Be sure the correct custom group is selected—this group will become the parent group.
 - b. Click **Add Child Group**.

The Add Child Group window opens, displaying a default child group name such as Group-0.
 - c. Replace the default child group name.
 - d. Click **Add**.

The new child group appears in the Custom Group list tree under the parent group.

TIP: Custom groups can be created in a hierarchy up to eight levels deep. Each layer can contain up to 32 peer containers under a single parent container.

4. To assign devices to a custom group:
 - a. Select a custom group, either a parent or child group, and then click **Assign Devices**.

The Assign Devices To Custom Group window opens, displaying a list of discovered network devices, their IP addresses, and their platforms. Platforms include junos-ex, junos-qfx, junos-qf, wlc-ap, and mssos. These are devices that can be added to the group.
 - b. Select one or more devices by adding a check mark and then click **Add**.

The devices are listed under the appropriate group in the Custom Groups List.

NOTE: A device can be part of a group at only one level in a hierarchy.

5. To add a rule that will automatically add devices to a parent or child custom group:
 - a. Select a custom group, either a parent or child group, that will have devices added to it automatically when a specific rule has been met.
 - b. Click **Add/Edit Rule(s)**.

The Add/Edit Rules window opens.
 - c. Click **Add Rule**.

A rule statement is displayed with three columns—two columns display the words *Please select...*. The third column is blank.

- d. From the first *Please select...* option in the rule statement, select the basis for the rule. You are indicating that automatic additions to the list will be based on either **Device Type**, **Firmware Version**, **Device Role**, **Serial Number**, **SKU/Model**, **Management IP**, or **Location**.
- e. From the second *Please select...* option in the rule statement, select an available operator, either **Equals**, **Not Equals**, **Like**, **Regex**, or **Contains**—the operators presented depend on the basis you selected in the first column. For example, if the basis for the rule is **SKU/Model**, then the only operator options are **Equals** and **Not Equals**. If the basis for the rule is **Location**, then your only option is to click **Select** for a list of locations.

TIP: The **Equals** operation matches all characters of the matching criteria. The **Like** operation matches the first few characters of the matching criteria.

- f. For the third option in the rule statement, provide a matching criteria. Matching criteria are indicated in the third column of the list shown in [Table 52](#).

TIP: Some rules have no third option.

Table 52: Three Options of a Rule Statement

Rule Basis	Operator	Matching Criteria
Device Role	Equals	Access
		Aggregation
		Core
		Unassigned (available only for EX Series switches for which logical category can be defined)

Table 52: Three Options of a Rule Statement (continued)

Rule Basis	Operator	Matching Criteria
Device Type	Equals or Not Equals	Access Point
		QFabric Member
		QFabric
		Switch
		Virtual Chassis
		Wireless Controller
Serial Number	Equals or Contains	You provide serial numbers or letters
SKU or Model	Equals or Contains	You provide model numbers or letters
Management IP	Equals or Regex TIP: Regex, a regular expression, consists of a sequence of characters that forms a search pattern.	You provide IP address or regular expression TIP: For example, (?<=\.){2,}(?=[A-Z]) is a regular expression.
Location	Select a previously configured location: NOTE: For directions to configure locations, see “Setting Up the Location View” on page 224 . i. Click <i>Please Select</i> . ii. From the Select Location window, select a location. iii. Click OK .	
Firmware Version	Equals or Contains	You provide a full or partial firmware version for devices.

- g. Click **OK**.

Rules are executed when new devices are discovered. Devices that match the defined rules are added to the group dynamically once discovery is complete.

TIP: If you add more than one rule to a Custom Group, then all rules must be met for a device to join the group.

6. The option **Associate devices based on the rules for the custom groups while saving group information** is enabled by default. When a device property change occurs, rules are processed and devices are added to the group, if the group has a rule for those actions. If you disable the option, the rule engine will be activated only when there is some change in the device property.
7. Click **Done**.
A status window opens with either the message *Data saved successfully* or with an error message. Click **OK**.
8. To edit a rule, select the appropriate custom group and then click **Add/Edit Rule**. When you edit a rule, devices in the group that no longer qualify because of the rule change are not automatically removed from the group. You must remove those devices manually. If more devices are now qualified to join the group because of your rule edit, the devices are added to the group on the next device notification change to the network.

TIP: To delete a device from the group, select the device and then click **Delete**. To delete an entire Custom Group, select the group and then click **Delete**. You are asked to confirm the deletion—click **OK**.

RELATED DOCUMENTATION

[Creating and Managing Port Groups | 495](#)

[Assigning a VLAN Profile to Devices or Ports | 531](#)

[Assigning Device Common Settings to Devices | 331](#)

[Assigning and Unassigning Devices to a Location | 234](#)

[Understanding Custom Device Groups | 273](#)

[Network Director Documentation home page](#)

Configuring Quick Templates

IN THIS CHAPTER

- Understanding Quick Templates | 282
- Configuring and Managing Quick Templates | 284

Understanding Quick Templates

Quick templates is a way to create a base build for devices. This feature enables you to use a CLI-based text editor to define your network configuration in the form of a template that you can apply to multiple devices in your network in addition to the profile assignment feature. Because quick templates are driven by Device Management Interface (DMI) schema, you can use them to set all the configuration parameters for any supported device.

By using these quick templates, you can configure, for example, routing protocols such as BGP, OSPF, ISIS, or even static routes by specifying the device configuration. You can append or add the system commands or the user-defined commands in the form of the variables in the CLI-based text editor. The user-defined commands support variables in the format `$(variable_name)`, which must be populated with data when you apply a template to a device.

The variable name defined for each CLI must be unique. Otherwise, you cannot assign different values to those variables even though they are used in different CLIs. For example, if a variable say `$(description)` is used in two CLIs `set vlans $(name) description $(description)` and `set snmp description $(description)`, you will not be able to define different values to the descriptions. To define different values, you must change the variable name for one of the commands.

The [Table 53](#) shows data types supported for the values entered for variables.

Table 53: Variable Data Types

Data Type	Description
Container	Holds other data types.
String	Contains character strings.

Table 53: Variable Data Types (continued)

Data Type	Description
Integer [Number]	Specifies a numeric value without a fractional component.
Boolean	Has two possible values: true and false. True if checked and False if unchecked.
Enumeration	Defines a variable to be a set of predefined constants. The variable is equal to one of the values that have been predefined for it.
Choice	Provides a radio button. Check the radio button to use the configuration option in the template.
String - Key [column in a table]	Identifies the uniqueness of the record in the table. If the table has a key specified , only one record with the given key could exist.

The Save option in the Create Quick Templates page enables you to save and also validate a template. If there are any conflicts in the configuration, you must resolve the conflicting variables in the configuration elements manually, before you deploy the configuration to the devices. Upon successful validation (and after you apply a template to a device), you can deploy the configurations (specified in the templates) to the devices. You can choose to deploy the configuration immediately, or at a later time. Depending upon the approval mode selected for your deployment, you can either deploy the changes directly or you can get an approval from the approver before deploying the changes. For more information about types of approval modes supported for deployments in Network Director, see [“Setting Up User and System Preferences” on page 108](#).

Benefits of Quick Templates

- Configuring a large number of devices can be tedious and time-consuming. Quick templates can apply necessary configurations on multiple devices at the same time, helping you save time and effort.
- Modifying configurations across multiple devices may lead to configuration errors. Deploying configuration using quick templates simplifies device configurations and reduces configuration errors.

RELATED DOCUMENTATION

[Setting Up User and System Preferences | 108](#)

[Deploying Configuration to Devices | 1180](#)

Configuring and Managing Quick Templates

IN THIS SECTION

- [Creating a Quick Template | 285](#)
- [Applying Templates to Devices | 286](#)
- [Editing a Quick Template | 287](#)
- [Deleting a Quick Template | 287](#)
- [Cloning a Quick Template | 287](#)
- [Using the Quick Template Details Window | 288](#)
- [Viewing Deployed Quick Templates | 288](#)

You can create and manage custom templates for your device configurations that are deployable through Network Director. Unlike other features that support implementation of only some of the device configurations, quick templates enables you to set up all the configuration parameters for any supported device because it is Device Management Interface (DMI) schema-driven.

Each device type is described by a unique data model that contains all the configuration data for that device. The Schema window shows the device family that you select while you create a template and the DMI schema that lists all the possible fields and attributes for a type of device. The latest schema describe the new features associated with recent device releases. After you create a quick template, you can add or delete device configuration details to and from quick templates by loading the configuration data from the schema. You need to apply these templates to devices manually.

If you click the **More tips** link you are guided on the variable and the command syntax usages. It also provides instructions on how to issue sub-commands. When defining your network configuration in quick templates by using a particular command, ensure that you define the sub-commands individually. Stating sub-commands as a single command causes errors. For example, the commands **set snmp location sunnyvale** and **set snmp contact admin@example.com** are valid when defined individually. However, if you combine these commands into the single command **set snmp location sunnyvale contact admin@example.com** schema validation treats the end command **contact** as an extra entry and throws an error.

To avoid any conflicts with the profile configurations while creating the template, a warning message **Please don't create any Profile conflict configuration** is displayed to indicate that you must not create a configuration as part of the template if the same configuration is available as part of the profile configuration.

The Templates page in the Quick Templates workspace lists the device templates created, in a tabular view. The [Table 54](#) lists the columns in the table along with a description:

Table 54: Quick Templates

Column	Description
Creation Time	Date and time when the template was created.
Template Name	Name of the quick template.
Device Family	Name of the device family for which the template is created. Selecting the option Common indicates that the template is applicable for all the device families.
OS Version	Junos OS version of the device family selected.
Description	Description of the quick template.
Last Updated Time	Date and time when the template was last modified.
Last Updated By	User name of the person who created the template.

This topic describes:

Creating a Quick Template

Quick templates enable you create a template to define configurations for your devices. You can create and deploy quick templates from the Wired workspace.

To create a quick template:

1. Click the Build Mode icon in the Network Director banner.
2. Select **Wired > Tasks > Manage Quick Templates** in the Tasks pane.
The Manage Quick Template page appears.
3. Click **Add**.
The Create Quick Template page opens.
4. Specify the following details:
 - Name—Type a name for the quick template. The quick template name is required. The quick template name must be unique and limited to 63 characters.
 - Description—Type a description for the quick template. The description is optional and limited to 255 characters.

- **Device Family**—From the Device Family list, select an appropriate device family. Selecting the option **Common** in device family creates a generic template, which can be applied to any device family. Therefore, specify only the most common settings such as system, SNMP, or track group settings that are applicable to all the platforms. If you want to define the settings that are specific to a platform select the appropriate platform from the device family instead of the Common option. For the list of device families supported by Network Director, see the latest [Network Director Release Notes](#).
- **OS Version**—From the OS Version list, select an appropriate DMI Schema version running on that platform. If you are unable to locate the DMI schema for a device family, you can update the DMI schema version on the Junos Space server. For more information about updating the DMI schema on the Junos Space server, see Junos Space documentation.

The Schema window displays the device family and the OS version selected in this step.

5. Type or paste the Junos commands in the form of variables in the CLI-based text editor provided in the CLI Commands section. For information on the type of supported variables, see [“Understanding Quick Templates” on page 282](#). Alternatively, you can navigate through the configuration option levels (at the left side) in Schema and double-click the configuration option you want to add to the quick template. The selected configuration option is displayed in the CLI Commands CLI-based text editor. The configuration options available here depend on the device family you selected.
6. Optionally, you can modify the configuration in the CLI Commands text area by using the tool bar functionalities such as undo, redo, cut, copy, paste, and find.
7. Click **Save**.

The template you created is displayed in the quick templates table.

Applying Templates to Devices

After you create a template, you can define your device configuration to be managed by using the quick templates, and apply these templates to the multiple devices.

To assign a template to a device:

1. Select the check box against the quick template for which you want to assign the profile.
2. Click **Assign**.
The Assign Quick Template : template names page opens.
3. Choose at least one device to which the profile needs to be assigned.
4. Click **Next**.

5. Choose a device and specify the quick template variables in Configure attributes page and click **Save**.
For example, when you configure a VLAN interface in a quick template, you can specify the variables VLAN and interface names for that template for a selected device.
6. Optionally, you can apply the settings specified here to all the selected devices of a device family by selecting the check box against the option **Apply above settings to all other selected devices**.
7. Click **Next** and then click **Finish**.
8. Review the profile association with the quick template and then click **Finish**.

Editing a Quick Template

You can edit a quick template to modify configurations for your devices.

To edit a quick template:

1. Select the check box against the quick template that you want to modify.
2. Click **Edit**.
The Edit Quick Template : template name page opens.

3. Make the required changes to the quick template and click **Save**.

Deleting a Quick Template

To delete a quick template:

1. Select the check box against the quick template that you want to delete.
2. Click **Delete**.
The Delete Quick Templates window opens.

3. Click **Yes** to delete the quick template; else click **No**.

Cloning a Quick Template

A cloned quick template is a copy of an existing quick template. You can use the quick template as a master copy to create clone of that template. When you clone a quick template, you create a copy of the entire device configuration, including its settings, and other contents. Cloning a quick template saves time if you

are deploying device configuration that are similar to the master copy, rather than creating a template and defining configurations multiple times.

To create a copy of an existing template:

1. Select the check box against the quick template you want to clone.
2. Click **Clone**.

The cloned template named master template-clone is shown in the list of templates.

Using the Quick Template Details Window

Use the Quick Template Details window to view the details of the quick template. [Table 55](#) describes the fields in this window.

Table 55: Quick Template Details

Field	Description
Name	Displays the name of the quick template.
Description	Provides a description of the quick template.
Device Family	Displays the device family for which quick template is created.
OS Version	Displays the Junos OS version for the selected device family.
CLI Commands	Displays the CLI commands configured for the device family.

Viewing Deployed Quick Templates

You deploy the device configurations defined in a quick template after you have applied the template to a device. The View Deployed Templates option enables an administrator or an operator to view the list of templates that are deployed to the devices.

You can mouse over the template name to view the date and time when the template was created and last modified.

The View Deployed Templates page lists the deployed templates device in a tabular view. The [Table 56](#) lists the columns in the table along with a description.

Table 56: View Deployed Template

Column	Description
Template Name	Indicates the name of the template whose configuration is deployed to the system.
Creation Time	Indicates the date and time when the template was created.
Last Updated Time	Indicates the date and time when the template was last modified.
User Name	Indicates the user name of the person who created the template.

Depending upon the type of approval mode configured—Manual Approval or Auto Approval mode— you can either deploy the device configurations defined in the template directly or by pursuing an approval from a configuration approver for the device changes.

To view the list of quick templates that are deployed to a device:

1. Click the Build Mode icon in the Network Director banner.

2. Select a device in the View pane.

The View Deployed Templates option appears under Wired>Tasks.

3. Click **View Deployed Templates**.

The Deployed Templates For Device: device name page displays listing the templates applied for that device.

RELATED DOCUMENTATION

[Understanding Quick Templates | 282](#)

[Deploying Configuration to Devices | 1180](#)

[Network Director Documentation home page](#)

Configuring Device Settings

IN THIS CHAPTER

- [Understanding Device Common Settings Profiles | 290](#)
- [Creating and Managing Device Common Settings | 291](#)
- [Assigning Device Common Settings to Devices | 331](#)

Understanding Device Common Settings Profiles

Network Director enables you to configure device-level settings for switches and wireless LAN controllers in the Device Common Settings profile. Once you create the profiles, you can assign the profiles to a switch or a controller and you can deploy the profiles using the Deploy mode tasks.

Network Director also creates Device Common Settings profiles when it discovers devices. It creates a Device Common Settings profile for each device it discovers, importing the device-level settings from the device into the profile.

While configuring the profiles, you can specify the basic settings, which includes the profile name, device user list, and time settings. Apart from the basic settings, you can optionally specify the management and protocol settings too.

RELATED DOCUMENTATION

[Creating and Managing Device Common Settings | 291](#)

[Assigning Device Common Settings to Devices | 331](#)

[Understanding Network Configuration Profiles | 197](#)

[Network Director Documentation home page](#)

Creating and Managing Device Common Settings

IN THIS SECTION

- [Managing Device Common Settings | 291](#)
- [Creating a Device Common Settings Profile | 293](#)
- [Specifying Basic Settings for Device Common Settings | 295](#)
- [Specifying Management Settings for EX Switching Device Common Settings | 298](#)
- [Specifying Management Settings for Wireless Device Common Settings | 302](#)
- [Specifying Management Settings for Campus Switching ELS Device Common Settings | 306](#)
- [Specifying Management Settings for Data Center Non-ELS Device Common Settings | 310](#)
- [Specifying Management Settings for Data Center ELS Device Common Settings | 314](#)
- [Specifying Protocol Settings for EX Switching Device Common Settings | 317](#)
- [Specifying DNS Settings for Wireless Device Common Settings | 321](#)
- [Configuring Wireless Dynamic Authorization Client \(DAC\) Settings | 321](#)
- [Specifying DHCP Relay/DNS Settings for Campus Switching ELS Device Common Settings | 322](#)
- [Specifying Protocol Settings for Data Center Switching Non-ELS Device Common Settings | 325](#)
- [Specifying Protocol Settings for Data Center Switching ELS Device Common Settings | 327](#)
- [Reviewing and Saving a Device Common Settings Configuration | 330](#)
- [What to Do Next | 330](#)

Use the Manage Device Common Settings page to create new device common settings for switching and wireless devices and to manage the existing device common settings.

This topic describes:

Managing Device Common Settings

From the Manage Device Common Settings page, you can:

- Create a new Device Common Settings profile by clicking **Add**. For directions, see [“Creating a Device Common Settings Profile” on page 293](#).
- Modify an existing Device Common Settings profile by selecting it and clicking **Edit**.
- Assign a Device Common Settings profile to a device by selecting a profile and clicking **Assign**. For directions, see [“Assigning Device Common Settings to Devices” on page 331](#).

- Modify an existing assignment of a Device Common Settings profile by selecting the profile and clicking **Edit Assignment**.
- View information about a Device Common Settings profile by either double-clicking the profile name or by selecting the profile and clicking **Details**.
- Delete a Device Common Settings profile by selecting a profile and clicking **Delete**.

TIP: You cannot delete common settings profiles that are in use—that is, assigned to devices or used by other profiles. To see the current assignments for a profile, select the profile and click **Details**.

- Clone a Device Common Settings profile by selecting a profile and clicking **Clone**.

Table 57 describes the device information available on the Manage Device Common Settings page. This page lists all Device profiles defined for your network, regardless of your current selected scope in the network view.

Table 57: Manage Device Common Settings Settings

Field Name	Action
Profile Name	Name given to the profile when the profile was created.
Family Type	The device family; EX Series switch, Campus Switching ELS, wireless LAN controller (WLC), Data Center Switching.
Description	Description of the Device profile entered when the profile was created.
Assignment State	Displays the assignment state of the profile. A profile can be: <ul style="list-style-type: none"> • Unassigned—When the profile is not assigned to any device • Deployed—When the profile is assigned to a device and is deployed from Deploy mode • Pending Deployment—When the profile is assigned to a device, but not yet deployed in the network. For deployment directions, see “Deploying Configuration to Devices” on page 1180.
Assigned to	Displays the number of devices to which the profile assignment is done.
Creation Time	Date and time when the profile was created.
Last Updated Time	Date and time when the profile was last modified.
User Name	The username of the person who created or modified the profile.

TIP: All columns might not be displayed. To show or hide fields listed in the Manage Authorization Profiles table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

Creating a Device Common Settings Profile

In Network Director, as an administrator, you can configure Device Common Settings profiles by using the Create Device Profile page for either switches or wireless LAN controllers. You can view the summary of the configurations before saving the Device profile.

At minimum, you must specify the Device profile and profile name in the workflow. You can include additional configuration such as:

- Device users
- Management services
- Multicast, spanning-tree protocol (STP)
- Domain Name Server
- DHCP servers, DHCP Relay servers, Login Banner, and Global PoE settings for switches

You can create profiles on the basis of the device family and each Device profile is specific to a device family. After you create a Device profile, you assign the profiles to different devices.

NOTE: You can assign only one profile to a device. However, you can assign the same profile to multiple devices.

To create a Device profile:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.

3. From the Tasks pane, select the type of network (Wired or Wireless), the appropriate functional area (Wired or Wireless), and select the name of the profile that you want to create. For example, to create a RADIUS profile for a wireless device, click **Wireless > Profiles > RADIUS**. The appropriate Manage Profile page opens.

4. Click **Add** to add a new profile.

If you chose to create a profile for the wired network, Network Director opens the Device Family Chooser window.

- a. From the Device Family Chooser, select the device family for which you want to create a profile. The available device families are **Switching (EX)**, **Campus Switching ELS** (Enhanced Layer 2 Software), **Data Center Switching Non-ELS** and **Data Center Switching ELS**.

- b. Click **OK**.

The Create Device Common Settings wizard for the selected device family is displayed. It consists of four sections, Basic Settings, Management Settings, Protocol Settings, and Review.

If you chose to create a profile for the wireless network, Network Director opens the Create Device Common Settings for Wireless wizard.

5. Specify the basic settings. Complete the Basic Setting wizard page as described in both the online help and in [“Specifying Basic Settings for Device Common Settings” on page 295](#).
6. When you have completed the basic settings, either click **Next** or click **Management Settings** at the top of the wizard window.
7. Complete the Management Settings described in both the online help and in the sections [“Specifying Management Settings for EX Switching Device Common Settings” on page 298](#), [“Specifying Management Settings for Wireless Device Common Settings” on page 302](#), [“Specifying Management Settings for Campus Switching ELS Device Common Settings” on page 306](#), [“Specifying Management Settings for Data Center Non-ELS Device Common Settings” on page 310](#) and [“Specifying Management Settings for Data Center ELS Device Common Settings” on page 314](#).
8. When you have completed the management settings, click **Next**.
9. Complete the protocol settings as described both online help and in the sections [“Specifying Protocol Settings for EX Switching Device Common Settings” on page 317](#), [“Specifying DNS Settings for Wireless Device Common Settings” on page 321](#), [“Specifying DHCP Relay/DNS Settings for Campus Switching ELS Device Common Settings” on page 322](#), [“Specifying Protocol Settings for Data Center Switching Non-ELS Device Common Settings” on page 325](#) and [“Specifying Protocol Settings for Data Center Switching ELS Device Common Settings” on page 327](#)

10. When you have completed the protocol settings, either click **Next** or click **Review** at the top of the wizard window.
11. You can either save your profile or make changes to your profile from the Review page. For more information, see [“Reviewing and Saving a Device Common Settings Configuration”](#) on page 330.
12. Click **Finish** to save the Device profile configuration.

The system saves the Device profile and displays the Manage Device Common Settings page. Your new or modified Device profile is listed in the table.

Specifying Basic Settings for Device Common Settings

To configure the basic settings for any Device Common Settings profile, enter the settings described in [Table 58](#). Mandatory settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 58: Device Profile Basic Settings

Field	Action
Profile Name	Type a name for the profile. You can use up to 32 characters for profiles created for wireless devices. Profile name must not contain special characters or spaces. Note that profiles that are automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character.
Description	Type a description of the profile containing up to 256 characters.
Login Banner for EX Series switches, Campus Switching ELS, and Data Center Switching	Enter the banner text—this text is displayed in the banner when you log in to the device.
Country Code for wireless LAN controllers only	Select the country code for the wireless LAN controllers. Country code settings are required on the primary wireless seed controller. TIP: Do not set the country code if you plan to provision the Device profile for active secondary and member nodes that will be part of a cluster.
AP Security Mode for wireless LAN controllers only	You can indicate that access point security is Required , Optional , or None .

Device Users

Table 58: Device Profile Basic Settings (continued)

Field	Action
Task: Add a Device User	<p>To add a device user:</p> <ol style="list-style-type: none"> Click Add under Device Users. The Add User window opens. Provide a username and password. Confirm the password. Enter a combination of 6 through 128 alphanumeric characters and special characters. The password is case sensitive and must be a combination of at least two different types of characters or a combination of upper case and lower case letters. TIP: Do not create a user with the name <i>root</i>. Select a role for the user: <ul style="list-style-type: none"> For switches, the role options are: Operator, Read-only, Super-user, or Unauthorized. Operators have clear, network, reset, trace, and view privileges. Super-Users have all privileges. For wireless controllers, the role options are: Framed, Administrative, or NAS-Prompt. Framed users have network user access only. Administrative users have access to the controller, including the enabled (configuration) mode. NAS-Prompt users have administrative access to the controller, excluding enabled mode. Click OK. The user is added to the list of Device Users. TIP: To edit an entry, select a row from the Device Users table and click Edit to modify the information. To delete an entry select a row from the Device Users table and click Delete to delete the user.
<h3>Time Settings</h3> <p>Time settings apply to all platforms. However, the setting for offset applies exclusively to wireless.</p>	
Time Zone	Select a country and time zone from the list. For wireless, you can also change the setting for Offset.

Table 58: Device Profile Basic Settings (continued)

Field	Action
Add a Time Server	<p>To add a time server:</p> <ol style="list-style-type: none"> 1. Click Add under Time Server. The Add Time Server window opens. 2. Provide an IP address and, optionally for switches only, mark the corresponding time server as Preferred. 3. Click OK. The server is added to the list of Time Servers. <p>TIP: To edit the settings of a time server, select it and then click Edit.</p>

To configure management settings, click **Next** or click **Management Settings** at the top of the wizard window. To skip the management settings and protocol settings, click **Review** at the top of the wizard window.

Management Settings are described in both the online help and in the sections [“Specifying Management Settings for EX Switching Device Common Settings” on page 298](#), [“Specifying Management Settings for Wireless Device Common Settings” on page 302](#), [“Specifying Management Settings for Campus Switching ELS Device Common Settings” on page 306](#), [“Specifying Protocol Settings for Data Center Switching Non-ELS Device Common Settings” on page 325](#) and [“Specifying Management Settings for Data Center Non-ELS Device Common Settings” on page 310](#).

Specifying Management Settings for EX Switching Device Common Settings

To configure the management settings for an EX switching Device profile:

1. Enter the settings described in [Table 59](#). All settings are optional. Default values are applied to the configuration if you skip the management settings configuration.

Table 59: Device Profile Management Settings for EX Switching

Task	Action
Enable Services	<p>You can enable one or more network protocol services for this Device profile: FTP, TELNET, HTTPS, or HTTP.</p> <p>NOTE: HTTP and HTTPS are not available for EX9200 Series switches.</p>
Configure PoE	<p>To add Power over Ethernet (PoE) configuration for EX switching, enable Configure PoE and provide these settings:</p> <p>NOTE: PoE configuration will be added only to switches that support PoE.</p> <ol style="list-style-type: none"> a. Using the arrows, adjust the Guard Band value from 0 through 19 watts. A guard band reserves a specified amount of power from the PoE power budget for the switch or line card in case of a spike in PoE consumption. For switches with multiple PoE line cards, such as the EX6200 switch, the guard band wattage is set to the specified value on all line cards, unless a line card has been explicitly configured with a different value. <p>TIP: The valid guard band range (in watts) for EX6200 and EX8200 switches is 0 through 15. Any value outside this range causes the profile deployment to fail.</p> b. Select a Management Mode for PoE, either Class or Static: <ul style="list-style-type: none"> • Class Management—In class PoE management mode, the maximum power for an interface is determined by the class of the connected powered device. • Static Management—In the static PoE management mode, you specify the maximum power for each PoE interface. The PoE controller then allocates this amount of power to the interface from its total budget. c. For PoE Global, you can indicate Enable All, Disable All, or None. <p>NOTE: If you deselect Configure PoE, PoE is disabled and the global PoE settings supported by this profile (poe guard-band, poe fpc all guard-band, poe management, poe fpc all management, and poe interface all) are deleted from the switch when the profile is deployed on the switch.</p>

Syslog Settings

Optionally, expand the Syslog Settings and provide the following system logging settings.

Table 59: Device Profile Management Settings for EX Switching (continued)

Task	Action
Enable Device Logging for Switches	<p>To enable device logging for switches:</p> <ol style="list-style-type: none"> a. Under Enable Device Log, click Add. The Add Log window opens. b. Select the log type for switching, either Console, File, User, or Host. <ul style="list-style-type: none"> • Console logging sends system log messages to the console. • File logging sends system log messages to the file you specify in File Name. • User logging sends system log messages to the terminal session of the user specified in User Name. You will also need to provide the name of the user. • Host logging sends system log messages to the server specified in Host. Host can be either an IP address or host name. c. Under Services, click Add. The phrase <i>Click to enter value</i> appears in both the Service column and Severity Filter column. d. Click the phrase <i>Click to enter value</i> in the Service column. A list box replaces the phrase in the Service column. e. From the Service list, select a logging service: Any, Authorization, Change-log, Conflict-log, Daemon, DFC, External, Firewall, FTP, Interactive-commands, Kernel, NTP, PFE, Security or User. f. Click the phrase <i>Click to enter value</i> in the Severity Filter column. A list box replaces the phrase in the Severity Filter column. g. Select an available severity filter from the list, either Alert, Any, Critical, Emergency, Error, Info, None, Notice, or Warning. The filter is added to the list of Severity Filters. The filter is activated when the corresponding service is triggered. h. Click OK. The log is added to the Enable Device Log list.
Edit Logging Settings	Select a Log Type from the Enable Device Log list and click Edit to change the configuration.

Table 59: Device Profile Management Settings for EX Switching (continued)

Task	Action
Delete Logging Settings	Select a Log Type from the Enable Device Log list and click Delete to remove the server configuration.

To configure protocol settings, either click **Next** or click **Protocol Settings**. To use the default protocol settings, skip to final review by clicking **Review** at the top of the wizard window.

Protocol Settings options are described in the section [“Specifying Protocol Settings for EX Switching Device Common Settings”](#) on page 317,

Specifying Management Settings for Wireless Device Common Settings

To configure the management settings for a wireless Device profile:

1. Enter the settings described in [Table 60](#). All settings are optional. Default values are accepted in the configuration settings if you skip the management settings options.

Table 60: Device Profile Management Settings for Wireless

Task	Action
Enable Services	You can enable one or more of the listed network protocol services for this Device profile: TELNET, HTTPS, HTTP, TFTP, and SSH . You can also change the default port numbers after clicking them.

Syslog Settings

Optionally, expand **Syslog Settings** and enable a device log and/or configure a syslog server.

Enable Device Logging for Wireless	<p>To enable device logging for wireless devices:</p> <ol style="list-style-type: none"> Under Enable Device Log, click Add. The Add Log window opens. Select a wireless log type, either Console, Session, or Trace. <ul style="list-style-type: none"> • Console logging sends system log messages to the console. • With session logging, each main session event—create, close, and deny—creates a log entry. • Trace operations record more detailed information about the operations, including packet forwarding and routing information. Select a Severity Filter for this log from the list, either Alert, Critical, Debug all, Emergency, Error, Info, Notice, or Warning. Click OK. The device log appears in the Enable Device Log list with its log type and severity filter. The filter will be activated when the corresponding service is triggered. <p>TIP: To edit logging settings, select an entry from the Enable Device Log table and click Edit. To delete a logging setting, select an entry from the Enable Device Log table and click Delete.</p>
------------------------------------	---

Table 60: Device Profile Management Settings for Wireless (continued)

Task	Action
Add a Syslog Server	<p>To add a syslog server to the common settings:</p> <ol style="list-style-type: none"> a. Under Syslog Servers, click Add. The Add Server window opens. b. Type the IP address of the server. c. Select a Severity Filter from the list, either Alert (default), Critical, Debug All, Emergency, Error, Info, Notice, or Warning. TIP: The filter will be activated when the corresponding service is triggered. d. Click OK. The server is added to the list of Syslog Servers. <p>TIP: To edit a syslog server, select a server and click Edit.</p> <p>To delete a syslog server, select a server from the Syslog Servers list, and then click Delete to remove the server.</p>

Service Settings

Optionally, expand **Service Settings** to configure the Web Portal setting to enable wireless WebAAA globally and/or provide the timeout settings for CLI Management sessions (SSH, Telnet, and Console).

Table 60: Device Profile Management Settings for Wireless (continued)

Task	Action
Web Portal (wireless only)	<p>Check Web Portal to enable a Web Portal for a wireless device. WebAAA provides a way to authenticate any user or device by using a Web browser. A common application of WebAAA is to control access for guests on your network. When a user requests access to an SSID or attempts to access a Web page before logging onto the network, MSS displays a login page in the user's browser. For more information, see "Understanding Web Portals" on page 905</p> <hr/> <p>SSL Mode: A Secure Socket Layer (SSL) connection means that any data that you send over the Internet is encrypted.</p> <ul style="list-style-type: none"> • None—Use no encryption for Web Portal. • Partial—Use SSL for Web Portal login but after successful authentication, users access the rest of the pages without SSL. • Full—Use SSL for all Web Portal communication. <hr/> <p>Force HTML: Check Force HTML to use HTML for all Web Portal pages. This requires users to login through a Web portal page. The current implementation of Web portal on the controller includes specific handling of Apple's iOS devices. When a request for this URL is sent and the requesting device does not get the anticipated response, the iOS device automatically opens a Web browser interface and then opens the Web Portal login page.</p>
Idle Time Out for Services	<p>Idle time out for service sessions such as Telnet and SSH. By default, MSS automatically terminates a console or a Telnet session that is idle for more than one hour. You can specify from 0 to 86400 seconds (one day). If you specify 0, the idle timeout is disabled. The timeout interval is in 30-second increments. For example, the interval can be 0, or 3 seconds, or 60 seconds, or 90 seconds, and so on. If you enter an interval that is not divisible by 30, the controller rounds up to the next 30-second increment.</p>
Console Time Out for Remote Connections	<p>Sets the timeout for the CLI console. You can specify from 0 to 86400 seconds (one day). If you specify 0, the console timeout is disabled.</p>

To configure wireless DNS settings, either click **Next** or click **DNS Settings**. To skip the DNS settings, click **Review** at the top of the wizard window.

DNS settings are described in the section ["Specifying DNS Settings for Wireless Device Common Settings" on page 321](#).

Specifying Management Settings for Campus Switching ELS Device Common Settings

To configure the management settings for an ELS campus switching device common setting profile:

1. Enter the settings described in [Table 61](#). All settings are optional—default values are applied to the configuration if you skip the management settings.

Table 61: Management Settings for ELS Switching Device Profile

Task	Action
Enable Services	You can enable one or more network protocol services for this Device profile: FTP, Telnet, HTTPS, or HTTP . By default, none are selected.
Configure PoE	<p>To add Power over Ethernet (PoE) configuration for ELS Switching, enable Configure PoE and provide these settings:</p> <p>NOTE: PoE configuration will be added only to switches that support PoE.</p> <ol style="list-style-type: none"> Using the arrows, adjust the Guard Band value from 0 through 19 watts. A guard band reserves a specified amount of power from the PoE power budget for the switch or line card in case of a spike in PoE consumption. For switches with multiple PoE line cards, such as the EX6200 switch, the guard band wattage is set to the specified value on all line cards, unless a line card has been explicitly configured with a different value. <p>TIP: The valid guard band rang (in watts) for EX6200 and EX8200 switches is 0 through 15. Any value outside this range causes the profile deployment to fail.</p> Select a Management Mode for PoE, either Class or Static: <ul style="list-style-type: none"> • Class Management—In class PoE management mode, the maximum power for an interface is determined by the class of the connected powered device. • Static Management—In the static PoE management mode, you specify the maximum power for each PoE interface. The PoE controller then allocates this amount of power to the interface from its total budget. For PoE Global, you can indicate Enable All, Disable All, or None. <p>NOTE: If you deselect Configure PoE, PoE is disabled and the global PoE settings supported by this profile (poe guard-band, poe fpc all guard-band, poe management, poe fpc all management, and poe interface all) are deleted from the switch when the profile is deployed on the switch.</p>

Syslog Settings

Optionally, expand the system logging section and configure device logging.

Table 61: Management Settings for ELS Switching Device Profile (continued)

Task	Action
Enable Device Logging for ELS Switches	<p>To enable device logging for ELS switches:</p> <ol style="list-style-type: none"> a. Under Enable Device Log, click Add. The Add Log window opens. b. Select the log type for ELS switching, either Console, File, User, or Host (default). <ul style="list-style-type: none"> • Console logging sends system log messages to the console. • File logging sends system log messages to the file you specify for File Name. • User logging sends system log messages to the terminal session of the user you specify for User Name. You will also need to provide the name of the user. • Host logging sends system log messages to the server you specify for Host. Host can be either an IP address or host name. c. Under Services, click Add. The phrase <i>Click to enter value</i> appears in both the Service column and Severity Filter column. d. Click on the phrase <i>Click to enter value</i> in the Service column. A list box replaces the phrase in the Service column. e. From the Service list, select a logging service: Any, Authorization, Change-log, Conflict-log, Daemon, DFC, External, Firewall, FTP, Interactive-commands, Kernel, NTP, PFE Security or User. f. Click on the phrase <i>Click to enter value</i> in the Severity Filter column. A list box replaces the phrase in the Severity Filter column. g. Select a Severity Filter from the list, either Alert, Any, Critical, Emergency, Error, Info, Notice, or Warning. h. Click OK. The filter is added to the list of Enabled Device Logs with entries in the Log Type column and filter name column. The filter will be activated when the corresponding log type is triggered.
Task: Edit Logging Settings	Select an entry from the Enable Device Log table and click Edit to change the settings.

Table 61: Management Settings for ELS Switching Device Profile *(continued)*

Task	Action
Task: Delete Logging Settings	Select an entry from the Enable Device Log table and click Delete to remove the server settings.

To configure DHCP Relay and DNS, either click **Next** or click **DHCP Relay/DNS Settings**. To skip the protocol settings, click **Review** at the top of the wizard window.

DHCP Relay and DNS options are described in the section [“Specifying DHCP Relay/DNS Settings for Campus Switching ELS Device Common Settings”](#) on page 322.

Specifying Management Settings for Data Center Non-ELS Device Common Settings

To configure the management settings for a Data Center Management Non-ELS Device profile:

1. Enter the settings described in [Table 62](#). All settings are optional—default values are applied to the configuration if you skip the management settings.

Table 62: Device Profile Management Settings for Data Center Non-ELS

Field	Action
Enable Services	You can enable one or more network protocol services for this Device profile: <ul style="list-style-type: none">• For Data Center Switching, FTP, Telnet, HTTP.

System Logging Settings

Optionally, expand the Syslog Settings and provide the following system logging settings:

Table 62: Device Profile Management Settings for Data Center Non-ELS (continued)

Field	Action
Enable Device Logging for Switches	<p>To enable device logging for switches:</p> <ol style="list-style-type: none"> a. Under Enable Device Log, click Add. The Add Log window opens. b. Select the log type for switching: Console, File, User, or Host. <ul style="list-style-type: none"> • Console logging sends system log messages to the console. • File logging sends system log messages to the file you specify in File Name. • User logging sends system log messages to the terminal session of the user specified in User Name. You will also need to provide the name of the user. • Host logging sends system log messages to the server specified in Host. Host can be either an IP address or host name. c. Under Services, click Add. The phrase <i>Click to enter value</i> appears in both the Service column and Severity Filter column. d. Click on the phrase <i>Click to enter value</i> in the Service column. A list box replaces the phrase in the Service column. e. From the Services list, select a logging service: any, Authorization, Change-log, Conflict-log, Daemon, DFC, External, Firewall, FTP, Interactive-commands, Kernel, NTP, PFE or Security. f. Click on the phrase <i>Click to enter value</i> in the Severity Filter column. A list box replaces the phrase in the Severity Filter column. g. Select a Severity Filter from the list: Alert, Critical, Debug All, Emergency, Error, Info, Notice, or Warning. The filter is added to the list of Severity Filters. The filter is activated when the corresponding service is triggered. h. Click OK.
Edit Logging Settings	Select a Log Type from the Enable Device Log table and click Edit to change the information.

Table 62: Device Profile Management Settings for Data Center Non-ELS (continued)

Field	Action
Delete Logging Settings	Select a Log Type from the Enable Device Log table and click Delete to remove the server information.

To configure protocol settings, either click **Next** or click **Protocol Settings**. To skip the protocol settings, click **Review** at the top of the wizard window.

Protocol Settings options are described in the section [“Specifying Protocol Settings for Data Center Switching Non-ELS Device Common Settings”](#) on page 325.

Specifying Management Settings for Data Center ELS Device Common Settings

To configure the management settings for a Data Center ELS Device profile:

1. Enter the settings described in [Table 63](#). All settings are optional—default values are applied to the configuration if you skip the management settings.

Table 63: Device Profile Management Settings for Data Center ELS

Field	Action
Enable Services	You can enable one or more network protocol services for this Device profile:

System Logging Settings

Optionally, expand the Syslog Settings and provide the following system logging settings:

Table 63: Device Profile Management Settings for Data Center ELS (continued)

Field	Action
Enable Device Logging for Switches	<p>To enable device logging for switches:</p> <ol style="list-style-type: none"> a. Under Enable Device Log, click Add. The Add Log window opens. b. Select the log type for switching: Console, File, User, or Host. <ul style="list-style-type: none"> • Console logging sends system log messages to the console. • File logging sends system log messages to the file you specify in File Name. • User logging sends system log messages to the terminal session of the user specified in User Name. You will also need to provide the name of the user. • Host logging sends system log messages to the server specified in Host. Host can be either an IP address or host name. c. Under Services, click Add. The phrase <i>Click to enter value</i> appears in both the Service column and Severity Filter column. d. Click on the phrase <i>Click to enter value</i> in the Service column. A list box replaces the phrase in the Services column. e. From the Services list, select a logging service: any, Authorization, Change-log, Conflict-log, Daemon, DFC, External, Firewall, FTP, Interactive-commands, Kernel, NTP, PFE or Security. f. Click on the phrase <i>Click to enter value</i> in the Severity Filter column. A list box replaces the phrase in the Severity Filter column. g. Select a Severity Filter from the list: Alert, Critical, Debug All, Emergency, Error, Info, Notice, or Warning. The filter is added to the list of Severity Filters. The filter is activated when the corresponding service is triggered. h. Click OK.
Edit Logging Settings	Select a Log Type from the Enable Device Log table and click Edit to change the information.

Table 63: Device Profile Management Settings for Data Center ELS (continued)

Field	Action
Delete Logging Settings	Select a Log Type from the Enable Device Log table and click Delete to remove the server information.

To configure protocol settings, either click **Next** or click **DHCP/DNS Settings**. To skip the DHCP/DNS settings, click **Review** at the top of the wizard window.

DHCP/DNS Settings options are described in the section [“Specifying Protocol Settings for Data Center Switching ELS Device Common Settings”](#) on page 327.

Specifying Protocol Settings for EX Switching Device Common Settings

To configure the protocol settings for an EX Switching Device profile, enter the settings described in [Table 64](#). All settings are optional.

Table 64: Device Profile Protocol Settings for EX Switching

Field	Action
Enable Storm Control	
	Select this option to enable storm control on a switch.
Spanning Tree Settings	

Table 64: Device Profile Protocol Settings for EX Switching (continued)

Field	Action
Spanning Tree Protocol Settings for switches only	<p>Select one of spanning-tree protocol (STP) settings for switches: STP, RSTP (default), MSTP, or None of these.</p> <ul style="list-style-type: none"> Spanning Tree Protocol—With STP configured, the switches use the IEEE 802.1D 2004 specification, force version 0. This configuration runs a version of RSTP that is compatible with classic, basic STP as defined in the 802.1D 1998 specification. Rapid Spanning Tree Protocol—RSTP provides faster reconvergence time than the original STP both by identifying certain links as point-to-point and by using protocol handshake messages rather than fixed timeouts. VLAN Spanning Tree Protocol (VSTP) and RSTP can be configured concurrently. You can selectively configure up to 253 VLANs by using VSTP; the remaining VLANs will be configured by using RSTP. VSTP and RSTP are the only spanning-tree protocols that can be configured concurrently on a switch. Multiple Spanning Tree Protocol—MSTP enables multiple VLANs to be mapped to the same spanning-tree instance, reducing the number of spanning-tree instances needed to support a large number of VLANs. MSTP provides multiple forwarding paths for data traffic and enables load-balancing. It improves the fault tolerance of the network because a failure in one instance, or forwarding path, does not affect other instances. <p>You can also select the Enable VSTP check box to enable VSTP.</p>
Multicast Settings	
Enable IGMP	Selecting this option enables Internet Group Management Protocol (IGMP) on all the interfaces for the selected device. Default is disabled. IGMP is a communications protocol used by both hosts and adjacent routers on IP networks to establish multicast group memberships.
Enable IGMP Snooping	Enables IGMP snooping on all VLANs. Default is enabled.
Enable DHCP Relay	
Select this option to display the DHCP Relay settings.	

Table 64: Device Profile Protocol Settings for EX Switching (continued)

Field	Action
Add DHCP Relay to Device Profile	<p data-bbox="557 331 998 363">To add DHCP Relay to this Device profile:</p> <ol data-bbox="557 394 1349 829" style="list-style-type: none"><li data-bbox="557 394 976 426">1. Select Legacy DHCP Relay (default).<li data-bbox="557 478 1349 829">2. Add one or more DHCP servers to the Device Common Settings profile:<ol data-bbox="597 531 966 829" style="list-style-type: none"><li data-bbox="597 531 966 615">a. Click Add under DHCP Servers. The Add Server window opens.<li data-bbox="597 667 846 699">b. Type an IP Address.<li data-bbox="597 751 735 783">c. Click OK. The server is added to the list of DHCP Servers.

Table 64: Device Profile Protocol Settings for EX Switching (continued)

Field	Action
Add Extended DHCP Relay to a Device Profile	<p>To add Extended DHCP Relay to this Device profile:</p> <ol style="list-style-type: none"> 1. Select Extended DHCP Relay instead of Legacy DHCP Relay. 2. Add one or more DHCP Server Groups to the Device Common Settings profile: <ol style="list-style-type: none"> a. Click Add under Add DHCP Servers Group. The Add Server Group window opens. b. Provide a name for the server group. c. Optionally, make this an active server group by checking Active Group. d. Add servers to the group by clicking Add under DHCP Servers. The phrase <i>Click to enter value</i> appears in the IP Address column. e. Select <i>Click to enter value</i> and then enter an IP Address. f. Click OK. The server is added to the DHCP server group list. g. Add a relay interface group by clicking Add under Add Relay Interface Group. The Add DHCP Relay Interface window opens. h. Type a DHCP interface group name. i. Select a server group from the Server Group list. j. Click OK. The group is added to the Relay Interface Group list.

Click either **Next** or **Review**, to see the Review page. For review directions, see [“Reviewing and Saving a Device Common Settings Configuration” on page 330.](#)

Specifying DNS Settings for Wireless Device Common Settings

To configure the DNS settings for a wireless Device profile, enter the settings described in [Table 65](#). All settings are optional.

Table 65: Device Profile DNS Settings for Wireless

Task	Action
DNS Settings	
Add a Domain Name Server	<p>To add a Domain Name Server (DNS) for wireless common settings:</p> <ol style="list-style-type: none"> 1. Click Add under Domain Name Servers. The Add Server window opens. 2. Enter the IP address of the DNS server. 3. Select a preference for the server, either Primary (default) or Secondary. TIP: You can add only one primary IP address, but can add several secondary IP addresses. 4. Click OK. The IP address and preference are now listed in the list of Domain Name Servers. <p>TIP: To modify an IP address, select it from the Domain Name Servers table and click Edit.</p> <p>To delete an IP address, select it from the Domain Name Servers table and click Delete.</p>

Click either **Next** or **DAC Settings** to see the next section of the wizard. For Dynamic Authorization Client configuration directions, see [“Configuring Wireless Dynamic Authorization Client \(DAC\) Settings”](#) on [page 321](#).

Configuring Wireless Dynamic Authorization Client (DAC) Settings

Dynamic Authorization Client (DAC) is a dynamic RADIUS extension that enables administrators supporting a RADIUS server to disconnect a user and change the authorization attributes of an existing user session. The DAC is the component sending the Disconnect and CoA requests to the Dynamic Authorization Server (DAS). Though the DAC often resides on the RADIUS server, it can be located on a separated host, such as a Routing Engine.

Table 66: Dynamic Authorization Client (DAC) Settings

Field	Directions
RADIUS DAS Port	The Dynamic Authorization Server (DAS) is the component residing on the NAS that processes the Disconnect and Change-of-Authorization (CoA) requests sent by the Dynamic Authorization Client (DAC). The Dynamic Authorization Server Port is the UDP where the DAS listens for Disconnect and CoA requests sent by the DAC. Default port used is 3799.
Task: Add a dynamic authorization client.	<p>Provide the following settings:</p> <ul style="list-style-type: none"> ● DAC Name— Type a name for the dynamic authorization client. ● IP Address—Type the IP Address of the dynamic authorization client. ● Key—Enter the authentication key used to communicate with the RADIUS server. ● Disconnect—When checked (default), a terminated client on the network is disconnected. When unchecked, a terminated client is re-authenticated. ● Change of Authorization—Request packets contain information for dynamically changing session authorizations. Typically, this is used to change data filters. The data filters can be of either the ingress or egress kind, and are sent in addition to the NAS and Session identification attributes ● Replay Protection—Drop the request with out-of-change or no timestamp. ● Replay Window—Maximum seconds that local and client times can differ. ● Wired Access Rule—four wired rule names for RADIUS DAC. ● SSID Selection—Indicate which SSIDs the client is allowed to access after authentication. Move up to four of the available SSIDs from the list on the left to the list on the right. <p>Click OK to add the client and close the Add DAC window.</p>
Task: Edit a dynamic authorization client	Select a client from the list and then click Edit . Make changes and then click OK .
Task: Delete a dynamic authorization client	Select a client from the list and then click Delete and then click OK .

Specifying DHCP Relay/DNS Settings for Campus Switching ELS Device Common Settings

To configure the DHCP relay and DNS settings for a Campus Switching ELS Device profile, enter the settings described in [Table 67](#). All settings are optional.

Table 67: Device Profile Protocol Settings for ELS Switching

Task	Action
<p>DHCP Relay</p> <p>DHCP relay enables a switch to relay DHCP requests from a client to a DHCP server when the client and server do not reside on the same VLAN. You define the client interfaces for DHCP relay as part of the process of assigning the profile to a device. Select Enable DHCP Relay to enable DHCP Relay and view the DHCP Relay configuration.</p>	
<p>Configure Legacy DHCP Relay for ELS Switches</p>	<p>To configure Legacy DHCP Relay for Campus Switching ELS:</p> <ol style="list-style-type: none"> 1. Select Legacy DHCP Relay (default). 2. Add DHCP servers for Legacy DHCP Relay: <ol style="list-style-type: none"> a. Click Add under DHCP Servers. The Add Server window opens. b. Enter the IP address of the DHCP server and then click OK. The DHCP server name appears in the list of DHCP servers. <p>TIP: You can add more than one DHCP server.</p>

Table 67: Device Profile Protocol Settings for ELS Switching (continued)

Task	Action
Configure Extended DHCP Relay for Campus Switching ELS	<p>To add Extended DHCP Relay to this Campus Switching ELS Device profile:</p> <ol style="list-style-type: none"> 1. Select Extended DHCP Relay. 2. Add one or more DHCP Servers Groups to the Device Common Settings profile: <ol style="list-style-type: none"> a. Click Add under Add DHCP Servers Group. The Add Server Group window opens. b. Provide a name for the server group. c. Optionally, make this an active server group by checking Active Group. d. Add servers to the group by clicking Add under DHCP Servers. The phrase <i>Click to enter value</i> appears in the IP Address column. e. Select <i>Click to enter value</i> and then enter an IP Address. f. Click OK. The server is added to the DHCP server group list. g. Add a relay interface group by clicking Add under Add Relay Interface Group. The Add DHCP Relay Interface window opens. h. Type a group name for the DHCP interface. i. Select a server group from the Server Group list. j. Click OK. The group is added to the Relay Interface Group list.

DNS Settings

Table 67: Device Profile Protocol Settings for ELS Switching (continued)

Task	Action
Add a domain name server	<p>To add a domain name server to the Campus Switching ELS Common Settings:</p> <ol style="list-style-type: none"> 1. Click Add under Domain Name Servers. The Add Server window opens. 2. Provide an IP address for the DNS server. 3. Click OK. The server is added to the Domain Name Servers list. <p>TIP: To edit a DNS server's settings, select it and then click Edit. To delete a DNS server, select it and then click Delete.</p>

Click either **Next** or **Review**, to see the Review page. For review directions, see ["Reviewing and Saving a Device Common Settings Configuration" on page 330](#).

Specifying Protocol Settings for Data Center Switching Non-ELS Device Common Settings

To configure the protocol settings for a Device profile, enter the settings described in [Table 68](#). All settings are optional.

Table 68: Device Profile Protocol Settings for Data Center Switching Non-ELS

Field	Action
Enable Storm Control	
	Select this option to enable storm control on a switch.

Table 68: Device Profile Protocol Settings for Data Center Switching Non-ELS (continued)

Field	Action
Spanning Tree Protocol Settings	<p>Select one of spanning-tree protocol (STP) settings for switches: STP, RSTP, MSTP, or None of these.</p> <ul style="list-style-type: none"> • Spanning Tree Protocol—With STP configured, the switches use the IEEE 802.1D 2004 specification, force version 0. This configuration runs a version of RSTP that is compatible with classic, basic STP as defined in the 802.1D 1998 specification. • Rapid Spanning Tree Protocol—RSTP provides faster reconvergence time than the original STP both by identifying certain links as point-to-point and by using protocol handshake messages rather than fixed timeouts. VLAN Spanning Tree Protocol (VSTP) and RSTP can be configured concurrently. You can selectively configure up to 253 VLANs using VSTP; the remaining VLANs will be configured using RSTP. VSTP and RSTP are the only spanning-tree protocols that can be configured concurrently on a switch. • Multiple Spanning Tree Protocol—MSTP enables multiple VLANs to be mapped to the same spanning-tree instance, reducing the number of spanning-tree instances needed to support a large number of VLANs. MSTP provides multiple forwarding paths for data traffic and enables load-balancing. It improves the fault tolerance of the network because a failure in one instance, or forwarding path, does not affect other instances. <p>Select the Enable VSTP check box to enable VSTP.</p>
Multicast Settings	<p>Select a multicast setting for switches:</p> <ul style="list-style-type: none"> • Enable IGMP—Enables IGMP on all the interfaces for the selected device. (Default is disabled.) • Enable IGMP Snooping—Enables IGMP snooping (monitoring) on all VLANs. (Default is enabled.)
DCBX Settings	<p>Select the Data Center Bridging Capability Exchange (DCBX) protocol features that you want to enable:</p> <ul style="list-style-type: none"> • Enable DCBX—DCBX is a discovery and exchange protocol for conveying configuration and capabilities among neighbors to ensure consistent configuration across the network. It is an extension of Link Layer Discovery Protocol (LLDP). • Enable LLDP Snooping—LLDP is a discovery and exchange protocol for conveying configuration and capabilities among neighbors to ensure consistent configuration across the network. Snooping is monitoring traffic, in this case LLDP traffic.
Enable DHCP Relay	<p>Select Enable DHCP Relay to enable DHCP relay. DHCP relay enables a switch to relay DHCP requests from a client to a DHCP server when the client and server do not reside on the same VLAN. You will define the client interfaces for DHCP relay as part of the process of assigning the profile to a device.</p>

Table 68: Device Profile Protocol Settings for Data Center Switching Non-ELS (continued)

Field	Action
Add DHCP Relay Servers	<p>To add DHCP servers:</p> <ol style="list-style-type: none"> 1. Click Add. The Add Server window opens. 2. Enter the IP address of the DHCP server and then click OK. The DHCP server name appears in the list of DHCP servers. 3. You can enter more than one DHCP server. <p>TIP: Select an IP Address from the DHCP Servers table and click Edit if you want to modify the IP address. Select an IP Address from the DHCP Servers table and click Delete to remove from the table.</p>
DNS Settings	<p>To configure DNS:</p> <ol style="list-style-type: none"> 1. Click Add to add a DNS server. The Add Server window opens. 2. Enter the IP address of the DNS server. 3. Click OK. <p>TIP: Select an IP Address from the Domain Name Servers table and click Edit if you want to modify the IP address. Select an IP Address from the Domain Name Servers table and click Delete to remove from the table.</p>

Click either **Next** or **Review**, to see the Review page. For review directions, see [“Reviewing and Saving a Device Common Settings Configuration” on page 330](#).

Specifying Protocol Settings for Data Center Switching ELS Device Common Settings

To configure the protocol settings for a Data Center Switching ELS Device profile, enter the settings described in [Table 69](#). All settings are optional.

Table 69: Device Profile Protocol Settings for Data Center Switching ELS

Task	Action
DCBX Settings	
Select the Data Center Bridging Capability Exchange (DCBX) protocol features that you want to enable	
Enable DCBX	Select Enable DCBX . DCBX is a discovery and exchange protocol for conveying configuration and capabilities among neighbors to ensure consistent configuration across the network. It is an extension of Link Layer Discovery Protocol (LLDP).
Enable LLDP	Select Enable LLDP . LLDP is a discovery and exchange protocol for conveying configuration and capabilities among neighbors to ensure consistent configuration across the network.
DHCP Relay	
DHCP relay enables a switch to relay DHCP requests from a client to a DHCP server when the client and server do not reside on the same VLAN. You define the client interfaces for DHCP relay as part of the process of assigning the profile to a device. Select Enable DHCP Relay to enable DHCP Relay and view the DHCP Relay configuration.	
Configure Legacy DHCP Relay for ELS Switches	To configure Legacy DHCP Relay for Campus Switching ELS: <ol style="list-style-type: none"> 1. Select Legacy DHCP Relay (default). 2. Add DHCP servers for Legacy DHCP Relay: <ol style="list-style-type: none"> a. Click Add under DHCP Servers. The Add Server window opens. b. Enter the IP address of the DHCP server and then click OK. The DHCP server name appears in the list of DHCP servers. <p>TIP: You can add more than one DHCP server.</p>

Table 69: Device Profile Protocol Settings for Data Center Switching ELS (continued)

Task	Action
Configure Extended DHCP Relay for Campus Switching ELS	<p>To add Extended DHCP Relay to this Campus Switching ELS Device profile:</p> <ol style="list-style-type: none"> 1. Select Extended DHCP Relay. 2. Add one or more DHCP Servers Groups to the Device Common Settings profile: <ol style="list-style-type: none"> a. Click Add under Add DHCP Servers Group. The Add Server Group window opens. b. Provide a name for the server group. c. Optionally, make this an active server group by checking Active Group. d. Add servers to the group by clicking Add under DHCP Servers. The phrase <i>Click to enter value</i> appears in the IP Address column. e. Select <i>Click to enter value</i> and then enter an IP Address. f. Click OK. The server is added to the DHCP server group list. g. Add a relay interface group by clicking Add under Add Relay Interface Group. The Add DHCP Relay Interface window opens. h. Type a group name for the DHCP interface. i. Select a server group from the Server Group list. j. Click OK. The group is added to the Relay Interface Group list.

DNS Settings

Table 69: Device Profile Protocol Settings for Data Center Switching ELS (continued)

Task	Action
Add a domain name server	<p>To add a domain name server to the Campus Switching ELS Common Settings:</p> <ol style="list-style-type: none"> 1. Click Add under Domain Name Servers. The Add Server window opens. 2. Provide an IP address for the DNS server. 3. Click OK. The server is added to the Domain Name Servers list. <p>TIP: To edit a DNS server's settings, select it and then click Edit. To delete a DNS server, select it and then click Delete.</p>

Click either **Next** or **Review**, to see the Review page. For review directions, see [“Reviewing and Saving a Device Common Settings Configuration” on page 330](#).

Reviewing and Saving a Device Common Settings Configuration

From this page, you can save or make changes to Device Common Settings:

- To make changes to the settings, click the **Edit** associated with the configuration you want to change.

Alternatively, you can also click appropriate sections of the workflow at the top of the page that corresponds to the configuration you want to change.

When you have completed your modifications, click **Review** to return to this page.

- To save a new profile or to save modified settings to an existing profile, click **Finish**.

The Manage Device Common Settings page is displayed with the new or modified profile listed

What to Do Next

Once the Device Common Settings profile is created, you must assign the profile to the required device by using the Manage Device Profile page and then deploy the Device profile by using the **Deploy** mode. To assign a Device Common Settings profile to a device, see [“Assigning Device Common Settings to Devices” on page 331](#). For information about deploying your configurations, see [“Deploying Configuration to Devices” on page 1180](#).

NOTE: A device can have only one Device profile assigned to it. However, you can assign the same Device profile to multiple devices.

RELATED DOCUMENTATION

[Assigning Device Common Settings to Devices | 331](#)

[Deploying Configuration to Devices | 1180](#)

[Network Director Documentation home page](#)

Assigning Device Common Settings to Devices

IN THIS SECTION

- [Assigning Device Common Settings | 332](#)
- [Editing the Assignments of the Device Common Setting | 334](#)

Once a Device Common Settings profile is created or discovered (system-created profile), you must assign it to devices using the steps described in this topic. You can assign a Device profile to a either single device, a series of single devices, or a Custom Group of devices (see [“Creating Custom Device Groups” on page 276](#)).


NOTE: A device can have only one Device Common Settings profile assigned to it.

You must have one or more device profiles created or discovered before you can assign a device profile to a device. When you deploy an assigned device profile, the configuration is pushed onto the device.

This topic describes:

Assigning Device Common Settings

To assign device common settings to either a single device, a series of single devices, or members of a Custom Group:

1. Click  in the Network Director banner.
2. Select **Device Common Settings** from the Profile and Configuration Management menu in the Tasks pane.

The Manage Device Common Settings page is displayed. The page displays all the device profiles that you configured as well as the system-created profiles detected during device discovery.
3. Select an undeployed profile from the list of profiles and then click **Assign**.

The Assign Device Profile page for the selected device family appears with a wizard consisting of three parts, Device Selection, Profile Assignment, and Review. Device Selection is displayed.
4. Expand the Device Selection object tree and select one or more objects to receive the device profile. You must place a check next to a device to select it—simply highlighting the device does not select it.

NOTE: If Network Director fails to read the configuration of one or more devices after device discovery, those devices are not displayed in the Device Selection list. You will not be able to assign profiles to those devices. The Manage Jobs page in System mode displays details of the device discovery jobs. Use the information displayed on this page to take appropriate corrective steps to enable Network Director to reread the configuration of the failed device. For more information, see [“Discovering Devices in a Physical Network” on page 204](#).

5. Click either **Next** or click **Profile Assignment** from the wizard workflow.

The Profile Assignment page opens, displaying your selections, including their Device (name), Type, Assigned To, and Attributes. The Assigned To column now has the entry DEVICE and the Attributes column has the entry Undefined.
6. Click **Define** in the **Attributes** column in the Assignments table to configure the attributes.

The Configure Attributes window opens, listing all the Layer 3 interfaces available on the device.
 - a. Select the Layer 3 interfaces that are required for DHCP relay from the Available list and using the right arrow, move them to the Selected list. You can reorder the interfaces using the UP and DOWN arrows.

- b. Click **Save** to save the interface list and close the Configure Attributes window.
7. You can view the assignment details for the selected device and also remove any assignments:
 - To view the assignment details, select the device and click **View Assignments**.

The Profile Details page for selected device appears. Expand the **Device** name to view the details of the assignment. The assignment status displays the status whether the device is deployed or is pending device update, and so on.
 - To delete a device common setting assignment for a device, select the device from the Assignments table and click **Remove**.
8. Click **Next** or click **Review** from the wizard workflow to review the assignments. On the Review page, click **Edit** to edit the profile assignment.
9. Click **Finish** once you are done reviewing the profile assignment.

The Create Profile Assignments Job Details window appears with a status report for the profile assignment job—click **OK** to close this window. If you have assigned the profile to a large number of objects, the profile assignment job can take some time to complete. Instead of waiting for the Job Details dialog box to report job completion status, you can close it and check the details of the profile assignment job at a later time using the Manage Job task in System mode.

NOTE: If any assignment fails, the profile assignment job fails and none of the assignments are created. Check the details for the profile assignment job for information about why the assignment failed.

An assigned Device profile has the Assignment State *Pending Deployment* in the Manage Device Common Settings list. Deploy any device profile in this state by following the directions “[Deploying Configuration to Devices](#)” on page 1180.

To view the details of a profile, select the profile from the Manage Device Common Settings page and then click **Details**.

SEE ALSO

[Creating and Managing Device Common Settings | 291](#)

[Deploying Configuration to Devices | 1180](#)

[Network Director Documentation home page](#)

Editing the Assignments of the Device Common Setting

Use the Edit Assignments page to change device common setting assignments. To edit an existing assignment:

1. Select a profile from the **Manage Device Common Settings** page and click **Edit Assignment**.

The Edit Assignments page for the selected device appears.

2. Expand the **Devices** cabinet and make the desired change from the **Operation** column of the table.

3. Click **Define** from the **Attributes** column of the table to modify the attributes.

The Configure attributes page is displayed listing all the Layer 3 interfaces available on the device.

- Select the Layer 3 interfaces that are required for DHCP relay from the Available box and using the right arrow, move them to the Selected box.

You can rearrange the order of the interfaces using the Up and Down arrows.

- Click **Save** after you are done with selecting the interfaces.

4. Click **Apply** once you are done with the changes.

The Manage Device Common Settings page is displayed.

RELATED DOCUMENTATION

[Creating and Managing Device Common Settings | 291](#)

[Creating Custom Device Groups | 276](#)

[Network Director Documentation home page](#)

Configuring Authentication, Authorization, and Access for Your Network

IN THIS CHAPTER

- Understanding Central Network Access Using RADIUS and TACACS+ | 335
- Creating and Managing RADIUS Profiles | 339
- Creating and Managing LDAP Profiles | 345
- Understanding Access Profiles | 351
- Creating and Managing Access Profiles | 352
- Understanding Authentication Profiles | 381
- Creating and Managing Authentication Profiles | 383
- Understanding Wireless Authorization Profiles | 395
- Creating and Managing Wireless Authorization Profiles | 395
- Assigning Wireless Authorization Profiles to Controllers | 404

Understanding Central Network Access Using RADIUS and TACACS+

IN THIS SECTION

- Why Do I Want Remote Authentication ? | 336
- Where Is RADIUS Installed on the Network? | 337
- How Is TACACS+ Installed on the Network? | 337
- A Comparison of RADIUS and TACACS+ | 338

Remote Access Dial In User Service (RADIUS) and Terminal Access Controller Access-Control System Plus (TACACS+) are two common security protocols used to provide centralized access into networks. RADIUS was designed to authenticate and log remote network users, while TACACS+ is most commonly used for administrator access to network devices like routers and switches. Both protocols provide centralized Authentication, Authorization, and Accounting (AAA) management for computers that connect and use a network service.

- *Authentication* - Who is allowed to gain access to the network? Traditionally authorized users provide a username and password to verify their identity for both RADIUS and TACACS+.
- *Authorization* - What services can a user access once they are authenticated? It is unlikely that you want your finance people to have access to the developer database. Visitors may have access only to the Internet, while only IT staff can access the entire passwords database.
- *Accounting* - What services did each user access and for how long? Accounting records record the user's identification, network address, point of attachment and a unique session identifier—these statistics are tracked and added to the user's record. This is useful when time on the system is billed to individuals or departments.

Why Do I Want Remote Authentication ?

Remote authentication enables you to keep your username and passwords in one place, on a central server. The advantage to using RADIUS or TACACS+ on this central server is that you don't configure changes on each separate network device when a user is added or deleted, or when a user changes a password. You only make one change to the configuration on the server and then devices continue to access the server for authentication. Although authentication is the most well known function of RADIUS and TACACS+, there are two additional functions provided, authorization and accounting.

NOTE: Instead of using a flat database on the RADIUS server, you can refer to external sources such as SQL, Kerberos, LDAP, or Active Directory servers to verify user credentials.

Why Not Just Rely on Firewalls and Filters for Access Control?

Routers and firewalls usually control access to services using filters based on source and/or destination IP addresses and ports. This means that restrictions are applied to devices and not to individual clients. For example if I enable traffic from 10.1.0.255 to access a particular web server, then anyone who is sitting at the machine with the address of 10.1.0.255 automatically has access to this server. Using RADIUS or TACACS+, that same person sitting at the machine with the address of 10.1.0.255 also has to provide a username and password to access a service.

What About Using LDAP For Authentication?

Lightweight Directory Access Protocol (LDAP) is a client/server protocol used to access and manage directory information. It reads and edits directories over IP networks and runs directly over TCP/IP using

simple string formats for data transfer. Directory servers include information about various entities on your network, such as user names, passwords, rights associated with user names, metadata associated with user names, devices connected to the network, and device configuration.

Use LDAP to obtain directory information, such as email addresses and public keys. If you want to make directory information available over the Internet, this is the way to do it. LDAP works well for captive portal authentication. However, LDAP does not implement 802.1X security easily. 802.1X was essentially designed with RADIUS in mind, so 802.1X challenge/response protocols like MSCHAPv2 work well with RADIUS.

Where Is RADIUS Installed on the Network?

RADIUS includes three components: an authentication server, client protocols, and an accounting server. The RADIUS server portion of the protocol is usually a background process running on a UNIX or Microsoft Windows server.

With RADIUS, the term client refers to a network access device (NAD) that provides the client part of the RADIUS service—wireless access points, a modem pool, a switch, a network firewall, or any other device that needs to authenticate users can be configured as a NAD to recognize and process connection requests from outside the network edge. When a NAD receives a user's connection request, it may perform an initial access negotiation with the user to obtain identity/password information. Then the NAD passes this information to the RADIUS server as part of an authentication/authorization request.

NOTE: RADIUS requires that each network client device be configured.

How Is TACACS+ Installed on the Network?

TACACS+ logon authentication protocol uses software running on a central server to control access by TACACS-aware devices on the network. The server communicates with switches or other TACACS-aware devices automatically—these devices do not require further configuration if they are TACACS-aware. The TACACS+ protocol is supported by most enterprise and carrier-grade devices.

Install the TACACS+ Service as close as possible to the user database, preferably on the same server. TACACS+ needs to be closely synchronized with your Domain, and any network connection issues, DNS problems, or even time discrepancies can cause a critical service failure. Installing TACACS+ on the same server as the user database can also improve performance.

TACACS+ servers should be deployed in a fully trusted internal network. If you keep your TACACS+ service within your trusted network, you need to open only one port, TCP 49. There should not be any direct access from untrusted or semi-trusted networks.

NOTE: RADIUS is typically deployed in a semi-trusted network, and TACACS+ uses internal administrative logins, so combining these services on the same server could potentially compromise your network security.

A Comparison of RADIUS and TACACS+

Table 70: RADIUS and TACACS+

	RADIUS	TACACS+
Primary Use	Authenticate and log remote network users	Provide administrator access to network devices like routers and switches
Authentication and Authorization	Authentication and Authorization checking are bundled together. When the client device requests authentication from the server, the server replies with both authentication attributes and authorization attributes. These functions can not be performed separately.	All three AAA functions (authentication, authorization, and accounting) can be used independently. Therefore, one method such as kerberos can be used for authentication, and a separate method such as TACACS+ can be used for authorization.
Accounting	The accounting features of the RADIUS protocol can be used independently of RADIUS authentication or authorization.	
Protocol	User Datagram Protocol (UDP)/IP with best-effort is used for delivery on ports 1645/1646, 1812/1813	TCP used for delivery on port 49. Also has multiprotocol support for AppleTalk Remote Access (ARA) protocol, NetBIOS Frame Protocol Control protocol, Novell Asynchronous Services Interface (NASI), and X.25 PAD connection.
Encryption applied to	Password	Username and password
802.1X Security	If you want to use 802.1x port-based network access control, you have to use the RADIUS client because the TACACS+ client does not support that feature.	
Model	client/server	
Recommended Environment	semi-trusted	trusted

RELATED DOCUMENTATION

[Creating and Managing RADIUS Profiles | 339](#)

[Network Director Documentation home page](#)

Creating and Managing RADIUS Profiles

IN THIS SECTION

- [Managing RADIUS Profiles | 339](#)
- [Creating RADIUS Profiles | 340](#)
- [Specifying Settings for a RADIUS Profile | 341](#)
- [What to Do Next | 344](#)

Remote Authentication Dial In User Service (RADIUS) is a networking protocol that provides centralized authentication, authorization, and accounting (AAA) management for computers to connect and use a network service. By default, RADIUS servers are used for both accounting and authentication. From Network Director, you can create and manage RADIUS profiles that configure RADIUS server settings.

TIP: In addition to your RADIUS server, you can configure an LDAP server for either wireless and EX Series ELS switch authentication—for directions, see [“Creating and Managing LDAP Profiles” on page 345](#).

This topic describes:

Managing RADIUS Profiles

From the Manage RADIUS Profiles page, you can:

- Create a new profile by clicking **Add**. For directions, see [“Creating RADIUS Profiles” on page 340](#).
- Modify an existing profile by selecting it and clicking **Edit**.
- View information about a profile by selecting the group and clicking **Details** or by clicking the profile name.
- Delete profiles by selecting the profile and clicking **Delete**.

TIP: You cannot delete profiles that are in use—that is, assigned to objects or used by other profiles. To see the current assignments for a profile, select the profile and click **Details**.

- Clone a profile by selecting the profile and clicking **Clone**.

Table 71 describes the information provided about RADIUS profiles on the Manage RADIUS Profiles page. This page lists all RADIUS profiles defined for your network, regardless of your current selected scope in the network view.


Table 71: RADIUS Profile Information

Field	Description
RADIUS Profile Name	Name given to the RADIUS profile when it was created.
Server Address	IP address of the RADIUS server.
Server Port	UDP port being used by the RADIUS server.
Creation Time	Date and time when this profile was created.
Update Time	Date and time when this profile was last modified.
User Name	The username of the user who created or modified the profile.

TIP: All columns may not be currently displayed. To show or hide fields in the table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

Creating RADIUS Profiles

To create a RADIUS profile:

1. Click  in the Network Director banner.
2. Under Select View, select either **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View** or **Topology View**.

3. From the Tasks pane, select the type of network (Wired or Wireless), the appropriate functional area (System, AAA, or Wireless), and select the name of the profile that you want to create. For example, to create a radius profile for a wireless device, click **Wireless > AAA > Radius**. The Manage Profile page opens.
4. Click **Add** on the Manage RADIUS Profiles page.
The Create RADIUS Profile page appears.
5. Enter settings for the RADIUS profile on the Create RADIUS Profile page as described in [“Specifying Settings for a RADIUS Profile” on page 341](#).
6. Click **Done**.

Specifying Settings for a RADIUS Profile

Use the Create RADIUS Profile page to define authentication, authorization, and accounting settings for a RADIUS server.

[Table 72](#) describes the RADIUS profile settings.

Table 72: RADIUS Profile Settings

Field	Action
Server Name	Type a name for the server, using up to 64 alphanumeric characters and no special characters other than the underscore. The name must be unique among servers.
Server Address	Type the IP address of the RADIUS server.
Authentication Port (default is 1812)	Using the arrows, adjust the number of the UDP port to use for RADIUS authentication messages. The default UDP port is 1812, and the range is from 0 to 65535.
Secret	Provide a password for the RADIUS server.

Advanced Settings

You can change the advanced settings for a RADIUS server, or you can use the default settings.

Table 72: RADIUS Profile Settings (continued)

Field	Action
Accounting Port (default is 1813)	Using the arrows, adjust the number of the UDP port to use for RADIUS accounting messages. The default UDP port is 1813, and the range is from 0 to 65535.
Retry Count (default is 3)	Using the arrows, adjust the retry count until it reflects the number of times Network Director retries connecting to the RADIUS server when the RADIUS server is unavailable.
Timeout (default is 5 seconds)	Using the arrows, adjust the timeout value. Timeout indicates how many seconds Network Director allows for RADIUS server connection before giving an unreachable error.
Dead Time (default is 5 seconds)	Using the arrows, adjust the number of seconds before Network Director checks a RADIUS server that was previously unresponsive. The default value is 5 seconds.
Use MAC as Password	Enable this option if you want each client device to use its MAC address as its password for the RADIUS server. If you enable Use MAC As Password, then the Authorization Password field becomes unavailable.
Authorization Password	If you are not using MAC addresses as passwords for the RADIUS server, provide a common password here.

Table 72: RADIUS Profile Settings (continued)

Field	Action
MAC Address Format	<p>Select None, Hyphens, Colons, One-Hyphen, or Raw to determine the MAC address format used with the RADIUS server. For example:</p> <ul style="list-style-type: none"> ● None—For unicast IPv4, an example MAC address is 0123456789ab. For unicast IPv6, an example MAC address is 20010db8000000000000ff0000428329. ● Hyphens—For unicast IPv4, an example MAC address using hyphens is 01-23-45-67-89-ab. For unicast IPv6, an example MAC address using hyphens is 2001-0db8-0000-0000-0000-ff00-0042-8329. ● Colons—For unicast IPv4, an example MAC address using colons is 01:23:45:67:89:ab. For unicast IPv6, an example MAC address using colons is 2001:0db8:0000:0000:0000:ff00:0042:8329. ● One-Hyphen: IPv6 unicast addresses other than those that start with binary 000 are logically divided into two parts: a 64-bit (sub-)network prefix, and a 64-bit interface identifier used to identify a host's network interface. The hyphen is placed between the two parts. ● Raw: The IPv6 address is represented by all numbers—no sections containing all zeros are skipped and then represented by a double colon. For example, this is a raw IPv6 address: 2001:0000:0234:C1AB:0000:00A0:AABC:003F.

Table 72: RADIUS Profile Settings (continued)

Field	Action
Authentication Protocol (Default is PAP)	<p>Select PAP, CHAP, MSCHAP-V2, or None to determine an authentication protocol for the RADIUS server. These authentication protocols work as follows:</p> <ul style="list-style-type: none"> • PAP: stands for Password Authentication Protocol and is used by Point to Point Protocols to validate users before allowing them access to server resources. Almost all network operating system remote servers support PAP. However, PAP transmits unencrypted ASCII passwords over the network and is therefore not secure. Use it as a last resort when the remote server does not support the stronger authentication. • CHAP: stands for Challenge Handshake Authentication Protocol and authenticates a user or network host to an authenticating entity. CHAP provides protection against replay attacks by the peer through the use of an incrementally changing identifier and of a variable challenge-value. CHAP requires that both the client and server know the plaintext of the secret password—it is never sent over the network. CHAP provides better security than PAP does. • MSCHAP-V2: stands for Microsoft's implementation of the Challenge Handshake Authentication Protocol version 2 on the router for password-change support. This feature provides users accessing a router the option of changing the password when the password expires, is reset, or is configured to be changed at the next login. The MS-CHAP variant does not require either peer to know the plaintext of the secret password. MSCHAP-V2 is used as an authentication option with RADIUS servers used for Wi-Fi security using the WPA-Enterprise protocol.
Server Priority (default is 1)	<p>Enter a server priority to indicate the order in which RADIUS servers are accessed. Entering a one means that this server is checked first.</p>

Click **OK** to add the RADIUS server to the EX Switching Access profile. You can add more RADIUS servers if needed.

If you have multiple RADIUS servers, you can prioritize them in the Authentication Server Order section, using the arrows.

Click **Done** to create the RADIUS server profile.

The RADIUS server name appears in the list of RADIUS servers on the Manage RADIUS Profiles page.

What to Do Next

Link the RADIUS server to an Access profile. For directions, see [“Creating and Managing Access Profiles” on page 352](#).

NOTE: Assigned settings from any profile, including this one, have lower priority than settings made directly to a controller or an access point. For more information, see [“Adding and Managing an Individual Access Point”](#) on page 1156 and [“Configuring a Controller”](#) on page 1037.

RELATED DOCUMENTATION

[Creating and Managing Access Profiles | 352](#)

[Creating and Managing LDAP Profiles | 345](#)

[Understanding Central Network Access Using RADIUS and TACACS+ | 335](#)

[Network Director Documentation home page](#)

Creating and Managing LDAP Profiles

IN THIS SECTION

- [Managing LDAP Profiles | 346](#)
- [Creating LDAP Profiles | 347](#)
- [Specifying Settings for an LDAP Profile | 348](#)
- [What to Do Next | 350](#)

Lightweight Directory Access Protocol (LDAP) is an Internet protocol that email and other programs use to look up information from a server. Use LDAP to look up encryption certificates, pointers to printers and other services on a network, in addition to providing a single logon where one user password is used for different services. LDAP authentication is appropriate for any kind of directory-like information where fast lookups and infrequent updates are used. From Network Director, you can create and manage LDAP profiles for both wireless and EX Switching ELS.

TIP: In addition to an LDAP server, you can configure a RADIUS server for both authentication and accounting purposes—for directions, see [“Creating and Managing RADIUS Profiles” on page 339](#).

This topic describes:

Managing LDAP Profiles

From the Manage LDAP Profiles page, you can:

- Create a new LDAP profile by clicking **Add**. For directions to add an LDAP profile, see [“Creating LDAP Profiles” on page 347](#).
- Modify an existing LDAP profile by selecting it and clicking **Edit**.
- View information about a profile by selecting the group and clicking **Details** or by clicking the profile name.
- Delete LDAP profiles by selecting the profile and clicking **Delete**.

TIP: You cannot delete profiles that are in use—that is, assigned to objects or used by other profiles. To see the current assignments for a profile, select the profile and click **Details**.

- Clone an LDAP profile by selecting a profile and clicking **Clone**.

[Table 73](#) describes the information provided about LDAP profiles on the Manage LDAP Profiles page. This page lists all LDAP profiles defined for your network, regardless of your current selected scope in the network view.

Table 73: LDAP Profile Information

Field	Description
LDAP Name	Name given to the LDAP server profile when it was created.
Server Address	IP address of the LDAP server.
Server Port	UDP port being used by the LDAP server.
Domain Name	Domain using the LDAP server.
Creation Time	Date and time when this profile was created.


Table 73: LDAP Profile Information (continued)

Field	Description
Update Time	Date and time when this profile was last modified.
User Name	The username of the user who created or modified the profile.

TIP: All columns of information may not be displayed. To show or hide fields in the table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

Creating LDAP Profiles

To create an LDAP profile:

1. Click  **Build** in the Network Director banner.
2. Under Select View, select either **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View** or **Topology View**.

3. From the Tasks pane, select the type of network (Wired or Wireless), the appropriate functional area (System, AAA, or Wireless), and select the name of the profile that you want to create. For example, to create a radius profile for a wireless device, click **Wireless > AAA > Radius**. The Manage Profile page opens.
4. Click **Add** to add a new profile.
The Create LDAP Profile page for the selected device family is displayed.
5. Enter settings for the LDAP profile as described in [“Specifying Settings for an LDAP Profile” on page 348](#).
6. Click **Done**.

Specifying Settings for an LDAP Profile

Use the Create LDAP Profile page to define LDAP directory information services over an IP network.

Table 74 describes the LDAP settings.

Table 74: LDAP Profile Settings

Field	Action
Server Name	Type a name for the server, using up to 64 alphanumeric characters and no special characters other than the underscore. The name must be unique among servers.
Server Address	Type the IP address of the LDAP server.
Server Port (default is 389)	Using the arrows, adjust the number of the UDP port to use for LDAP authentication messages. The default port is 389 for unencrypted LDAP servers and 636 for unencrypted LDAP servers.
Advanced LDAP Settings	
Fully Qualified Domain Name	Type a fully qualified domain name (FQDN)—this is the complete domain name for a specific computer, or host, on the Internet. The FQDN consists of two parts: the host name and the domain name. For example, an FQDN for a server might be ldap12.example.com. The host name is ldap12, and the host is located within the domain example.com. This domain name specifies all domain levels, including the top-level domain and the root zone. A fully qualified domain name is specified with a trailing dot, for example: ldap12.example.com.
Dead Time (default is 5 seconds)	Using the arrows, adjust the number of seconds before Network Director checks an LDAP server that was previously unresponsive. The default value is 5 seconds.
Timeout (default is 5 seconds)	Using the arrows, adjust the number of seconds Network Director tries to establish connection with RADIUS server before giving an unreachable error.

Table 74: LDAP Profile Settings (continued)

Field	Action
Bind Mode (default is SASL-MD5)	<p>Select either SASL-MD5 or SIMPLE-AUTH to establish authentication for an LDAP session.</p> <p>Simple Authentication and Security Layer (SASL) is a framework for authentication and data security in Internet protocols. It decouples authentication mechanisms from application protocols, in theory enabling any authentication mechanism supported by SASL to be used in any application protocol that uses SASL.</p> <p>SIMPLE-AUTH sends the user's domain name and password in plain text. The server then checks the password against the password attribute in the named entry.</p> <p>TIP: We recommend that connections using SIMPLE-AUTH be encrypted using Transport Layer Security (TLS).</p>
MAC Address Format (default is Hyphens)	<p>Select None, Hyphens, Colons, One-Hyphen, or Raw to determine the MAC address format used with the LDAP server. For example:</p> <ul style="list-style-type: none"> • None: For unicast IPv4, an example MAC address is 0123456789ab. For unicast IPv6, an example MAC address is 20010db8000000000000ff0000428329. • Hyphens: For unicast IPv4, an example MAC address with hyphens is 01-23-45-67-89-ab. For unicast IPv6, an example MAC address with hyphens is 2001-0db8-0000-0000-0000-ff00-0042-8329. • Colons: For unicast IPv4, an example MAC address with colons is 01:23:45:67:89:ab. For unicast IPv6, an example MAC address with colons is 2001:0db8:0000:0000:0000:ff00:0042:8329. • One-Hyphen: IPv6 unicast addresses other than those that start with binary 000 are logically divided into two parts: a 64-bit (sub-)network prefix, and a 64-bit interface identifier used to identify a host's network interface. The hyphen is placed between the two parts. • Raw: The IPv6 address is represented by all numbers—no sections containing all zeros are skipped and then represented by a double colon. For example, this is a raw IPv6 address: 2001:0000:0234:C1AB:0000:00A0:AABC:003F. <p>TIP: A media access control address (MAC address) is a unique identifier assigned to network interfaces for communications on the physical network segment. MAC addresses are used as a network address for most IEEE 802 network technologies, including Ethernet.</p>
Base Domain	<p>Base domains contain no extra dots. For example, example.com is a base domain, but www.example.com is not because it contains an extra dot.</p>

Table 74: LDAP Profile Settings (continued)

Field	Action
Domain Prefix (default is cn)	Enter a domain prefix to identify a subdomain. The subdomain name can be used to identify services, devices, or regions.
Use MAC as Password (default is unchecked)	Check this option if you want each client device to use its MAC address as its password for the LDAP server.
Authorization Password	If you are not using individual MAC addresses as passwords for the LDAP server, provide a common password here.

Click **Done** to create the LDAP Server profile. The profile appears on the list on the Manage LDAP Profiles page.

What to Do Next

Link the LDAP server to an Access profile for either wireless or for Campus Switching with ELS. For directions, see [“Creating and Managing Access Profiles” on page 352](#).

NOTE: Assigned settings from any profile, including this one, have lower priority than settings made directly to a controller or an access point. For more information, see [“Adding and Managing an Individual Access Point” on page 1156](#) and [“Configuring a Controller” on page 1037](#).

RELATED DOCUMENTATION

[Creating and Managing Access Profiles | 352](#)

[Creating and Managing RADIUS Profiles | 339](#)

[Network Director Documentation home page](#)

Understanding Access Profiles

Access profiles enable access configuration on the network—this consists of authentication configuration and accounting configuration. Network Director supports RADIUS, Lightweight Directory Access Protocol (LDAP), and local authentication as authentication methods, and RADIUS for accounting.

Authentication prevents unauthorized devices and users from gaining access to your network. Authentication controls access to your network using authentication methods such as 802.1X, MAC RADIUS, or captive portal. For 802.1X and MAC RADIUS authentication, end devices or users must be authenticated before they receive an IP address from a DHCP server. For captive portal authentication, the switch or the controller enables the end devices to obtain an IP address, after which these devices can forward packets such as DHCP, DNS, and ARP.

Accounting servers collect and send information used for billing, auditing, and reporting, such as:

- User identity
- Connection start and stop times
- Number of packets received and sent
- Number of transferred bytes

The accounting information is stored locally or on a remote RADIUS server. You can track sessions by using this information. As network users roam through a Network or Mobility Domain, accounting records can be used to track their network usage.

RADIUS is an authentication and accounting server used for validating users who attempt to access the wireless controller or switch. RADIUS is a distributed client-server system—the RADIUS client runs on the controller or the switch, and the server runs on a remote network system.

LDAP is an Internet protocol for accessing and updating information in an X.500-compliant directory. Network administrators for LDAP clients can connect to X.500 directory service and add, delete, modify, or search for information if they have the required access rights to the directory. LDAP is designed to run over TCP/IP and can access information in both X.500 directories and many non-X.500 directories.

NOTE: LDAP is supported as an authentication and accounting method for Campus Switching ELS and wireless devices.

With local authentication, you configure a password for each user allowed to log in to the controller or switch.

You can define one or more Access profiles. Each Access profile is specific to a device family. Use the Manage Access Profiles page to create, modify, view, and delete existing Access profiles.

RELATED DOCUMENTATION

[Creating and Managing Access Profiles | 352](#)

[Creating and Managing RADIUS Profiles | 339](#)

[Creating and Managing LDAP Profiles | 345](#)

[Network Director Documentation home page](#)

Creating and Managing Access Profiles

IN THIS SECTION

- [Managing Access Profiles | 352](#)
- [Creating an Access Profile | 354](#)
- [Specifying Basic Settings for an EX Series Switching or Data Center Switching Access Profile | 356](#)
- [Specifying RADIUS Accounting Settings for an EX Switching or Data Center Switching Access Profile | 358](#)
- [Specifying Basic Settings for a Wireless Access Profile | 361](#)
- [Specifying Server Group Settings for a Wireless Access Profile | 361](#)
- [Specifying Basic Settings for a Campus Switching ELS Access Profile | 372](#)
- [Specifying RADIUS and LDAP Settings for Campus Switching ELS | 372](#)
- [Reviewing and Modifying the Access Profile Settings | 380](#)
- [What To Do Next | 380](#)

Access profiles enable authentication configuration for both methods and servers. Network Director supports the configuration of RADIUS, Lightweight Directory Access Protocol (LDAP), and local authentication as authentication methods, and RADIUS as an accounting method.

Use the Manage Access Profiles page to create new Access profiles and manage existing Access profiles.

This topic describes:

Managing Access Profiles

From the Manage Access Profiles page, you can:

- Create a new Access profile by clicking **Add**. For directions, see [“Creating an Access Profile” on page 354](#).
- Modify an existing profile by selecting it and clicking **Edit**.

- View information about an Access profile, including the interfaces it is associated with, by either clicking the profile name or by selecting the profile and clicking **Details**.
- Delete an Access profile by selecting the Access profile and clicking **Delete**.

TIP: You cannot delete profiles that are in use—that is, assigned to objects or used by other profiles. To see the current assignments for an Access profile, select the Access profile and click **Details**.

- Clone a profile by selecting a profile and clicking **Clone**.

TIP: The default Access profile named *Juniper Networks-access-profile* is always available.

Table 75 describes the information provided about Access profiles on the Manage Access Profiles page. This page lists all Access profiles defined for your network, regardless of the scope you selected in the network view.

Table 75: Manage Access Profile Fields

Field	Description
Profile Name	Name given to the profile when the profile was created.
Description	Description of the profile that was entered when the profile was created. TIP: To display the entire description, you might need to resize the Description column by clicking the column border in the heading and dragging it.
Family Type	The device family on which the profile was created: EX Switching, Wireless, Campus Switching ELS, or Data Center Switching Non-ELS.
Creation Time	Date and time when the profile was created.
Last Updated Time	Date and time when the profile was last modified.
User Name	The username of the person who created or modified the profile.

TIP: All columns might not be displayed. To show or hide fields listed in the table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

Creating an Access Profile

In Network Director, you create an Access profile that is then used to authenticate network users. You can also specify servers to be used for user accounting purposes. You can create Access profiles for these kinds of hardware devices:

- EX Series Switches—configure Basic Settings and optional Accounting Settings.
- Wireless (WLC)—configure Basic Settings and Server Group Settings.
- EX Series switches with ELS—configure Basic Settings and Server Settings.
- Data Center Switching Non-ELS—configure Basic Settings and optional Accounting Settings.

To create an Access profile, follow these steps:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.

3. Under Select View, select either **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View** or **Topology View**.

4. From the Tasks pane, select the type of network (Wired or Wireless), the appropriate functional area (System, AAA, or Wireless), and select the name of the profile that you want to create. For example, to create a radius profile for a wireless device, click **Wireless > AAA > Radius**. The Manage Profile page opens.

5. Click **Add** to add a new profile.

If you chose to create a profile for the wired network, Network Director opens the Device Family Chooser window.

- a. From the Device Family Chooser, select the device family for which you want to create a profile. The available device families are **Switching (EX)**, **Campus Switching ELS** (Enhanced Layer 2 Software), **Data Center Switching Non-ELS** and **Data Center Switching ELS**.

- b. Click **OK**.

The Create Access Profile page for the selected device family is displayed.

If you chose to create a profile for the wireless network, Network Director opens the Create Access Profile for Wireless page.

6. Click **Add**.

The Device Family Chooser window opens.

7. From the Device Family Chooser, select the device family for which you want to create a profile. The available device families are **Switching EX**, **Wireless (WLC)**, **Campus Switching ELS**, and **Data Center Switching Non-ELS**.

8. Click **OK**.

The Create Access Profile wizard for the selected device family opens—it consists of two sections: Basic Settings and RADIUS and LDAP configuration.

9. Specify the access settings for the Access profile by doing one of the following:

- For either EX Series switches or Data Center Switching, specify the access settings described in online help, or in [“Specifying Basic Settings for an EX Series Switching or Data Center Switching Access Profile” on page 356](#) and [“Specifying RADIUS Accounting Settings for an EX Switching or Data Center Switching Access Profile” on page 358](#).
- For controllers, specify the access settings as described in online help and [“Specifying Basic Settings for a Wireless Access Profile” on page 361](#) and [“Specifying Server Group Settings for a Wireless Access Profile” on page 361](#).
- For Campus Switching ELS, specify the access settings as described in [“Specifying Basic Settings for a Campus Switching ELS Access Profile” on page 372](#) and [“Specifying RADIUS and LDAP Settings for Campus Switching ELS” on page 372](#).

10. Click either **Next** or **Review**. The Review page appears.

You can either save your profile or make changes to your profile from the Review page. For directions, see [“Reviewing and Modifying the Access Profile Settings” on page 380](#).

11. Click **Done** to save the Access profile.

The system saves the Access profile and then displays the Manage Access Profiles page. Your new or modified Access profile is listed in the table of Access profiles.

Specifying Basic Settings for an EX Series Switching or Data Center Switching Access Profile

Basic settings for EX Series switching or data center switching Access profile include the profile name, authentication server order, and the RADIUS authentication details.

To configure the basic settings for an EX Series switch or data center switching Access profile, enter the settings described in [Table 76](#). Required settings are indicated in the user interface by a red asterisk (*) that appears next to the field label.

Table 76: Access Profile Basic Settings for EX Series Switches and Data Center Switching

Field	Action
Access Profile Details	
Profile Name	Type a unique name that identifies the profile. You can use up to 64 characters for profiles created for wired devices. Profile name must not contain special characters or spaces. Note that profiles that are automatically created by Network Director as part of device discovery or out-of-band changes may contain the underscore (_) character.
Description	Type the description of the profile.
Revert Interval	Specify the number of seconds the switch waits after an authentication server becomes unreachable. The switch rechecks the connection to the server when the specified interval expires. Default is 3 seconds.
RADIUS Servers: Authentication	
View	Select a server entry from the list and then click View to see the details of that entry.

Table 76: Access Profile Basic Settings for EX Series Switches and Data Center Switching (continued)

Field	Action
Task: Create and add a new RADIUS server configuration	<p>To both create and add a RADIUS server configuration to this Access profile for authentication:</p> <ol style="list-style-type: none"> 1. Click Add > Create RADIUS. The Create RADIUS Server window opens. 2. Complete these fields: <ul style="list-style-type: none"> • Server Name—Type the name of the RADIUS server that you want to create. • Server Address—Type the IP address of the RADIUS server. • Authentication Port—The default RADIUS authentication port is 1812. You can change the port number by using the up and down arrows. • Secret—Provide a password. If the password contains spaces, enclose it in quotation marks. The secret password used by the switch must match the one used by the server. 3. Expand the RADIUS server and change any of these configurations: <ul style="list-style-type: none"> • Accounting Port—You can change the default port number (1813) by using the up and down arrows. • Retry Count—Specify the number of times that a device attempts to contact the LDAP authentication server. The default retry count is 3. You can change this value by using the up and down arrows to 1 through 10 times. • Timeout (seconds)—Specify the number of seconds the switch waits to receive a response from a RADIUS server. The default timeout is 5 seconds. You can change this value, using the up and down arrows, to 1 through 90 seconds. 4. Click OK. The RADIUS server is automatically added to the list of authentication servers assigned to this Access profile. 5. If you have more than one RADIUS server listed, you can use the arrows to reorder the list priority so that the most preferred RADIUS server is listed first.

Table 76: Access Profile Basic Settings for EX Series Switches and Data Center Switching (continued)

Field	Action
Task: Add a previously configured RADIUS server for authentication	<p>The RADIUS tab is selected by default for server configuration and configured RADIUS servers are listed on this Server Settings page. To add a previously configured RADIUS server to this Access profile for authentication:</p> <ol style="list-style-type: none"> 1. Click Add > Select RADIUS. A list of available configured RADIUS servers is displayed. Servers in this list were either automatically discovered or created by using the directions in “Creating and Managing RADIUS Profiles” on page 339. 2. Select one or more RADIUS servers from the list of Available servers and use the arrows to move the server to the Selected list. 3. Click OK. The RADIUS server is added to the list of authentication servers to be used with this Access profile. 4. If you have more than one RADIUS server listed, you can use the arrows to reorder the login priority so that the most preferred RADIUS server is listed first.
Task: Delete a server	<p>To delete a RADIUS server from this Access profile:</p> <ol style="list-style-type: none"> 1. Select a RADIUS server from the list. 2. Click Delete. The RADIUS server is removed from the list of authentication servers to be used with this Access profile.

Proceed to the RADIUS Accounting settings for EX Switching Access profiles by clicking either **Accounting Settings** or **Next**. These settings are described in [“Specifying RADIUS Accounting Settings for an EX Switching or Data Center Switching Access Profile” on page 358](#).

Specifying RADIUS Accounting Settings for an EX Switching or Data Center Switching Access Profile

Configure the settings listed in [Table 77](#) for the Access profile Accounting Settings page. Accounting settings are optional in an Access profile. You can also specify accounting settings later by modifying an existing Access profile.

Table 77: Accounting Settings for an EX Switching and Data Center Switching Access Profile

Task	Description
View	Select a RADIUS server entry from the list and then click View to see the details of that entry.
Create a new RADIUS server for both authentication and accounting	<p>To both create and add a RADIUS server configuration to this Access profile for both authentication and accounting:</p> <p>NOTE: A RADIUS profile must be configured for authentication in addition to accounting.</p> <ol style="list-style-type: none"> 1. Click Add > Create RADIUS. The Add RADIUS Server window opens. 2. Complete these settings: <ul style="list-style-type: none"> • Server Name—Type the name of the RADIUS server that you want to create. • Server Address—Type the IP address of the RADIUS server. • Authentication Port—The default RADIUS authentication port is 1812. You can change the port number by using the up and down arrows. • Secret—Provide a password. If the password contains spaces, enclose it in quotation marks. The secret password used by the switch must match that used by the server. 3. Expand the Advanced Settings section and change any default settings, including the accounting port: <p>NOTE: If you do not change the accounting configuration, default values are used.</p> <ul style="list-style-type: none"> • Accounting Port—The default RADIUS accounting port is 1813. You can change the port number by using the up and down arrows. • Retry Count—Specify the number of times that a device attempts to contact the RADIUS server. The default retry count is 3. You can change this value by using the up and down arrows to 1 through 10 times. • Timeout (seconds)—Specify the number of seconds the switch waits to receive a response from the RADIUS server. The default timeout is 5 seconds. You can change this value, using the up and down arrows, to 1 through 90 seconds. 4. Click OK. The RADIUS server is automatically added to the list of RADIUS accounting servers assigned to this Access profile. 5. If you have more than one RADIUS server listed, you can use the arrows to reorder the list priority so that the most preferred RADIUS server is listed first.

Table 77: Accounting Settings for an EX Switching and Data Center Switching Access Profile (continued)

Task	Description
Add a previously configured RADIUS server for accounting	<p>A RADIUS server must already be configured before you can add that server for accounting. If the server was previously configured only for authentication, default accounting settings are applied. To add a RADIUS server for accounting:</p> <ol style="list-style-type: none"> Expand the Accounting Settings section of the Server Settings page. This is where RADIUS accounting is configured. A list of configured RADIUS servers is displayed. Click Add > Select RADIUS. A list of eligible RADIUS servers is displayed. Servers on this list were either automatically discovered, created following the directions “Creating and Managing RADIUS Profiles” on page 339, or created on this page following the directions <i>Create and add a new RADIUS server configuration</i>. If the server was configured only for authentication, default accounting settings were applied—you can use those default settings. Select a RADIUS server from the list of Available servers and then use the arrows to move it to the list of Selected servers. Click OK. The RADIUS server is added to the list of accounting servers to be used with this Access profile. If the RADIUS server was previously configured only for authentication, default accounting settings are applied. If you have more than one RADIUS server listed, you can use the arrows to reorder the login priority so that the most preferred RADIUS server is listed first for accounting.
Delete a server	<p>To delete a server from this Access profile:</p> <ol style="list-style-type: none"> Select a server from the list. Click Delete. The server is removed from the list of servers to be used with this Access profile.

Proceed to the Access profile review by clicking either **Review** or **Next**.

Specifying Basic Settings for a Wireless Access Profile

To configure the basic settings for a wireless Access profile, enter the settings described in [Table 78](#). Required settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 78: Access Profile Basic Settings for Wireless

Field	Action
Profile Name	Type a unique name that identifies the profile. Use up to 32 characters for wireless profile names. Profile names must not contain special characters or spaces. Note that profiles automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character.
Description	Type the description of the profile.

Proceed to the server group settings for wireless Access profiles by either clicking **Server Group Settings** from the wizard or by clicking **Next**. These settings are described in [“Specifying Server Group Settings for a Wireless Access Profile” on page 361](#).

Specifying Server Group Settings for a Wireless Access Profile

You can add RADIUS or LDAP servers to an Access profile for wireless authentication. You can add a RADIUS profile for accounting.

To configure RADIUS or LDAP servers for a wireless Access profile, enter the settings described in [Table 79](#). Required settings are indicated in the user interface by a red asterisk (*) that appears next to the field label.

1. Table 79: Server Group Settings

Field	Action
Configure Server Group	
Group Name	<p>Type a unique name that identifies the profile.</p> <p>You can use up to 32 characters for profiles created for wireless devices. Profile name must not contain special characters or spaces. Note that profiles that are automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character.</p>
Server Group Type	<p>Select the type of server group that you want to create:</p> <ul style="list-style-type: none"> • Accounting—Only RADIUS servers perform user accounting services, and they must also be configured for authentication to do accounting. • Authentication—Both RADIUS and LDAP servers perform user authentication services. • Both—Only RADIUS servers perform both user accounting services and authentication. <p>Server Type: Select either RADIUS or LDAP.</p> <p>TIP: LDAP servers do not perform accounting functions, so if you select Accounting for the Server Group Type, the only option here is RADIUS.</p>
Enable Load Balance	<p>Select to enable load-balancing for the servers that are part of the given server group.</p> <p>Load balancing enables the controller to distribute authentication requests across the authentication servers in a server group. Distributing the authentication process across multiple authentication servers significantly reduces the load on individual servers while increasing resiliency on a system-wide basis.</p>
Enable Command Audit	<p>Wireless controllers can log all CLI commands and all events. Select Enable Command Audit to capture and send every valid command to a RADIUS server log when the accounting command is enabled. The following information is captured: Timestamp, TTY Port, Username, Source IP address, Command issued, Command status (success or failure).</p>
RADIUS or LDAP Server Configuration	
The heading for this section depends on whether you selected RADIUS or LDAP for Server Type .	
View	Select any server entry from the list and then click View to see the details of that entry.

Table 79: Server Group Settings (continued)

Field	Action
Task: Configure a new RADIUS server for this Access profile	

Table 79: Server Group Settings (continued)

Field	Action
	<p>When doing RADIUS Server Configuration, both create and add a RADIUS server to this Access profile by following these steps:</p> <ol style="list-style-type: none"> a. Click Add > Create RADIUS. <ul style="list-style-type: none"> The Create RADIUS Server window opens. b. Complete these fields: <ul style="list-style-type: none"> ● Server Name—Type the name of the RADIUS server that you want to create. ● Server Address—Type the IP address of the RADIUS server. ● Authentication Port—The default RADIUS authentication port is 1812. You can change the port number by using the up and down arrows. ● Secret—Provide a password. If the password contains spaces, enclose it in quotation marks. The secret password used by the switch must match that used by the server. c. Optionally, if you selected either Accounting or Both for a Server Group Type, expand Advanced Settings and change the values for any of these fields: <ul style="list-style-type: none"> ● Accounting Port—Using the arrows, adjust the number of the UDP port to use for RADIUS accounting messages. The default UDP port is 1813, and the range is from 0 through 65535. ● Retry Count—Using the arrows, adjust the retry count until it reflects the number of times Network Director retries connecting to the RADIUS server when the RADIUS server is unavailable. Default is 3. ● Timeout—Using the arrows, adjust the timeout value. Timeout indicates how many seconds Network Director allows for RADIUS server connection before giving an unreachable error. Default is 5. ● Dead Time—Using the arrows, adjust the number of seconds before Network Director checks a RADIUS server that was previously unresponsive. The default value is 5 seconds. ● Use MAC Address as Password—Enable this option if you want each client device to use its MAC address as its password for the RADIUS server. If you enable Use MAC As Password, then the Authorization Password field becomes unavailable. ● Authorization Password—If you are not using MAC addresses as passwords for the RADIUS server, provide a common password here. ● MAC Address Format—Select None, Hyphens, Colons, One-Hyphen, or Raw to determine the MAC address format used with the RADIUS server. For descriptions and examples of these formats, see “Creating and Managing RADIUS Profiles” on page 339.

Table 79: Server Group Settings (continued)

Field	Action
	<ul style="list-style-type: none"> ● Authentication Protocol—Select PAP, CHAP, MSCHAP-V2, or None to determine an authentication protocol for the RADIUS server. These authentication protocols work as follows: <ul style="list-style-type: none"> PAP stands for Password Authentication Protocol and is used by Point to Point Protocols to validate users before allowing them access to server resources. Almost all network operating system remote servers support PAP. However, PAP transmits unencrypted ASCII passwords over the network and is therefore not secure. Use it as a last resort when the remote server does not support the stronger authentication. CHAP stands for Challenge Handshake Authentication Protocol and authenticates a user or network host to an authenticating entity. CHAP provides protection against replay attacks by the peer through the use of an incrementally changing identifier and of a variable challenge-value. CHAP requires that both the client and server know the plaintext of the secret password—it is never sent over the network. CHAP provides better security than PAP does. MSCHAP—stands for Microsoft's implementation of the Challenge Handshake Authentication Protocol version 2 on the router for password-change support. This feature provides users accessing a router the option of changing the password when the password expires, is reset, or is configured to be changed at the next login. The MS-CHAP variant does not require either peer to know the plaintext of the secret password. MSCHAP-V2 is used as an authentication option with RADIUS servers used for Wi-Fi security using the WPA-Enterprise protocol. <p>d. Click OK.</p> <p>The Create RADIUS Server window closes and the RADIUS server is automatically added to the list of authentication servers assigned to this Access profile.</p> <p>e. Optionally, if you have more than one RADIUS server listed, use the arrows to reorder the list priority so that the most preferred RADIUS server is listed first.</p>

Table 79: Server Group Settings (continued)

Field	Action
<p>Task: Add a previously configured RADIUS server for authentication</p>	<p>To add a previously configured RADIUS server to this Access profile:</p> <p>TIP: RADIUS servers are created by following the steps in “Creating and Managing RADIUS Profiles” on page 339.</p> <ol style="list-style-type: none"> a. Click Add > Select RADIUS. The Select RADIUS Server window opens. b. Select a RADIUS server from the available column—any RADIUS server that you created by using Add > Create RADIUS is listed here. c. Click the right arrow to move the highlighted RADIUS server from the Available column to the Selected column. d. Click OK. The Select RADIUS Server window closes and the RADIUS server is added to the list of authentication servers to be used with this Access profile.

Table 79: Server Group Settings (continued)

Field	Action
Task: Configure a new LDAP server for this Access profile	

Table 79: Server Group Settings (continued)

Field	Action
	<p>To both create and add an LDAP server to this Access profile:</p> <ol style="list-style-type: none"> a. Click Add > Create LDAP. <ul style="list-style-type: none"> The Create LDAP Server window opens. b. Complete these fields for an LDAP server: <ul style="list-style-type: none"> ● Server Name—Type the name of the LDAP server that you want to create. ● Server Address—Type the IP address of the LDAP server. ● Server Port—The default LDAP authentication port is 389. You can change the port number by using the up and down arrows. c. Optionally, expand the Advanced Settings section and change any of the following advanced LDAP server configuration: <ul style="list-style-type: none"> ● FQ Domain Name—A fully qualified domain name specifies an exact location in the tree hierarchy of the Domain Name System (DNS), including all domain levels such as the top-level domain and the root zone. A fully qualified domain name can be interpreted only one way. ● Dead Time—When a server does not respond for this number of seconds, it is removed from the list of authentication servers for this Access profile. The default dead time is 5 seconds. You can change this value by using the up and down arrows. ● Timeout—Adjust the length of time (default is 5 seconds) that elapses with no connection before Network Director gives an unreachable LDAP server error. You can change this value by using the up and down arrows to 1 through 90 seconds. ● Bind Mode—When an LDAP session is created (LDAP client connects to a server) the authentication state of the session is set to anonymous. BIND mode establishes the authentication state for a session and sets the LDAP protocol version. The default is Simple bind—you can change this to SASL-MD5. With Simple bind, the users' credentials are sent to the LDAP Directory Service in clear text. ● MAC Address Format—The default address format is None, which means that the MAC address is stated in a single stream (for example, 12ae53ef5676), with no subgrouping of the numbers. You can change this setting to Hyphens (for example, 12-ae-53-ef-56-76), Colons (for example, 12:ae:53:ef:56:76), One-Hyphen, or Raw. ● Base Domain—The top level of the LDAP directory tree is the base, referred to as the base DN. Enter a base domain name, for example, DC=eng, DC=Juniper Networks, or DC=com. This string indicates where to load users and groups.

Table 79: Server Group Settings (*continued*)

Field	Action
	<ul style="list-style-type: none"> ● Domain Prefix—AD or NT domains use the NetBIOS name. Default is cn. ● Use MAC as Password—Select this option to use the MAC address of devices as the password for authentication purposes. ● Authorization Password—If MAC addresses are not used as passwords, provide a password to be used for authentication purposes. <p>d. Click OK.</p> <p>The server is automatically added to the list of authentication servers assigned to this Access profile.</p> <p>e. Optionally, if you have more than one server listed, use the arrows to reorder the list priority so that the most preferred server is listed first.</p> <p>f. Click OK.</p> <p>The server is added to the list of authentication servers to be used with this Access profile.</p>
Task: Add a previously configured LDAP server for authentication	<p>To add a previously configured LDAP server to this Access profile:</p> <p>a. Click Add > Select LDAP.</p> <p>b. Select an LDAP server from the Available column—any LDAP server that you created using Add > Create LDAP is listed here in addition to any created with “Creating and Managing LDAP Profiles” on page 345.</p> <p>c. Click the right arrow to move the highlighted LDAP server from the Available column to the Selected column.</p> <p>d. Click OK.</p> <p>The LDAP server is added to the list of authentication servers to be used with this Access profile.</p>
Task: Delete any server	<p>To delete a server from this Access profile:</p> <p>a. Select a server from the list.</p> <p>b. Click Delete.</p> <p>The server is removed from the list of authentication servers to be used with this Access profile.</p>

The system adds the server details to the Server Configuration table.

NOTE: Use the UP and DOWN arrows to reorder the server groups. User authentication is first attempted with the server listed first. If that authentication fails, the next method on the list is used.

Proceed to the review for wireless Access profiles by either clicking **Review** or by clicking **Next**. For directions for this section, see [“Reviewing and Modifying the Access Profile Settings”](#) on page 380.

Specifying Basic Settings for a Campus Switching ELS Access Profile

To configure the basic settings for a Campus Switching ELS Access profile:

1. Complete the basic settings and authentication order on the Create Access Profile for Campus Switching ELS page, as described in both the online help and in [Table 80](#). Required settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 80: Access Profile Basic Settings for Campus Switching ELS

Field	Action
Access Profile Details	
Profile Name	Type a unique name that identifies the profile. You can use up to 32 characters for profiles created for wireless devices. Profile name must not contain special characters or spaces. Note that profiles that are automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character.
Description	Type the description of the profile.
Authentication Order	
Server settings depend on which authentication is done first, RADIUS or LDAP.	
Authentication Order	Indicate whether to authenticate first with configured RADIUS servers or with configured LDAP servers by selecting the method from <i>Based On</i> . By default, RADIUS authentication using no password is selected for initial authentication. You can change this to RADIUS authentication with a password by selecting Password . Select LDAP to authenticate first with configured LDAP servers. TIP: LDAP is not supported for Data Center or EX Switching devices.

Proceed to the Server Settings for Campus Switching ELS Access profiles by clicking either **Server Settings** or **Next**. The settings are described in “[Specifying RADIUS and LDAP Settings for Campus Switching ELS](#)” on page 372.

Specifying RADIUS and LDAP Settings for Campus Switching ELS

Configure either a RADIUS server, an LDAP server, or both, on the Server Settings page. A RADIUS server can provide both user accounting services and user authentication but you must be using the RADIUS

server for authentication in order to use it for accounting. An LDAP server provides only user authentication. The server settings in this section determine the options used for the access servers in this Access profile.

Configure the Server settings for a Campus Switching ELS Access profile by following the directions in [Table 81](#).

Table 81: Authentication and Accounting Server Settings for ELS Campus Switching

Task	Action
<p>AAA: Authentication Server</p>	
<p>RADIUS servers are selected for configuration by default. RADIUS servers can do both authentication and accounting.</p>	
<p>View configured servers in this profile</p>	<p>Select a server entry from the list and then click View to see the details of that entry.</p>

Table 81: Authentication and Accounting Server Settings for ELS Campus Switching (continued)

Task	Action
Create and add a new RADIUS server for authentication	<p>The RADIUS tab is selected by default for AAA Authentication Server configuration. To configure a RADIUS accounting server and add it to this Access profile:</p> <ol style="list-style-type: none"> 1. Click Add > Create RADIUS on the RADIUS tab. The Create RADIUS Server window opens. 2. Provide the following RADIUS authentication server information: <ul style="list-style-type: none"> • Server Name • Server Address • Authentication Port—The default RADIUS authentication port is 1812. You can change the port number by using the up and down arrows. • Secret—Provide the authentication secret password. If the password contains spaces, enclose it in quotation marks. The secret password used by the local router must match the one used by the server. 3. Optionally, expand the Advanced Settings for a RADIUS server and change any of these configurations: <ul style="list-style-type: none"> • Accounting Port—You can change the default accounting port number (1813) by using the up and down arrows. • Retry Count—Specify the number of times that a device attempts to contact the LDAP authentication server. The default retry count is 3. You can change this value by using the up and down arrows to 1 through 10 times. • Timeout (seconds)—Specify the number of seconds the switch waits to receive a response from a RADIUS server. The default timeout is 5 seconds. You can change this value, using the up and down arrows, to 1 through 90 seconds. 4. Click OK. The Create RADIUS Server window closes and the RADIUS server is automatically added to the list of RADIUS servers assigned to this Access profile. 5. If you have more than one RADIUS server listed, you can use the arrows to reorder the list priority so that the most preferred RADIUS server is listed first.

Table 81: Authentication and Accounting Server Settings for ELS Campus Switching (continued)

Task	Action
Add a previously configured RADIUS server for authentication	<p>The RADIUS tab is selected by default for server configuration and configured RADIUS servers are listed on this Server Settings page. To add a previously configured RADIUS server to this Access profile:</p> <ol style="list-style-type: none">1. Click Add > Select RADIUS on the RADIUS tab. The Select RADIUS Server window opens, displaying a list of available RADIUS servers is displayed. Servers on this list were either automatically discovered, created following the directions "Creating and Managing RADIUS Profiles" on page 339, or created on this page following the directions in <i>Create and add a new RADIUS server configuration</i>.2. Select one or more RADIUS servers from the list of previously configured RADIUS servers.3. Click OK. The Select RADIUS Server window closes and the RADIUS server is added to the list of RADIUS authentication servers to be used with this Access profile.4. Optionally, if you have more than one RADIUS server listed, use the arrows to reorder the login priority so that the most preferred RADIUS server is listed first.

Table 81: Authentication and Accounting Server Settings for ELS Campus Switching (continued)

Task	Action
Add a previously configured RADIUS server for accounting	<p>A RADIUS server can provide both authentication and accounting. To configure accounting settings for a RADIUS server:</p> <p>TIP: In order to provide accounting, authentication must also be configured.</p> <ol style="list-style-type: none"> Expand the RADIUS Accounting Servers section of the Server Settings. A list of RADIUS servers configured for accounting is displayed. Click Add > Select RADIUS. The Select RADIUS Server window opens, displaying a list of eligible RADIUS servers is displayed. Servers on this list were either automatically discovered, created following the directions “Creating and Managing RADIUS Profiles” on page 339, or created on this page following the directions <i>Create and add a new RADIUS server configuration</i>. Select one or more RADIUS servers from the list of previously configured RADIUS servers. Click OK. The Select RADIUS Server window closes and the RADIUS server is added to the list of RADIUS Accounting Servers to be used with this Access profile. Optionally, if you have more than one RADIUS server listed, use the arrows to reorder the login priority so that the most preferred RADIUS server is listed first.

Table 81: Authentication and Accounting Server Settings for ELS Campus Switching (continued)

Task	Action
Create and add a new RADIUS server for both authentication and accounting	<p>RADIUS is the only server selection available for accounting. To configure a RADIUS server for both authentication and accounting, and add it to this Access profile:</p> <ol style="list-style-type: none"> 1. Under RADIUS Accounting Server, click Add > Create RADIUS. The Create RADIUS Server window opens. 2. Provide the following RADIUS authentication server information: <ul style="list-style-type: none"> • Server Name • Server Address • Authentication Port—The default RADIUS authentication port is 1812. You can change the port number by using the up and down arrows. • Secret—Provide the authentication secret password. If the password contains spaces, enclose it in quotation marks. The secret password used by the local router must match that used by the server. 3. Expand the Advanced Settings and change any of these configurations: <ul style="list-style-type: none"> • Accounting Port—You can change the default port number (1813) by using the up and down arrows. • Retry Count—Specify the number of times that a device attempts to contact the LDAP authentication server. The default retry count is 3. You can change this value by using the up and down arrows to 1 through 10 times. • Timeout (seconds)—Specify the number of seconds the switch waits to receive a response from a RADIUS server. The default timeout is 5 seconds. You can change this value, using the up and down arrows, to 1 through 90 seconds. 4. Click OK. The Create RADIUS Server window closes and the RADIUS server is automatically added to the list of RADIUS Accounting Servers assigned to this Access profile. 5. If you have more than one RADIUS accounting server listed, you can use the arrows to reorder the list priority so that the most preferred RADIUS server is listed first.

Table 81: Authentication and Accounting Server Settings for ELS Campus Switching (continued)

Task	Action
Create and add a new LDAP authentication server	<p>TIP: LDAP servers can be configured for wireless and for Campus Switching ELS.</p> <p>To configure a new LDAP authentication server and add it to this Access profile:</p> <ol style="list-style-type: none"> 1. Click the LDAP tab to display the LDAP settings. 2. Provide a Base Distinguished Name for the LDAP server. LDAP APIs reference an LDAP object by its distinguished name (DN), which is a sequence of relative distinguished names (RDN) connected by commas—for example, DC=eng, DC=Juniper Networks, DC=com. You can do an LDAP query to determine the DN for the LDAP server. 3. Click Add > Create LDAP. The Create LDAP Server window opens. 4. Provide the following LDAP server information: <ul style="list-style-type: none"> • Server Name • Server Address • Server Port—The default LDAP server port is 389. You can change the port number by using the up and down arrows. 5. Optionally provide the following Advanced LDAP server information after expanding the Advanced Settings section: <ul style="list-style-type: none"> • Timeout (seconds)—Specify the number of seconds the switch waits to receive a response from a RADIUS server. The default timeout is 5 seconds. You can change this value, using the up and down arrows, to 1 through 90 seconds. • Retry—Specify the number of times that a device attempts to contact the LDAP authentication server. The default retry count is 5. You can change this value by using the up and down arrows to 1 through 10 times. 6. Click OK. The Create LDAP Server window closes and the LDAP server is added to the list of LDAP servers.

Table 81: Authentication and Accounting Server Settings for ELS Campus Switching (continued)

Task	Action
Add a previously configured LDAP server for authentication	<p>TIP: LDAP servers can be configured for wireless and for Campus Switching ELS.</p> <p>To add a previously configured LDAP authentication server to this Access profile:</p> <ol style="list-style-type: none"> 1. Click the LDAP tab to display the LDAP settings. 2. Provide a Base Distinguished name for the LDAP server. LDAP APIs reference an LDAP object by its distinguished name (DN), which is a sequence of relative distinguished names (RDN) connected by commas. You can do an LDAP query to determine the DN for the LDAP server. 3. Click Add > Select LDAP. The Select LDAP Server window opens, displaying a list of configured LDAP servers displayed. Servers on this list were either automatically discovered, or created following the directions “Creating and Managing LDAP Profiles” on page 345, or created by clicking Add > Create LDAP on this page. 4. Select one or more LDAP servers from the list. 5. Click OK. The Select LDAP Server window closes and selected LDAP servers are added to the list of LDAP authentication servers to be used with this Access profile. 6. Optionally, use the arrows to reorder the LDAP servers so that the most preferred LDAP server is listed first. <p>TIP: LDAP is not supported for Data Center or EX Switching devices.</p>
Delete a server	<p>To delete any server from this Access profile:</p> <ol style="list-style-type: none"> 1. Select a server from the list. 2. Click Delete. The server is removed from the list of servers to be used with this Access profile.

Proceed to the review for wireless Access profiles by either clicking **Review** or by clicking **Next**. For directions for this section, see [“Reviewing and Modifying the Access Profile Settings” on page 380](#).

Reviewing and Modifying the Access Profile Settings

From this page, you can save or make changes to a Access profile:

- To make changes to the profile, click **Edit** associated with the configuration to be changed.

Alternatively, you can click the appropriate sections in the profile workflow at the top of the page that corresponds to the configuration to be changed.

When you are finished with your modifications, click **Review** to return to this page.

- To save a new profile or to save modified settings to an existing profile, click **Finish**.

You will be returned to the Manage Access Profiles page. Your new or modified Access profile is listed in the table of Access profiles.

What To Do Next

After you create an Access profile, you can do one of the following:

- For wireless devices, link the Access profile to an Authentication profile that is created for the same device family. For more information see [“Creating and Managing Authentication Profiles” on page 383](#).
- For switching devices, configure Access profile as a attribute while assigning Port profiles to interfaces. For more information see [“Creating and Managing Port Profiles” on page 414](#).

NOTE: Assigned settings from any profile, including this one, have lower priority than settings made directly to a controller or an access point. For more information, see [“Adding and Managing an Individual Access Point” on page 1156](#) and [“Configuring a Controller” on page 1037](#).

RELATED DOCUMENTATION

[Understanding Access Profiles | 351](#)

[Creating and Managing LDAP Profiles | 345](#)

[Creating and Managing RADIUS Profiles | 339](#)

[Creating and Managing Authentication Profiles | 383](#)

[Creating and Managing Port Profiles | 414](#)

[Network Director Documentation home page](#)

Understanding Authentication Profiles

IN THIS SECTION

- [802.1X Authentication | 381](#)
- [MAC RADIUS Authentication | 382](#)
- [Captive Portal Authentication | 382](#)
- [Last Resort Authentication | 382](#)

Authentication profiles include the authentication method and authentication parameters to be used for client authentication. Available authentication methods are 802.1X (dot1x), MAC-RADIUS, captive portal, and last-resort. You can configure last-resort authentication only on wireless devices. 802.1X is the default authentication method for all device types but you can change this or add additional authentication types. If you configure multiple authentication methods on a single interface, the system tries the first method listed and then falls back to another method if the first method is unsuccessful.

You can create one or more Authentication profiles to specify different authentication methods based on client devices or sessions.

Each Authentication profile is specific to a device family. After you create an Authentication profile, you can include it in a WLAN Service profile or a Port profile. The Authentication profile specified in a WLAN profile or a Port profile is used to authenticate all the users and devices that connect to that WLAN or on the port.

802.1X Authentication

Newer equipment supports the IEE standard called 802.1X. 802.1X is basically an Enterprise, per-user (username and password) authentication mechanism – it is both the newest and strongest authentication you can use. Since 802.1X authentication is the most secure authentication option, it is preferable to the older PSK authentication, Web Portals, MAC authentication, or open authentication, which really means no authentication.

802.1X authentication involves three entities, a supplicant, an authenticator, and an authentication server. The supplicant is a client device, such as a laptop, that wishes to attach to a network. The authenticator would be either a switch or an access point. The authentication server is usually a RADIUS server, which can interpret 802.1X EAP modes.

- *Single supplicant mode* authenticates only the first end device that connects to an authenticator port. All other end devices connecting to the authenticator port after the first has connected successfully, whether they are 802.1X-enabled or not, are permitted free access to the port without further authentication.

If the first authenticated end device logs out, all other end devices are locked out until an end device authenticates.

- *Single-secure supplicant mode* authenticates only one end device to connect to an authenticator port. No other end device can connect to the authenticator port until the first logs out.
- *Multiple supplicant mode* authenticates multiple end devices individually on one authenticator port. If you configure a maximum number of devices that can be connected to a port through port security, the lesser of the configured values is used to determine the maximum number of end devices allowed per port.

MAC RADIUS Authentication

A media access control address (MAC address) is a unique identifier assigned to network interfaces for communications on the physical network segment. MAC addresses are used as a network address for most IEEE 802 network technologies, including Ethernet. A client's MAC address can be used for authentication by mapping a password to the client's entry in the MAC address table. MAC authentication can be done either locally or with a RADIUS server.

Captive Portal Authentication

Captive Portals are frequently used to authenticate hotspots, forcing all users to use the configured logon web page. Many companies use captive portals to authenticate guest users for temporary use of the company network. The Captive Portal has one password for all users, which should be changed frequently.

Last Resort Authentication

You can configure last-resort authentication only on wireless devices.

RELATED DOCUMENTATION

[Creating and Managing Authentication Profiles | 383](#)

[Network Director Documentation home page](#)

Creating and Managing Authentication Profiles

IN THIS SECTION

- [Managing Authentication Profiles | 383](#)
- [Creating an Authentication Profile | 384](#)
- [Specifying Authentication Settings for Switches | 386](#)
- [Specifying Authentication Settings for Wireless | 390](#)
- [What To Do Next | 394](#)

Authentication profiles enable specification of the authentication method and authentication parameters to be used for authenticating clients and users who connect to a WLAN or an access port switch.

Use the Manage Authentication Profiles page to create new Authentication profiles and manage existing Authentication profiles.

To display the Manage Authentication Profiles page: In Build mode, select Authentication from Profile and Configuration Management in the Tasks pane. The Manage Authentication Profiles page appears.

This topic describes:

Managing Authentication Profiles

From the Manage Authentication Profiles page, you can:

- Create a new Authentication profile by clicking **Add**. For directions, see [“Creating an Authentication Profile” on page 384](#).
- Modify an existing profile by selecting it and clicking **Edit**.
- View information about a profile, including the interfaces it is associated with, by clicking the profile name or by selecting the profile and clicking **Details**.
- Delete an Authentication profile by selecting a profile and clicking **Delete**.

TIP: You cannot delete profiles that are in use—that is, assigned to objects or used by other profiles. To see the current assignments for a profile, select the profile and click **Details**.

- Clone a profile by selecting a profile and clicking **Clone**.

Table 82 describes the information provided about Authentication profiles on the Manage Authentication Profiles page. This page lists all Authentication profiles defined for your network, regardless of the scope you selected in the network view.

Table 82: Manage Authentication Profile Fields

Field	Description
Profile Name	Name given to the profile when the profile was created.
Family Type	The device family on which the profile was created.
Description	Description of the profile that was entered when the profile was created. TIP: To display the entire description, you might need to resize the Description column by clicking the column border in the heading and dragging it.
Creation Time	Date and time when this profile was created.
Update Time	Date and time when this profile was last modified.
User Name	The username of the user who created or modified the profile.

TIP: All columns might not be displayed. To show or hide fields in the Manage Authentication Profiles table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

Creating an Authentication Profile


In Network Director, you can create an Authentication profile to configure methods to be used to authenticate users. You can also specify details about the accounting servers to be used for accounting purposes.

For an Authentication profile, you must specify the following:

- A profile name
- At least one access rule

After you create an Authentication profile, you can include it in a WLAN profile or a Port profile. The Authentication profile specified in a WLAN profile or a Port profile acts as the default profile for all the users and devices that connect to that WLAN or on the port.

To create an Authentication profile:

1. Click  in the Network Director banner.
2. Under Select View, select either **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View** or **Topology View**.

3. From the Tasks pane, select the type of network (Wired or Wireless), the appropriate functional area (System, AAA, or Wireless), and select the name of the profile that you want to create. For example, to create a radius profile for a wireless device, click **Wireless > AAA > Radius**. The Manage Profile page opens.
4. Click **Add** to add a new profile.

If you chose to create a profile for the wired network, Network Director opens the Device Family Chooser window.

- a. From the Device Family Chooser, select the device family for which you want to create a profile. The available device families are **Switching (EX)**, **Campus Switching ELS** (Enhanced Layer 2 Software), **Data Center Switching Non ELS** and **Data Center Switching ELS**.

- b. Click **OK**.

The Create Authentication Profile page for the selected device family is displayed.

If you chose to create a profile for the wireless network, Network Director opens the Create Authentication Profile for Wireless page.

5. Specify authentication settings by doing one of the following:
 - For EX Series switches, Campus Switching Enhanced Layer 2 Software, or Data Center Switching profiles, specify the settings as described in [“Specifying Authentication Settings for Switches” on page 386](#).
 - For controllers, specify the settings as described in [“Specifying Authentication Settings for Wireless” on page 390](#).
6. Click **Done** to save the Authentication profile.

The system saves the Authentication profile and displays the Manage Authentication Profiles page. Your new or modified Authentication profile is listed in the table of Authentication profiles.

Specifying Authentication Settings for Switches

To configure an Authentication profile for switching devices, enter the Create Authentication Profile page settings described in [Table 83](#) for creating Authentication profiles on switches. Required settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 83: Authentication Profile Settings for Switches

Field	Action
Profile Name	Type the name of the profile. You can use up to 64 characters for profiles created for wired devices. Profile name must not contain special characters or spaces. Note that profiles that are automatically created by Network Director as part of device discovery or out-of-band changes may contain the underscore (_) character.
Description	Type a short description for the profile.
802.1X Authenticator	
Enable 802.1X	802.1X authentication is enabled by default for a switching profile. 802.1X authentication works by using an Authenticator Port Access Entity (the switch) to block all traffic to and from a supplicant (end device) at the port until the supplicant's credentials are presented and matched on the Authentication server (a RADIUS server). When authenticated, the switch stops blocking traffic and opens the port to the supplicant. Network access can be further defined using VLANs. NOTE: If you disable 802.1X authentication, several related settings become unavailable.
Enable MAC-RADIUS	Select to enable MAC-RADIUS based authentication for this profile. MAC RADIUS authentication enables LAN access to permitted MAC addresses. When a new MAC address appears on an interface, the switch consults the RADIUS server to check whether the MAC address is a permitted address. If the MAC address is configured on the RADIUS server, the device is allowed access to the LAN. TIP: You can combine 802.1X and MAC-RADIUS authentication.

Table 83: Authentication Profile Settings for Switches (*continued*)

Field	Action
Supplicant Mode	<p>Specify the mode authentication supplicants use, either Single, Multiple, or Single-Secure.</p> <ul style="list-style-type: none"> • Single—Allows only one host for authentication. • Single-Secure—Allows only one end device to connect to the port. No other end device is enabled to connect until the first logs out. • Multiple—Allows multiple hosts for authentication. Each host is checked before being admitted to the network.
Guest VLAN	<p>Click Select and then select the VLAN to which an interface is moved when no 802.1X supplicants are connected on the interface. The VLAN specified must already exist on the switch.</p>
Reject VLAN	<p>Click Select and then select the VLAN to which an interface is moved when the switch receives an Extensible Authentication Protocol Over LAN (EAPoL) Access-Reject message during the authentication process between the switch and the RADIUS authentication server.</p>
Server Fail Type	<p>Specify the server fail fallback action the switch takes when all RADIUS authentication servers are unreachable, either None, Deny, Permit, Use cache, or VLAN Name.</p> <ul style="list-style-type: none"> • Deny—Force fail the supplicant authentication. No traffic will flow through the interface. • Permit—Force succeed the supplicant authentication. Traffic will flow through the interface as if it were successfully authenticated by the RADIUS server. • Use cache—Force succeed the supplicant authentication only if it was previously authenticated successfully. This action ensures that already authenticated supplicants are not affected. • VLAN Name—Move supplicant on the interface to the VLAN specified by this name. This action is allowed only if it is the first supplicant connecting to an interface. If an authenticated supplicant is already connected, then the supplicant is not moved to the VLAN and is not authenticated. If you select this option, also provide a Fail VLAN name.

Captive Portal

A Captive Portal is a special web page used for authentication by turning a web browser into an authentication mechanism.

Enable Captive-Portal	<p>Enable this option to display the captive portal setting for supplicant mode. When this option is enabled, additional captive portal settings are also available under Advanced Settings.</p>
-----------------------	--

Table 83: Authentication Profile Settings for Switches (*continued*)

Field	Action
Supplicant Mode	<p>Specify the mode to be used for Captive Portal supplicants, either Single, Multiple, or Single-Secure.</p> <ul style="list-style-type: none"> • Single—Allows only one host for authentication. • Multiple—Allows multiple hosts for authentication. Each host is checked before being admitted to the network. • Single-Secure —Allows only one end device to connect to the port. No other end device is allowed to connect until the first logs out.

To skip configuring the advanced settings and accept the default settings, click **Done**. You can now link the Authentication profile to a Port profile. For directions, see [“Creating and Managing Port Profiles” on page 414](#).

To configure advanced switch settings, click **Advanced Settings** and enter the Advanced Settings described in [Table 84](#).

Table 84: Authentication Profile Advanced Settings for Switches

Field	Action
802.1X Settings	
These settings are available only when 802.1X authentication is enabled for this Authentication profile. You can use the default settings or you can change them.	
Transmit Period (default is 30 seconds)	Specify how long, in seconds, the interface waits before retransmitting the initial EAPOL PDUs to the supplicant. The default is 30 seconds.
Maximum Requests (default is 2 requests)	Specify the maximum number of times an EAPOL request packet is transmitted to the supplicant before the authentication session times out. The default is 2 requests.
Retries (default is 3 retries)	Specify the number of times you want the switch to attempt to authenticate the port after an initial failure. The port remains in a wait state during the quiet period after the authentication attempt. The default is 3 retries.
Quiet Period (default is 60 seconds)	Specify the number of seconds the interface remains in the wait state following a failed authentication attempt by a supplicant before reattempting authentication. The default is 60 seconds.
No Reauthentication (default is unselected)	Select this check box if you do not want the switch to reauthenticate the supplicant after the Quiet Period elapses.

Table 84: Authentication Profile Advanced Settings for Switches (continued)

Field	Action
Reauthentication Interval (default is 3600 seconds)	If the No Reauthentication option is not checked, specify the number of seconds after which the authentication session times out. The default is 3600 seconds.
Supplicant Timeout (default is 30 seconds)	Specify how long the port waits for a response when relaying a request from the authentication server to the supplicant before resending the request. The default is 30 seconds.
RADIUS Server Timeout (default is 30 seconds)	Specify the length of time that the switch waits for a response from the RADIUS server. The default is 30 seconds.
MAC Restrict (Switches using MAC RADIUS only)	<p>When MAC-RADIUS is enabled in this Authentication profile, select this option to restrict authentication to MAC RADIUS only. When MAC-RADIUS restrict is configured, the switch drops all 802.1X packets. This option is useful when no other 802.1X authentication methods, such as guest VLAN, are needed on the interface, and eliminates the delay that occurs while the switch determines that a connected device is a non-802.1X-enabled host.</p> <p>Optionally enable Flap-On-Disconnect. When the RADIUS server sends a disconnect message to a supplicant, the switch resets the interface on which the supplicant is authenticated. If the interface is configured for multiple supplicant mode, the switch resets all the supplicants on the specified interface. This option takes effect only when the MAC Restrict option is also set.</p>

Captive Portal

If Captive Portal is enabled in this Authentication profile in the basic settings, you can either use the default advanced Captive Portal settings or change them as indicated.

Quiet Period (default is 60 seconds)	<p>Configure the time, in seconds, between when a user exceeds the maximum number of retries and when they can again attempt to authenticate.</p> <p>Range: 1 through 65,535</p> <p>Default: 60</p>
Retries (default is 3 retries)	<p>Configure the number of times the user can attempt to submit authentication information.</p> <p>Range: 1 through 65,535</p> <p>Default: 3</p>

Table 84: Authentication Profile Advanced Settings for Switches (continued)

Field	Action
Session Expiry (default is 3600 seconds)	Configure the maximum duration in seconds of a session. Range: 1 through 65,535 Default: 3600
Server Time Out (default is 30 seconds)	Configure the time in seconds an interface will wait for a reply when relaying a response from the client to the authentication server before timing out and invoking the server-fail action. Range: 1 through 65,535 Default: 30

Click **OK**.

The Advanced Settings window closes and you once again see the Create Authentication Profile for Switching page.

Click **Done**.

The Manage Authentication Profiles page reappears with your new Authentication profile listed.

You can now link the Authentication profile to a Port profile. For more details, see [“Creating and Managing Port Profiles” on page 414](#).

Specifying Authentication Settings for Wireless

While configuring an Authentication profile for wireless devices, you define one or more access rules. Each access rule is specific to an access type or authentication mechanism, such as 802.1X, MAC, Web, and open authentication. All authentication mechanisms are supported in a chain and are allowed in any sequence with one exception—Web authentication and Open authentication must not be configured simultaneously in one Authentication profile.

To configure an Authentication profile for wireless:

1. Enter the wireless authentication settings described in [Table 85](#). Required settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 85: Authentication Profile Wireless Settings

Field	Description
Profile Name	Type the name of the profile. You can use up to 32 characters for profiles created for wireless devices. Profile name must not contain special characters or spaces. Note that profiles that are automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character.
Description	Type the description of the profile.

2. Add at least one access rule by clicking **Add** under Access Rule.

The Add Access Rules window opens.

3. Enter the access rule settings described in [Table 86](#). Required settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 86: Wireless Access Rule Settings

Field	Description
Access Type	Select the type of access for the rule, either 802.1X Access (default), MAC Access , Web Access , or Open Access : <ul style="list-style-type: none"> ● 802.1X Access—Select to authenticate the client using 802.1X authentication method. For more information, see “Understanding the IEEE 802.11 Standard for Wireless Networks” on page 1076. ● MAC Access—Select to authenticate the client using MAC RADIUS authentication method. ● Web Access—Select to have the client log in to a web page before granting access to the SSID. ● Open Access—Select to automatically authenticate the client and enable access to the SSID requested by the client, without requiring a username and password from the client. <p>If you select open access as the access type, you must either <i>enable local accounting</i> or <i>specify an Access profile</i> to be able to save the access rule.</p> <p>The remaining options in this window vary, depending on which Access Type you choose.</p>

Table 86: Wireless Access Rule Settings (continued)

Field	Description
Matching Glob (all)	<p>Type the user glob for the access rule.</p> <p>A user glob is shorthand method for matching an authentication, authorization, and accounting (AAA) command to either a single user or a set of users.</p> <p>A user glob can contain up to 80 characters long and cannot include spaces or tabs. The double-asterisk (**) wildcard characters with no delimiter characters match all user names. The single-asterisk (*) wildcard character matches any number of characters up to, but not including, a delimiter character in the glob. Valid user glob delimiter characters are the at (@) sign and the period (.).</p> <p>NOTE: The matching glob value that you specify must be unique and cannot be used for any other access rules within the given authentication profile.</p>
EAP Type (801.X Access)	<p>If you selected 802.1X Access, you also need to indicate an EAP type. Extensible Authentication Protocol (EAP) is a generic point-to-point protocol that supports multiple authentication mechanisms. Select the EAP type that you want to use for this access rule:</p> <ul style="list-style-type: none"> ● PEAP Offload—Select if you want to offload all EAP processing from server groups. In this case, the RADIUS server is not required to communicate using the EAP protocols. ● Local EAP-TLS—Select if you want to use a local database to authenticate clients. EAP-TLS provides encryption and data integrity checking for the connection. Local EAP-TLS can only be used with Local Authentication. ● External Authentication Server(default)—Select if you want to use an external server for authentication. <p>NOTE: If you select to use an external server for authentication, you must not select the Enable Local Authentication.</p>
Enable Authentication	<p>Selected by default to enable authentication for this access rule.</p>
Enable Local Authentication	<p>Select to have users authenticated against the local database on the controller.</p> <p>TIP: Network Director displays this check box only if you selected Enable Authentication.</p>
MAC Prefix (MAC Access)	<p>If MAC Access is the access type, you can enable MAC prefix authentication.</p>

Table 86: Wireless Access Rule Settings (continued)

Field	Description
Enable Accounting	<p>Select to enable accounting for this access rule and display the accounting settings. Accounting collects and sends information used for billing, auditing, and reporting. Accounting can be done using RADIUS or by using local accounting.</p> <p>If Enable Accounting is selected, you can configure these additional accounting parameters:</p> <ul style="list-style-type: none"> ● Enable Local Accounting—Select if you want to enable local accounting. If you select local accounting, the accounting information is stored locally on the controller. ● Record Type—Select the local accounting mode to be used for this access rule: <ul style="list-style-type: none"> ● Start-Stop—When this mode is selected, a start record is generated when a user is first connected, and an update record is generated when a user roams from one wireless access point to another. A stop record is generated when a user terminates the session. ● Stop-Only—When this mode is selected, a stop record is generated when a user terminates the session.
Access Profile	<p>Specify an Access profile (default is None) to use for this access rule. Network Director displays the Access Profile field when you have enabled authentication, accounting, or both, for the given access rule. Specify an Access profile for each access rule, unless:</p> <ul style="list-style-type: none"> ● Open Access is the access type. ● 802.1X Access is the access type and Local EAP-TLS as the EAP type.

4. Click **OK** to save and add the wireless access rule to the list of access rules in the Create Authentication Profile page. You can create one or more access rules and authenticate each user or device group differently depending on your security policy and requirements.
5. You can configure a wireless web portal by providing the information listed in [Table 87](#).

Table 87: Wireless Web Portal Authentication

Web Portal Settings

A web portal is a web site that brings information together from diverse sources in a uniform way. The following Web Portal settings are valid only when the corresponding WLAN Service Profile has Fall Through Access set to **web-portal**. See [“Creating and Managing a WLAN Service Profile” on page 1090](#) for more information and directions.

Table 87: Wireless Web Portal Authentication (continued)

Web Portal Login Page	<p>Type the name of a Web portal login page, for example, <i>http://www.example.com</i>. Web traffic will be directed to this location for logging into the Web portal. For more information about Web login, see “Configuring Web Portal Web AAA” in the <i>Mobility System Software Configuration Guide</i>.</p> <p>TIP: Web Portal settings are valid only when Fall Through Access is set to WEB Portal in a WLAN Service Profile. When you link this Authentication Profile to a WLAN Service Profile with Fall Through Access set to WEB Portal, the settings are used.</p>
Web Portal Logout	<p>Optionally, enable Web Portal Logout, and provide the name of the Web Portal Logout Page, for example, <i>http://www.example.com</i>.</p>

6. Click **Done**.

The Manage Authentication Profiles for Wireless (WLC) page reappears with your new wireless Authentication profile listed. You can now link the Authentication profile to a WLAN profile. For more details, see [“Creating and Managing a WLAN Service Profile” on page 1090](#).

What To Do Next

After you create an Authentication profile, you can do the following:

- For switching devices, link the Authentication profile to a Port profile. For more details, see [“Creating and Managing Port Profiles” on page 414](#).
- For wireless devices, link the Authentication profile to a WLAN profile. For more details, see [“Creating and Managing a WLAN Service Profile” on page 1090](#).

NOTE: Assigned settings from any profile, including this one, have lower priority than settings made directly to a controller or an access point. For more information, see [“Adding and Managing an Individual Access Point” on page 1156](#) and [“Configuring a Controller” on page 1037](#).

RELATED DOCUMENTATION

[Understanding Authentication Profiles | 381](#)

[Creating and Managing Port Profiles | 414](#)

[Creating and Managing a WLAN Service Profile | 1090](#)

[Network Director Documentation home page](#)

Understanding Wireless Authorization Profiles

Once the user is authenticated, Network Director looks for the authorization attributes assigned to the user. Authorization profiles specify the access permissions and authorization attributes for authenticated users or devices. Authorization attributes specify the network resources available to the user. The VLAN profile is a mandatory attribute in order to place the user in the appropriate VLAN. You can provide further access controls using an Authorization profile by specifying encryption types, linking a class-of-service (CoS) profile, specifying the days and times during which the user can access the network, and so on.

You can create Authorization profiles only for wireless devices. After you create a Authorization profile, you can link it to a WLAN profile or assign it to a controller or a cluster. Once you assign an Authorization profile to a controller or a cluster, the VLAN, CoS, and Filter settings that you have defined within that profile are applied on the WLC or the cluster. These settings can then be used for dynamic users who are authenticated through RADIUS or LDAP servers.

RELATED DOCUMENTATION

[Creating and Managing Wireless Authorization Profiles | 395](#)

[Network Director Documentation home page](#)

Creating and Managing Wireless Authorization Profiles

IN THIS SECTION

- [Managing Authorization Profiles | 396](#)
- [Creating a Wireless Authorization Profile | 397](#)
- [Specifying Settings for a Wireless Authorization Profile | 399](#)
- [What To Do Next | 403](#)

Authorization profiles specify the access permission for authenticated users or devices.

Use the Manage Authorization Profiles page to create new wireless Authorization profiles and manage existing wireless Authorization profiles.

This topic describes:

Managing Authorization Profiles

From the Manage Authorization Profiles page, you can:

- Create a new wireless Authorization profile by clicking **Add**. For directions, see [“Creating a Wireless Authorization Profile” on page 397](#).
- Modify an existing Authorization profile by selecting it and clicking **Edit**.
- Associate an Authorization profile to specific devices or clusters by selecting it and clicking **Assign**. For directions, see [“Assigning Wireless Authorization Profiles to Controllers” on page 404](#).
- Change the current assignment of an Authorization profile by selecting it and clicking **Edit Assignment**. For directions, see [“Assigning Wireless Authorization Profiles to Controllers” on page 404](#).
- View information about an Authorization profile, including the interfaces it is associated with, by clicking the profile name or by selecting the profile and clicking **Details**.
- Delete an Authorization profile by selecting a profile and clicking **Delete**.

TIP: You cannot delete profiles that are in use—that is, assigned to objects or used by other profiles. To see the current assignments for an Authorization profile, select the Authorization profile and click **Details**.

- Clone a profile by selecting a profile and clicking **Clone**.

[Table 88](#) describes the information provided about Authorization profiles on the Manage Authorization Profiles page. This page lists all Authorization profiles defined for your network, regardless of the scope you selected in the network view.

Table 88: Manage Authorization Profile Fields

Field	Description
Profile Name	Name given to the profile when the profile was created.
Family Type	The device family on which the profile was created.
VLAN Profile	The VLAN profile associated with the Authorization profile. You specify a VLAN profile while creating an Authorization profile.

Table 88: Manage Authorization Profile Fields (continued)

Field	Description
VLAN Pool	The VLAN pool associated with the Authorization profile. You can specify a VLAN pool while creating an Authorization profile.
CoS Profile	The optional CoS profile associated with the Authorization profile.
Description	Description of the profile that was entered when the profile was created. TIP: To display the entire description, you might need to resize the Description column by clicking the column border in the heading and dragging it.
Assignment State	Displays the assignment state of the profile. A profile can be: <ul style="list-style-type: none"> • Unassigned—When the profile is not assigned to any object. • Deployed—When the profile is assigned and is deployed from Deploy mode. • Pending Deployment—When the profile is assigned, but not yet deployed in the network. For more information, see “Deploying Configuration to Devices” on page 1180.
Creation Time	Date and time when the profile was created.
Last Updated Time	Date and time when the profile was last modified.
User Name	The username of the person who created or modified the profile.

TIP: All columns might not be displayed—this is configurable. To show or hide fields listed in the Manage Authorization Profiles table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

Creating a Wireless Authorization Profile

In Network Director, you can create a wireless Authorization profile with access permissions for either wireless users or devices. You can also link a VLAN profile and a CoS profile to the Authorization profile to ensure that each user session is assigned to an appropriate VLAN and it gets the required class of service (CoS).

For an Authorization profile, you must specify the following:

- A profile name

- An associated VLAN Profile or VLAN pool profile

To create an Authorization profile for wireless users or devices:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.

3. In the Tasks pane, expand **Wireless**, expand **Profiles**, and then click **Authorization**.

The Manage Authorization Profiles page appears, displaying the list of currently configured wireless Authorization profiles.

4. Click **Add**.

The Create Authorization Profile for Wireless page appears.

5. Specify the settings as described in both the online help and in [“Specifying Settings for a Wireless Authorization Profile” on page 399](#).

6. Click **Done** to save the Authorization profile.

The system saves the Authorization profile and displays the Manage Authorization Profiles page. Your new or modified Authorization profile is listed in the table of Authorization profiles.

You will need this authorization profile to create a WLAN Service profile—for directions, see [“Creating and Managing a WLAN Service Profile” on page 1090](#).

Specifying Settings for a Wireless Authorization Profile

While creating an Authorization profile, you will have to specify a VLAN profile. Make sure that you have created a wireless VLAN profile before you attempt to create an Authorization profile. For directions, see [“Creating and Managing VLAN Profiles” on page 502](#).

1. Enter the settings described in [Table 89](#) to create an Authorization profile. Required settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 89: Authorization Profile Basic Settings (WLC)

Field	Action
Name	<p>Type a unique name that identifies the profile.</p> <p>You can use up to 32 characters for profiles created for wireless devices. Profile name must not contain special characters or spaces. Note that profiles that are automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character.</p>
Description	Type a short description for the profile.
VLAN Profile or VLAN Pool	<p>You can assign a VLAN profile or a VLAN pool profile to the selected controller. Enable either VLAN Profile or VLAN Pool. For directions, see “Creating and Managing VLAN Profiles” on page 502 or “Creating and Managing VLAN Pools” on page 535.</p> <p>Click the corresponding Select button and then select a VLAN profile or VLAN pool to include in the Authorization profile. When a VLAN profile or pool is applied to a port or a wireless access point, it is initiated when clients are connected and are authorized on the VLAN.</p>
CoS Profile	<p>Click Select and then select an optional CoS profile to include in the Authorization profile.</p> <p>CoS profiles enable the grouping of class of service parameters and the application of it to one or more network sessions. You can configure policers, classifiers, scheduler maps, rewrite rules and a traffic-control profile within a CoS Profile. For directions, see “Creating and Managing Wired CoS Profiles” on page 613.</p>
mDNS Profile	Select an mDNS Profile from the list for Apple TV, Internet printer, or Digital Auto Access Protocol (iTunes). mDNS Profiles are created by following the directions in “Creating and Managing mDNS Profiles” on page 997 . For more information, see “Understanding Bonjour” on page 995 .

Filters

Filters are computer programs that process and sort a data stream. For more information, see [“Understanding Filter Profiles” on page 540](#).

Table 89: Authorization Profile Basic Settings (WLC) (continued)

Field	Action
Ingress Filter	Click Select and then select a Filter profile to filter traffic that enters the controller from users through an access port, from a wired authentication port, or from the network through a network port. For directions, see “Creating and Managing Wired Filter Profiles” on page 542.
Egress Filter	Click Select and then select a Filter profile to filter traffic sent from the controller to users through an access port, from a wired authentication port, or from the network through a network port. For directions, see “Creating and Managing Wired Filter Profiles” on page 542.
Simultaneous Login	Restrict the number of concurrent sessions that a user can have on the network by selecting the number of concurrent sessions for users of this Authorization profile.
Service Type	Select the type of access that you want the users of the Authorization profile to have: <ul style="list-style-type: none"> • 2 (Framed)—Select to grant network user access. • 6 (Administrative)—Select to grant administrative access to the controller with authorization to access enabled (configuration) mode. The user must enter the enable command and the correct enable password to access enabled mode. • 7 (NAS-Prompt)—Select to grant administrative access to non-enabled mode only. In this mode, a user cannot enter the enable command nor the enable password to access the enabled mode.

To configure advanced settings, click **Advanced Settings**. To skip changing the default advanced settings and save the profile, click **Done**.

2. Enter the advanced settings described in [Table 90](#) to modify the default advanced settings for the Authorization profile.

Table 90: Authorization Profile Advanced Settings (WLC)

Field	Action
User Idle Timeout (default is 3600 seconds)	Specify the length of time that a user or device can remain idle before the controller disconnects the user or device.
Session Timeout (default is 180 seconds)	Specify the length of time a user or device can remain connected to the network before re-authenticating the session.

Table 90: Authorization Profile Advanced Settings (WLC) (continued)

Field	Action
Termination Action	<p>Select the action to be taken when the session expires:</p> <ul style="list-style-type: none"> • 0 (Disconnect)—Select to indicate that the session is to be terminated. • 1 (Re-authenticate)—Select to indicate that the user or device must reauthenticate when the session expires.
Uniform Resource Locator (URL)	<p>Specify the URL that the user is to be redirected after successful authentication.</p> <p>Use the following format: <i>http://www.example.com</i></p>

Accounting Interim Interval

Enable Updates	<p>Select Enable Updates to enable accounting updates for the Authorization profile.</p> <p>TIP: Accounting updates are applicable only if you have enabled accounting and selected START-STOP as the record type in the corresponding Authentication profile.</p> <p>Update Interval: If updates are enabled, you can modify the time in seconds between accounting updates.</p> <p>Specify a value from 180 (default) through 3600 seconds.</p> <p>NOTE: If both a RADIUS server and a controller supply a value for the Accounting Interim Interval, then the value from the controller takes precedence.</p>
-----------------------	--

Encryption Type

Encryption is the process of encoding messages or information in such a way that only authorized parties can read it. For more information, see [“Understanding Wireless Encryption and Ciphers” on page 899](#).

Table 90: Authorization Profile Advanced Settings (WLC) (continued)

Field	Action
Encryption Type	<p>Select the type of encryption supported for clients that use this Authorization profile. You can select a combination of encryption types. Clients who attempt to use an unauthorized encryption method are rejected. Network Director supports the following encryption types:</p> <ul style="list-style-type: none"> • AES-CCMP—AES-CCMP (Advanced Encryption Standard using Counter with CBC-MAC) is the standard encryption protocol for use with the WPA2 standard and is much more secure than the TKIP protocol. • TKIP—TKIP (Temporal Key Integrity Protocol) is a security protocol used in the IEEE 802.11 wireless networking standard. It uses the same underlying mechanism as the original WEP encryption, and consequently is vulnerable to the same attacks that WEP is vulnerable to. • None—No encryption is used.
Start and End Dates for Authorization	
Start Date	Select the date and 24-hour time from which users of this authorization profile are authorized to access the network.
End Date	Select the last date and 24-hour time that users of this authorization profile are authorized to access the network.
Time of Day Settings	
Time of Day	Indicate the time of day the user is permitted to log in to the network. The default is Any and the other options are Never and Day .

3. Click **OK** to save the advanced settings and close the Advanced Settings window.
4. Click **Done** to save the authorization profile and add it to the Manage Authorization Profiles list.

What To Do Next

After you have created an Authorization profile, you can:

- Associate it with a WLAN profile. For directions, see [“Creating and Managing a WLAN Service Profile” on page 1090](#).
- Assign it to one or more controllers. For directions, see [“Assigning Wireless Authorization Profiles to Controllers” on page 404](#).

RELATED DOCUMENTATION

[Understanding Wireless Authorization Profiles | 395](#)

[Assigning Wireless Authorization Profiles to Controllers | 404](#)

[Creating and Managing a WLAN Service Profile | 1090](#)

[Network Director Documentation home page](#)

Assigning Wireless Authorization Profiles to Controllers

IN THIS SECTION


- [Assigning Authorization Profiles | 404](#)
- [Editing Authorization Profile Assignments | 406](#)

After creating an Authorization profile, assign it to a wireless device (controller), a cluster of controllers, or the members of a group of controllers. A cluster is a logical grouping of wireless LAN controllers that ensures load balancing and seamless mobility within the wireless domain. Configure a mobility domain to be a cluster from the Create Mobility Domain page. For directions, see [“Creating a Mobility Domain for Wireless LAN Controllers” on page 1053](#). For directions to create device groups, see [“Creating Custom Device Groups” on page 276](#).

The following sections describe how to assign an Authorization profile and edit existing Authorization profile assignments.

Assigning Authorization Profiles

To assign an Authorization profile to controllers:

1. Click  in the Network Director banner.
2. Under Select View, select either **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Datacenter View** or **Topology View**.

3. Click **Authorization** under Profile and Configuration Management in the Tasks pane.

The Manage Authorization Profiles page appears.

4. Select an Authorization profile that you want to assign and click **Assign**.

The Assign Authorization Profile page appears displaying a hierarchal list of network objects, including the network, mobility domain, controllers and clusters that are already defined or discovered for your network.

5. Select a level and click **Next** to view the objects available at that level. If you select a network or a mobility domain, all the controllers or clusters that are part of that network or mobility domain are also selected.
6. Select one or more devices, groups, or clusters from the list.

NOTE: If Network Director fails to read the configuration of one or more devices after the device discovery, such devices are not displayed in the Assign Profile page. You will not be able to assign profiles to such devices. The Manage Jobs page in System mode displays details of the device discovery jobs. Use the information displayed on this page to take appropriate corrective steps to enable Network Director to reread the configuration of the failed device. For more information, see [“Discovering Devices in a Physical Network” on page 204](#).

7. If you want to view the existing assignment of a device or a cluster, select it and click **View Assignments**.

The Profile Details window opens displaying the device's current profile assignment.

Click **Close** to close the window.

8. If you want to remove an existing assignment from a device or a cluster, select it and click **Remove**. The system removes the assignment from the selected device or cluster.

9. Do one of the following to assign the Authorization profile to a device or a cluster:

- Click **Assign > Assign to Device** to assign the Authorization profile to the selected devices.
- Click **Assign > Assign to Cluster** to assign the Authorization profile to the selected clusters.

10. Click **Next** or **Review**.

The system displays the associations that you created. To modify any of these assignments, click **Edit** or **Profile Association**.

11. Click **Finish** to save the profile associations.

After you click Finish, the Create Profile Assignments Job Details window opens with a report on the status of the profile assignment job. If you have assigned the profile to a large number of objects, the


profile assignment job can take some time to complete. Instead of waiting for the Job Details window to report job completion status, you can close it and check the details of the profile assignment job at a later time using the Manage Job task in System mode.

NOTE: If any assignment fails, the profile assignment job fails and none of the assignments are created. Check the details for the profile assignment job for information about why the assignment failed.

Editing Authorization Profile Assignments

You can modify the Authorization profile assignments that you made to a device or a cluster using the Edit Assignments page.

To edit a profile assignment:

1. Click  in the Network Director banner.
2. Under Select View, select either **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Datacenter View** or **Topology View**.

3. Click **Authorization** under Profile and Configuration Management in the Tasks pane.
The Manage Authorization Profiles page appears.
4. Select the Authorization profile for which you want to edit assignments and click **Edit Assignment**.
The device categories to which the selected Authorization profile is assigned are displayed. Expand the device category to view the device details and the association status.
5. To delete an existing assignment from a device or a cluster, select it and then click **Delete**.
The Record Status field is updated for status of the record. A cross mark indicates a deleted record or assignment.
6. Click **Apply** to save the changes and close the page.

RELATED DOCUMENTATION

[Understanding Wireless Authorization Profiles | 395](#)

[Creating and Managing Wireless Authorization Profiles | 395](#)

[Network Director Documentation home page](#)

Configuring Interfaces and VLANs

IN THIS CHAPTER

- Understanding Port Profiles | 408
- Creating and Managing Port Profiles | 414
- Assigning and Unassigning Port Profiles from Interfaces | 476
- Managing Auto Assignment Policies | 484
- Creating Auto Assignments | 486
- Configuring Easy Config Setup | 489
- Understanding Port Groups | 495
- Creating and Managing Port Groups | 495
- Understanding VLAN Profiles | 499
- Creating and Managing VLAN Profiles | 502
- Assigning a VLAN Profile to Devices or Ports | 531
- Creating and Managing VLAN Pools | 535

Understanding Port Profiles

IN THIS SECTION

- Interface Settings Configured in the Port Profile | 409
- Interface Settings Configured by Referencing Other Profiles | 410
- Data Center Device Port Profile Settings | 410
- Default Port Profiles | 410

Port profiles provide a convenient way of provisioning interfaces on switches. You can either use predefined port profiles, or you can define your own custom port profile.

After you create a Port profile, you can assign it to interfaces on one or more switches, including aggregated interfaces. For the configuration created by the profile to take effect on the devices, you must use Deploy mode to deploy the configuration on the devices.

This topic describes:

Interface Settings Configured in the Port Profile

You can configure the following interfaces settings in a Port profile:

- Interface protocol family—You can configure an interface to be either an Ethernet switching interface, an IPv4 routing interface, or an IPv6 routing interface.
- Port mode—You can configure a switching interface port mode to be an access, trunk, or tagged-access interface for EX Series switches. Campus Switching ELS supports access mode and trunk mode. For more information about port modes, see [Ethernet Switching](#).
- PoE settings—The factory-default configuration of switches enables PoE on all interfaces that support PoE. For many implementations, no further configuration is necessary. You can, however, override the default settings for PoE interfaces in the Port profile. Most switch models have interfaces that support Power over Ethernet (PoE), but the EX9200 does not support PoE. For more information about PoE, see [Power over Ethernet \(PoE\)](#).

If you do not explicitly configure PoE in the Port profile, the existing PoE interface settings on the switch remain in effect. Device-wide PoE settings are configured in the Device Common Settings profile.

NOTE: PoE settings are not available for Data Center devices.

- Physical link settings—On switches, the autonegotiation of port speed and duplex mode is enabled by default. You can disable autonegotiation and set port speed and duplex mode in the Port profile. Other link settings you can configure include flow control, which is disabled by default, and maximum transmission unit (MTU).
- Storm control settings—You can optionally enable storm control settings on switches. Storm control monitors traffic levels drops broadcast, multicast, and unknown unicast packets when a specified traffic level—called the storm control level.
- RSTP settings—You can optionally enable RSTP settings on switches. RSTP this feature enables you to fine-tune STP by setting interfaces into edge, disable, or no-root-port states.
- Port security settings—You can optionally enable port security on switched access ports. Port security features help protect the access ports on your switch against address spoofing (forging) and Layer 2 denial-of-service (DoS) attacks. For more information about port security on switches, see [Port Security](#).

NOTE: For campus switching ELS devices, disabling port security settings option is not available.

Interface Settings Configured by Referencing Other Profiles

You can optionally configure other interface-related settings in the Port profile by referencing other profiles. These profiles are:

- CoS profile—Configures class-of-service settings on the interface. You can either select or create an in-line CoS profile while creating a port profile.
- Filter profile—Configures firewall filters (often called ACLs) on the interface.
- Authentication profile—(Switching interfaces only) Configures 802.1X authentication on an interface and configures related settings, such as captive portal authentication. You can either select or create an in-line authentication profile while creating a port profile.
- Access profile—Configures the access server settings used by all 802.1X authenticator interfaces on a switch. This profile is not included in the Port profile. Instead, you assign it to a device as part of the process of assigning a Port profile to the interfaces on the device.
- VLAN profile for Campus Switching ELS is mandatory. You can either select or create an in-line VLAN profile while creating a port profile.

If you want to use one or more of these profiles with the Port profile, be sure to create them before you create and assign the Port profile.

Data Center Device Port Profile Settings

Port profiles for data center switching devices include settings for the following features:

- Configuring Fibre Channel (FC) ports.
- Configuring Data Center Bridging Capability Exchange (DCBX) protocol for Ethernet interfaces.

Default Port Profiles

To help with the rapid provisioning of interfaces on switches, Network Director provides default Port profiles that contain settings for common uses of switch interfaces. You can modify or assign these default profiles to interfaces using the same method used for user-created profiles. [Table 91](#) describes the default Port profiles.

Table 91: Default Port Profiles

Profile Name	Description	Summary of Settings
Desktop_Port	Configures an untagged port that connects to desktop computer.	<ul style="list-style-type: none"> • Family Type—switching • Port Mode—access • Auto Negotiation—disabled • Flow Control—disabled • Maximum Bytes—disabled • Speed—no default provided • Link Mode—no default provided • Trust DHCP—no • MAC Limit—1 • MAC Limit Action—drop • CoS Profile—no default provided
Desktop_Phone_Port	Configures an untagged port that connects to a combined desktop and phone port.	<ul style="list-style-type: none"> • Family Type—switching • Port Mode—access • Auto Negotiation—disabled • Flow Control—disabled • Maximum Bytes—disabled • Speed—no default provided • Link Mode—no default provided • Trust DHCP—no • MAC Limit—2 • MAC Limit Action—drop • CoS Profile—juniper_CoS_template
Fibre_Channel_Port	Configures a Fibre Channel (FC) port. Available for data center switching profiles only.	<ul style="list-style-type: none"> • Port Type—Fibre Channel Port • Speed—4Gbps • Buffer to Buffer State Change Number—no default provided • Loopback Setting—no default provided

Table 91: Default Port Profiles (continued)

Profile Name	Description	Summary of Settings
FCoE_Transit_Port	Configures an Ethernet port for an FCoE transit switch. Available for data center switching profiles only.	<ul style="list-style-type: none"> ● Port Type—Ethernet Port ● CoS Profile—juniper_DC_Hier_CoS ● Family Type—switching ● Port Mode—trunk ● Filters—no default provided ● VLAN Options—no default provided ● DCBX Version—Auto ● Disable DCBX—disabled ● Disable Priority Flow Control—disabled ● ETS No Auto Negotiation—disabled ● Recommendation TVL—no default provided ● Auto Negotiation—disabled ● Flow Control—disabled ● Maximum Size—2500 ● Speed—no default provided ● Link Mode—no default provided ● Port Security—enabled ● Trust DHCP—disabled ● FCoE Trusted—enabled ● MAC Limit—no default provided ● MAC Limit Action—no default provided ● Allowed MAC List—no default provided
Server_Port	Configures a tagged port that connects to a server.	<ul style="list-style-type: none"> ● References the default CoS profile, juniper_CoS_template ● Sets protocol family to Ethernet switching ● Sets port mode to trunk ● Enables port security with trust DHCP enabled

Table 91: Default Port Profiles (continued)

Profile Name	Description	Summary of Settings
Switched_Downlink	Configures a tagged port that connects to endpoint devices in a branch environment or servers in a data center environment.	<ul style="list-style-type: none"> ● Family Type—switching ● Port Mode—trunk ● Auto Negotiation—disabled ● Flow Control—disabled ● Maximum Bytes—no default provided ● Speed—no default provided ● Link Mode—no default provided ● Trust DHCP—yes ● MAC Limit—no default provided ● MAC Limit Action—no default provided ● CoS Profile—juniper_CoS_template
Switched_Uplink	Configures a tagged port that connects a switch to another switch or larger network. For example, a port that connects an access switch to an aggregation switch.	<ul style="list-style-type: none"> ● Family Type—switching ● Port Mode—trunk ● Auto Negotiation—disabled ● Flow Control—disabled ● Maximum Bytes—no default provided ● Speed—no default provided ● Link Mode—no default provided ● Trust DHCP—yes ● MAC Limit—no default provided ● MAC Limit Action—no default provided ● CoS Profile—juniper_CoS_template
Wireless_Access_Port	Configures an untagged port that connects to a wireless access point.	<ul style="list-style-type: none"> ● Family Type—switching ● Port Mode—access ● Auto Negotiation—disabled ● Flow Control—disabled ● Maximum Bytes—no default provided ● Speed—no default provided ● Link Mode—no default provided ● Trust DHCP—no default provided ● MAC Limit—no default provided ● MAC Limit Action—no default provided ● CoS Profile—no default provided

RELATED DOCUMENTATION

[Creating and Managing Port Profiles | 414](#)

[Understanding Access Profiles | 351](#)

[Understanding Authentication Profiles | 381](#)

[Understanding Class of Service \(CoS\) Profiles | 609](#)

[Understanding Filter Profiles | 540](#)

[Network Director Documentation home page](#)

Creating and Managing Port Profiles

IN THIS SECTION

- [Managing Port Profiles | 415](#)
- [Creating Port Profiles | 418](#)
- [Specifying Settings for an EX Switching Port Profile | 418](#)
- [Specifying Settings for a Campus Switching ELS Port Profile | 432](#)
- [Specifying Settings for the Data Center Switching Non-ELS Port Profile | 446](#)
- [Specifying Settings for a Data Center Switching ELS Port Profile | 460](#)
- [What to Do Next | 475](#)

Port profiles provide a way to provision multiple switch interfaces, including Ethernet interfaces on EX Series switches, Campus Switching ELS, Data Center Switching devices, and Fibre Channel (FC) interfaces on Data Center Switching devices. In a Port profile, you can define a set of attributes to be shared by multiple interfaces. For example, you can create a Port profile for all access interfaces that connect to VoIP desk phones, configuring the appropriate class-of-service (CoS), authentication, and port security settings for these interfaces in the Port profile. You then assign the Port profile to those interfaces and deploy the resulting configuration on the interfaces.

Port profiles define only shared attributes. To enable you to configure specific attributes for an interface or a switch during the process of assigning a Port profile to an interface, the Create Port profile wizard provides two setup options: Quick Setup and Custom Setup. The Quick Setup option enables you to create initial configuration settings for a Port profile including selecting or create inline VLAN profile. The Custom Setup option enables you to configure all the advanced settings and create any inline sub-profiles. In Custom Setup option, apart from selecting the existing VLAN, CoS, and authentication sub-profiles, you can also create these sub-profiles.



CAUTION: Ports that are involved in EVPN-VXLAN is not configured through port profile. Else, it renders EVPN-VXLAN defunct.

NOTE:

- If you switch from Quick Setup to Custom Setup, all the configuration settings are saved. However, if you switch from Custom Setup to Quick Setup, all the advanced settings done in the Custom Setup are lost.

To manage or create Port profiles: In Build mode, select **Port** from Profiles in the Tasks pane. The Manage Port Profile page appears.

This topic describes:

Managing Port Profiles

Use the Manage Port Profiles page to manage existing Port profiles and to create new ones. Port profiles enable the definition and application of a common set of attributes to interfaces.

From the Manage Port Profiles page, you can:

- Create a new profile by clicking **Add**. For details, see [“Creating Port Profiles” on page 418](#).
- Modify an existing profile by selecting it and clicking **Edit**.
- Associate a Port profile to specific interfaces by selecting it and clicking **Assign**.

During the assignment process, you can choose to configure interface-specific settings, such as IP address.

- Change a Port profile's current interface assignments by selecting it and clicking **Edit Assignments**. This opens the Edit assignments for profile-name page, which displays the assignment state and other details of the interfaces in a grid layout. After editing an assignment, and click **Apply**. The Edit Profile Assignment Job Details window opens, which reports the status of the interface assignment that you edited.
- View information about a profile, including the interfaces it is associated with, by selecting the profile and clicking **Details** or by clicking the profile name, which opens the Profiles Details page. This page displays the profile details and the interface associations in a grid layout. It also has an option using which you can search profiles associated with a device and filter devices. Click **Show Filters** to filter an interface based on its IP address, serial number, type, or location or custom group.
- Perform the search for the following:
 - A Port profile for a specific device by specifying the device details in the search field.

- A port profile that is assigned to a specific port on a device. In this case, you must first enter the device details and then specify the port details in the search field to view the port profile.
- Port profiles that are assigned to interfaces that are part of the same VLAN. When you specify the VLAN name in the search field, all the Port profiles that are part of the same VLAN are listed in the table.
- Delete profiles by selecting the profiles and clicking **Delete**.

TIP: You cannot delete profiles that are in use—that is, assigned to objects or used by other profiles. To see the current assignments for a profile, click the profile name.

- Clone a profile by selecting a profile and clicking **Clone**.

Network Director provides a set of default Port profiles: Desktop Port, Desktop and Phone Port, Server Port, Switched Downlink, Switched Uplink, Wireless Access Port, and Custom Port. These profiles contain configuration appropriate for the named port role. You can manage these profiles the same way that you manage a user-created profile. For more information about these profiles, see [“Understanding Port Profiles” on page 408](#).

[Table 92](#) describes the information provided about Port profiles on the Manage Port Profiles page. This page lists all Port profiles defined for your network, regardless of your current selected scope in the network view.

Table 92: Manage Port Profiles Table


Column	Description
Profile Name	<p>Name given to the profile when the profile was created.</p> <p>Click the profile name to view profile details.</p> <p>A  next to the profile name indicates that the profile is assigned to a port using an auto assignment policy. For more details on auto assignment policies, see “Managing Auto Assignment Policies” on page 484.</p>
Family Type	<p>One of the following:</p> <ul style="list-style-type: none"> • EX—for EX Series switches • ELS—for Campus Switching ELS • Data Center Switching—for Data Center Switching devices • Data Center Switching ELS—for Data Center Switching ELS devices
Description	Description of the Port profile that was entered when the profile was created.


Table 92: Manage Port Profiles Table (*continued*)

Column	Description
Port Family	One of the following: <ul style="list-style-type: none"> • Switching—for Port profiles that configure Layer 2 interfaces • Routing—for Port profiles that configure Layer 3 interfaces • FIBRE—for Port profiles that configure Fibre Channel (FC) interfaces.
VLANs	Name of the member VLANs configured or referenced for that Port profile.
Assignment State	One of the following states: <ul style="list-style-type: none"> • Deployed—The profile has been assigned to interfaces and the configuration has been deployed on the devices. • Pending Deployment—The profile has been assigned to interfaces or its previous assignments have been changed, but the new or modified configuration has not yet been deployed on the devices. • Unassigned—The profile has not yet been assigned to interfaces.
Creation Time	Date and time when the profile was created.
Last Updated Time	Date and time when the profile was last modified.
Assigned to Devices	Number of devices to which the Port profile is assigned. Click on the link to view the profile details.
Assigned to Port	Number of ports to which the Port profile is assigned. Click on the link to view the profile details.
Assigned to	Number of port assignments and device associations for a profile.
User Name	The username of the user who created or modified the profile.

TIP: All columns might not be currently displayed. To show or hide fields in the table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

Creating Port Profiles

To create a Port profile for EX Series switches, Campus Switching ELS, Data Center Switching, or Data Center Switching ELS:

1. Click  in the Network Director banner.
2. Under Views, select one of the following views: **Logical View**, **Location View**, **Device View** or **Custom Group**.

TIP: Do not select **Datacenter View** or **Topology View**.

3. Click **Port** under Wired > Profiles in the Tasks pane.
The Manage Port Profile page appears.
4. Click **Add**.
The Select a Device Family page opens.
5. Select **Switching (EX)**, **Campus Switching ELS**, **Data Center Switching ELS**, or **Data Center Switching Non-ELS**.

The Create Port Profile page appears showing the Quick Setup and Custom Setup tabs for the selected family with the appropriate fields for configuring that family.

6. Select initial settings in the Quick Setup option and advanced settings in the Custom Setup option for the Port profile. For information about the Port profile settings, select the section for the type of port you are configuring:
 - [Specifying Settings for an EX Switching Port Profile on page 418](#)
 - [Specifying Settings for a Campus Switching ELS Port Profile on page 432](#)
 - [Specifying Settings for the Data Center Switching Non-ELS Port Profile on page 446](#)
 - [Specifying Settings for a Data Center Switching ELS Port Profile on page 460](#)

Specifying Settings for an EX Switching Port Profile

Use the Create Port Profile page to define a common set of port attributes, which you can then apply to a group of interfaces. These directions address creating a Port profile for EX Series switches.

TIP: You can reference a VLAN profile, CoS profile, Ingress Filter profile, Egress Filter profile, and an Authentication profile in a Port profile. You can either create these profiles in their respective profile pages before you create Port profiles or you can create these profiles as in-line sub-profiles while configuring Port profiles. You can also enable power over Ethernet (PoE).

After you create a Port profile, you assign it to individual interfaces or to members of a port group. During this process, you can also configure interface-specific attributes, such as IP address, and certain device-specific attributes, such as the Access profile to use for all ports on the device. You can assign only one Port profile to an interface.

[Table 93](#) describes the Quick Setup settings available in a Port profile. [Table 94](#) describes the Custom Setup settings. The defaults for these options depend on the Service Type you select.

Table 93: Port Profile Quick Setup Settings for an EX Switching Port Profile

Field	Action
Profile Name	A default name that corresponds to the Service Type is displayed—when you change the Service Type, this default profile name changes. You can also change the name of profile, using up to 64 alphanumeric characters and no special characters other than the underscore. The name must be unique among Port profiles.
Description	A default description of the preconfigured service types appears by default. You can change the description of the Port profile, which appears on the Manage Port Profiles page. You can use up to 256 characters.

Table 93: Port Profile Quick Setup Settings for an EX Switching Port Profile (continued)

Field	Action
Service Type	<p>Select one the preconfigured switching options, Desktop Port, Desktop Phone Port, Printer Port, Switched Uplink, Switched Downlink, Server Port, or Wireless Access Port. To create your own switching or routing service type, select Custom.</p> <p>TIP: No preconfigured routing Service Types are provided. You must create them using the Custom option.</p> <hr/> <p>Desktop Port default service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—no default provided • Family Type—switching • Port Mode—access • Power over Ethernet—disabled • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—disabled • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • Trust DHCP—disabled • MAC Limit—1 • MAC Limit Action—drop • Allowed MAC List—no default provided

Table 93: Port Profile Quick Setup Settings for an EX Switching Port Profile (*continued*)

Field	Action
	<p>Desktop Phone Port preconfigured service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—juniper_CoS_template • Family Type—switching • Port Mode—access • Power over Ethernet—disabled • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—disabled • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • Trust DHCP—disabled • MAC Limit—2 • MAC Limit Action—drop • Allowed MAC List—no default provided
	<p>Printer Port preconfigured service type has the following default settings:</p> <ul style="list-style-type: none"> • Family Type—switching • Port Mode—access • Power over Ethernet—no default provided • Auto Negotiation—enabled • Flow Control—enabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • Port Security—no default provided • Trust DHCP—no default provided • MAC Limit—no default provided • MAC Limit Action—no default provided • Allowed MAC List—no default provided

Table 93: Port Profile Quick Setup Settings for an EX Switching Port Profile (*continued*)

Field	Action
	<p>Switched Uplink preconfigured service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—juniper_CoS_template • Family Type—switching • Port Mode—trunk • Power over Ethernet—disabled • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • MAC Limit—no default provided • Trust DHCP—disabled • MAC Limit Action—no default provided • Allowed MAC List—no default provided
	<p>Switched Downlink preconfigured service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—juniper_CoS_template • Family Type—switching • Port Mode—trunk • Power over Ethernet—disabled • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • MAC Limit—no default provided • Port Security—enabled • Trust DHCP—enabled • MAC Limit Action—no default provided • Allowed MAC list—no default provided

Table 93: Port Profile Quick Setup Settings for an EX Switching Port Profile (continued)

Field	Action
	<p>Server Port preconfigured service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—juniper_CoS_template • Family Type—switching • Port Mode—trunk • Power over Ethernet—disabled • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • MAC Limit—no default provided • MAC Limit Action—no default provided • Allowed MAC list—no default provided
	<p>Wireless Access Port preconfigured service type has the following default settings:</p> <ul style="list-style-type: none"> • Family Type—switching • Port Mode—Access • Power over Ethernet—disabled • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • Trust DHCP—disabled • MAC Limit—no default provided • MAC Limit Action—no default provided • Allowed MAC List—no default provided

Table 93: Port Profile Quick Setup Settings for an EX Switching Port Profile (continued)

Field	Action
Family Type	<p>This setting cannot be changed if any preconfigured Service Type was selected. If you selected the Custom Service Type, indicate whether the interface operates as a Layer 2 (Switching) or a Layer 3 (Routing) interface.</p> <p>TIP: All preconfigured Service Types are for switching.</p> <p>If you select Routing, you configure an IP address on a per-interface basis when you assign the profile to individual interfaces.</p> <p>TIP: Service Type must be set to Custom to configure a routing interface.</p>
Port Mode	<p>This setting cannot be changed if any preconfigured Service Type was selected. If you selected the Custom Service Type, select the port mode for the EX Series switching interface, either Access, Trunk, or Tagged Access.</p> <ul style="list-style-type: none"> ● Access—Use for interfaces that connect to an end device, such as a desktop computer, an IP telephone, a wireless access point, a printer, or a security camera. The interface must belong to a single VLAN. Frames sent and received over the over the interface are untagged Ethernet frames. This is the default for a Desktop Port, Desktop Phone Port, and Wireless Access Port. ● Trunk—Use for interfaces that connect to a switch or router. Trunk interfaces can belong to more than one VLAN, enabling VLAN traffic to be multiplexed on a single physical interface. The Ethernet frames sent and received over the interface are tagged frames, in which IEEE 802.1Q tagging is used to segregate the traffic from each VLAN. This is the default for Switched Uplink, Switched Downlink, Server Port, and Wireless Access Port. ● Tagged Access—Use for access interfaces where VLAN tagging is required, typically when the interface connects to a server running virtual machines using virtual Ethernet port aggregator (VEPA) technology. The traffic generated by the server can contain an aggregation of VLAN packets from different virtual machines on that server, requiring that packets be tagged.
<p>VLAN Options</p> <p>Available VLAN options depend on the Service Type selected.</p>	
Member VLAN (available for Switched Uplink, Switched Downlink, Server Port)	<p>Click All if you want to assign an interface to all the VLANs.</p> <p>This option is enabled when Port Mode is Trunk or TaggedAccess.</p>

Table 93: Port Profile Quick Setup Settings for an EX Switching Port Profile (continued)

Field	Action
Member VLANs (available for Desktop Port, Desktop Phone Port, Switched Uplink, Switched Downlink, Server Port, Wireless Access Port, Custom Port)	Select a VLAN for the interface by clicking Select , selecting one of the listed filters, and then clicking OK . The VLAN is added to the Member VLANs list. Or Configure a VLAN by clicking Configure VLAN Settings and clicking Create . Enter the VLAN name and ID and click OK .
Voice VLAN (available for Desktop Phone Port, Custom Port)	Select a voice VLAN for the interface by clicking Select , selecting one of the listed filters, and then clicking OK . The VLAN is added to the Member VLANs list. Or Configure a VLAN by clicking Configure VLAN Settings and clicking Create . Enter the VLAN Name and ID and click OK .
Native VLAN (available for Switched Uplink, Switched Downlink)	Select a native VLAN for the interface by clicking Select , selecting one of the listed VLANs, and then clicking OK . The VLAN is added to the Member VLANs list. Or Configure a VLAN by clicking Configure VLAN Settings and clicking Create . Enter the VLAN name and ID and click OK .

After providing the information in the fields listed in the preceding, click **Done**.

To use default Port Profile Custom Setup settings, click **Done**. To configure Custom Setup settings, click **Custom Setup** and then provide the information in [Table 94](#) and then click **Done**.

Clicking **Done** in either case displays the dialog Do you want to assign Port Profile to Ports. Click **Yes** to create a profile assignment; else click **No** to navigate to the Manage Port Profile page and to create the Port assignment later on.

Table 94: Port Profile Custom Setup Settings

Field	Action
-------	--------

Advanced Settings

Expand Advanced Settings to configure link settings and port security. The Link Setting in Port profile is disabled by default. On enabling Link Settings, autonegotiation and flow control are enabled by default.

Table 94: Port Profile Custom Setup Settings (continued)

Field	Action
Enable Auto Negotiation	<p>Autonegotiation of link speed and duplex mode is enabled by default; clear to disable autonegotiation.</p> <p>If you disable autonegotiation, you must set link speed and link mode.</p> <p>You cannot disable autonegotiation if a link speed of 1 Gbps is configured. This configuration might be accepted, but autonegotiation is not disabled.</p> <p>NOTE: This setting is ignored when you assign a Port profile to an Aggregated Ethernet interface.</p>
Enable Flow Control	<p>Select to enable flow control on the interface, which permits the switch suspend packet transmission for a set period of time in response to a PAUSE frame sent by a congested switch.</p> <p>Flow control applies only to links operating at 1 Gbps, full-duplex mode.</p>
MTU	<p>Using the arrows, indicate the maximum transmission unit (MTU), which is the maximum size of Ethernet frames sent by the interface. To calculate the MTU, add 14 bytes overhead to the maximum payload you want sent.</p> <p>Range: 256 through 9216 bytes</p>
Speed	<p>Select the link speed.</p> <p>If you select a link speed when autonegotiation is enabled, autonegotiation remains enabled and the interface advertises the link speed that you specify as its maximum link speed.</p> <p>NOTE: This setting is ignored when you assign a Port profile to an Aggregated Ethernet interface.</p>
Link Mode	<p>Select the duplex mode, either Automatic, Full Duplex, or Half Duplex. Select Automatic to enable autonegotiation when autonegotiation is disabled.</p> <p>NOTE: This setting is ignored when you assign a Port profile to an Aggregated Ethernet interface.</p> <p>You cannot select Half Duplex with link speed set to Autonegotiation or 1 Gbps.</p>

Table 94: Port Profile Custom Setup Settings (continued)

Field	Action
<p>Storm Control Settings</p> <p>Enabling storm control on a switching device monitors traffic levels and drops broadcast, multicast, and unknown unicast packets when a specified traffic level—called the storm control level or storm control bandwidth—is exceeded, thus preventing packets from proliferating and degrading the LAN.</p> <p>You can customize the storm control level for a specific interface by explicitly configuring either bandwidth or level.</p> <p>NOTE: You cannot configure both bandwidth and level for the same interface.</p>	<p>Unit</p> <ul style="list-style-type: none"> <p>Percentage—Configures the storm control level as a percentage of the available bandwidth used by the combined applicable traffic streams that are subject to storm control on that interface.</p> <p>The level can be set from 0% to 100%, where 0% indicates that the entire traffic is being suppressed and 100% indicates no traffic is being suppressed, in other words there is no storm control.</p> <p>The default level is 80%.</p> <p>Kbps—Configures the storm control level as the bandwidth in kilobits per second (Kbps) of the applicable traffic streams on that interface.</p> <p>Set the bandwidth from 100 through 10,000,000 in Kbps. When you configure storm control bandwidth, the value you configure is rounded off internally to the closest multiple of 64 Kbps, and the rounded-off value represents the bandwidth that is actually used. For example, if you configure a bandwidth limit of 150 Kbps, storm control uses a bandwidth limit of 128 Kbps.</p> <p>Value</p> <p>Configures the traffic storm control threshold level value as a percentage of bandwidth or bandwidth in kilobits per second depending upon the specified unit.</p>
<p>Power over Ethernet (PoE)</p>	
<p>You can enable PoE and display the configuration options by enabling Configure Power over Ethernet.</p>	
<p>Configure Power over Ethernet</p>	<p>Enable to configure PoE settings.</p> <p>If you do not enable this option, Network Director does not send any PoE configuration commands to the device when the profile is deployed on the device. For example, if PoE is enabled on an interface, it remains enabled. On EX Series switches, the factory-default configuration enables PoE on all interfaces that support PoE.</p> <p>If you enable this option, the PoE settings in this profile is deployed on the interfaces that support PoE. If you assign this Port profile to interfaces that do not support PoE, the profile is deployed successfully on those interfaces, but the PoE settings do not take effect.</p>

Table 94: Port Profile Custom Setup Settings (continued)

Field	Action
Maximum Power (W)	<p>Use the arrows to adjust the maximum PoE power in watts allocated to a PoE port. The powered device cannot draw more power than the wattage specified. If it does so, PoE power to the port is shut down. Maximum power for PoE is 15.4W, Extended PoE is 18.6W and PoE+ is 30W.</p> <p>The Maximum Power setting has no effect when the PoE management mode for a switch or line card is class mode, which is the default mode. In class mode, the power allocated to a PoE port is determined either by LLDP negotiation with the powered device or by the PoE class of the powered device if LLDP is not supported.</p> <p>You must set the PoE power management mode for the switch or line card to static mode for the Maximum Power setting to take effect. Do this in the Device Common Settings profile.</p> <p>If you specify a maximum wattage that is greater than the maximum wattage that can be supplied by the port, your configuration is accepted when the Port profile is deployed on the port. However, the maximum wattage is set to the port's maximum supported wattage. The maximum supported wattages for PoE ports are:</p> <ul style="list-style-type: none"> • 15.4W for ports that support IEEE 802.3af only • 18.6W for IEEE 802.3af ports on switches that support enhanced PoE • 30W for ports that support IEEE 802.3at <p>Default: 15.4W</p>
Priority	<p>Select a power priority for the PoE port—either Low or High. If there is a shortage of PoE power on the switch, power to low priority ports is shut down before power to high priority ports. Within ports with the same assigned priority, power priority is determined by the port number—ports with a lower port number have a higher power priority.</p> <p>Default: low priority</p>
Disable PoE	<p>Select to disable PoE on the interfaces using this Port profile.</p>

Port Security (Switching Interfaces Only)

Select to enable port security (default); clear to disable port security.

When port security is enabled, you can configure port security options such as learned MAC address limits on an interface. When port security is disabled, no port security is applied to the interface, including the default port security options.

Table 94: Port Profile Custom Setup Settings (*continued*)

Field	Action
Trust DHCP	<p>Select to permit messages from a DHCP server to be received on the interface—this is the default. Clear to block all messages from a DHCP server from being received on the interface.</p> <p>TIP: For this port security feature to work, DHCP snooping must be enabled on the VLAN the interface belongs to. You can enable DHCP snooping on the VLAN in the VLAN profile. For directions, see “Creating and Managing VLAN Profiles” on page 502.</p>
MAC Limit	<p>Type the number of MAC address that can be dynamically learned on the interface.</p> <p>Range: 1 through 163,839</p> <p>Default: For Desktop Ports, 1. For Desktop Phone Ports, 2. For all others, none.</p>
MAC Limit Action	<p>Select the action to be taken if the MAC address limit is exceeded:</p> <ul style="list-style-type: none"> ● Drop—Drop any packet with a previously unlearned MAC address and generate a system log entry, and SNMP trap, or an alarm. This is the default for a Desktop Port and Desktop Phone Ports. ● Log—Accept packets with new MAC addresses and learn the addresses, but generate a system log entry, and SNMP trap, or an alarm. ● Shutdown—Shut down the interface and generate a system log message, SNMP trap, or an alarm. <p>If an interface is shut down because the MAC address limit has been exceeded, you must use the CLI command clear ethernet-switching port-error interface <i>name</i> to clear the error and bring the interface back into service.</p> <p>TIP: You can use the CLI to configure autorecovery on an interface that has been shut down by a MAC limit error.</p> <ul style="list-style-type: none"> ● None—No action. This selection effectively disables MAC address limiting on the interface. This is the default for Switched Uplink Ports, Switched Downlink Ports, and Server Ports.

Table 94: Port Profile Custom Setup Settings (continued)

Field	Action
<p>Allowed MAC List</p>	<p>Indicate the MAC addresses of devices that are allowed access to the interface in the Allowed MAC List. Any device whose MAC address does not match an address in the list is not allowed access to the interface. A list with no entries means that a client with any MAC address is permitted to access the interface.</p> <p>To enter a MAC address, click Add and then type the MAC addresses in the field provided. Enter MAC addresses as two-character hexadecimal numbers separated by colons. Click Save to save the entry.</p> <p>NOTE: Configuring an allowed MAC address list does not block the switch from receiving Layer 2 control packets—such as Link Layer Discovery Protocol (LLDP) packets—transmitted from MAC addresses that are not specified in the address list. Control packets do not undergo the MAC address check. However, the switch does not forward them to another destination.</p> <p>Default: No entries</p>
<p>RSTP Settings</p> <p>In addition to enabling or disabling the Spanning Tree Protocol (STP) as part of device profiles, this feature enables you to fine-tune STP by setting interfaces into edge, disable, or no-root-port states.</p>	<p>Edge</p> <p>RSTP defines the concept of an edge port, which is a designated port that connects to non-STP-capable devices, such as PCs, servers, routers, or hubs that are not connected to other switches. Because edge ports connect directly to end stations, they cannot create network loops and can transition to the forwarding state immediately. You can manually configure edge ports, and a switch can also detect edge ports by noting the absence of communication from the end stations.</p> <p>Disable</p> <p>Disables the RSTP on interface.</p> <p>NOTE: Configuring interfaces to one of these states is not mandatory for ELS switches. Hence, the option Disable is not applicable for ELS switches and therefore not supported.</p> <p>No Root Port</p> <p>Configures an interface to be a spanning-tree designated port. If the bridge receives superior STP bridge protocol data units (BPDUs) on a root-protected interface, that interface transitions to a root-prevented STP state (inconsistency state) and the interface is blocked. This blocking prevents a bridge from being elected the root bridge. When the bridge stops receiving superior STP BPDUs on the root-protected interface, interface traffic is no longer blocked.</p>

Table 94: Port Profile Custom Setup Settings (continued)

Field	Action
CoS Settings	<p>Click Select Cos Profile to choose from existing CoS profiles. The CoS configuration contained in the CoS profile is applied to the interfaces that the Port profile is assigned to when you deploy the configuration. Click OK. Some preconfigured Service Types have a default CoS profile—see the description for Service Types field for details.</p> <p>Or</p> <p>Click Configure CoS settings to configure CoS profile. See “Creating and Managing Wired CoS Profiles” on page 613 for steps to configure a CoS profile.</p>
Authentication Settings (Desktop Port, Desktop Phone Port, Wireless Access Port, Custom Port)	<p>Select the Authentication profile for the interface from a list of existing profiles by clicking Select, selecting one of the listed profiles, and then clicking OK. By assigning an Authentication profile to the Port profile, you can enable 802.1x and captive portal authentication on interfaces.</p> <p>If you do not specify an Authentication profile, the interface is an open port and no authentication is required to connect.</p> <p>NOTE: You cannot configure 802.1x authentication on aggregated Ethernet interfaces. If you attempt to deploy a Port profile that contains an Authentication profile on an aggregated Ethernet interface, the deployment fails.</p> <p>Or</p> <p>Click Configure Authentication Settings to configure 802.1x and captive portal authentications. See “Creating and Managing Authentication Profiles” on page 383 for steps to configure the authentication profile.</p>
Filter Settings (available for all Service Types, including Custom for routing)	<ul style="list-style-type: none"> ● Ingress Filter Select an Ingress Filter for the interface by clicking Select, selecting one of the listed filters, and then clicking OK. ● Egress Filters Select an Egress Filter for the interface by clicking Select, selecting one of the listed filters, and then clicking OK.
VRRP Settings (available when Service Type is Custom and Family Type is Routing)	<p>Select the VRRP profile for the interface from a list of existing profiles by clicking Select. Select one of the listed profiles, and then click OK.</p>

If you configured Custom Setup settings, click **Done**. Upon clicking **Done** displays the dialog **Do you want to assign Port Profile to Ports?**. Click **Yes** to create a profile assignment else click **No** to create the profile and navigate to the Manage Port Profile page to create the Port assignment later.

Specifying Settings for a Campus Switching ELS Port Profile

Use the Create Port Profile page to define a common set of port attributes in a Port profile. You can then apply the Port profile to interfaces on a group of Campus Switching ELS devices.

TIP: You can reference a VLAN profile, CoS profile, Ingress Filter profile, Egress Filter profile, and an Authentication profile in a Port profile. You can either create these profiles in their respective profile pages before you create Port profiles or you can create these profiles as in-line sub-profiles while configuring Port profiles. You can also enable power over Ethernet (PoE).

After you create a Port profile, you can assign it to individual interfaces or to members of a Port group. During this assignment process, you can also configure interface-specific attributes, such as IP address, and certain device-specific attributes, such as the Access profile to use for all ports on the device. You can assign only one Port profile to an interface.

[Table 95](#) describes the Quick Setup settings available in a Port profile. [Table 96](#) describes the Custom Setup settings. The defaults for these options depend on the Service Type you select.

Table 95: Port Profile Quick Setup Settings for Campus Switching ELS

Field	Action
Profile Name	Type the name of profile by using up to 64 alphanumeric characters and no special characters other than the underscore. The name must be unique among Port profiles.
Description	Type a description of the Port profile, which appears on the Manage Port Profiles page. You can use up to 256 characters.

Table 95: Port Profile Quick Setup Settings for Campus Switching ELS (continued)

Field	Action
Service Type	<p>Select one the preconfigured options Desktop Port, Desktop Phone Port, Printer Port, Switched Uplink, Switched Downlink, Server Port, or Wireless Access Port. To create your own service type, select Custom.</p> <hr/> <p>Desktop Port service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—juniper_CS_Hier_Ethernet_CoS for Hierarchical port scheduling; no default profile for Non-Hierarchical port scheduling • Family Type—switching • Port Mode—access • Power over Ethernet—disabled • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • Trust DHCP—disabled • MAC Limit—1 • MAC Limit Action—drop • Allowed MAC List—no default provided

Table 95: Port Profile Quick Setup Settings for Campus Switching ELS (continued)

Field	Action
	<p>Desktop Phone Port service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—juniper_CoS_template for Non-Hierarchical port scheduling; juniper_CS_Hier_Ethernet_CoS for Hierarchical port scheduling • Family Type—switching • Port Mode—access • Power over Ethernet—disabled • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • Trust DHCP—disabled • MAC Limit—2 • MAC Limit Action—drop • Allowed MAC List—no default provided
	<p>Printer Port preconfigured service type has the following default settings:</p> <ul style="list-style-type: none"> • Family Type—switching • Port Mode—access • Power over Ethernet—no default provided • Auto Negotiation—enabled • Flow Control—enabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • Port Security—no default provided • Trust DHCP—no default provided • MAC Limit—no default provided • MAC Limit Action—no default provided • Allowed MAC List—no default provided

Table 95: Port Profile Quick Setup Settings for Campus Switching ELS (continued)

Field	Action
	<p>Switched Uplink service type has the following default settings:</p> <ul style="list-style-type: none"> ● CoS Profile—juniper_CoS_template for Non-Hierarchical port scheduling; juniper_CS_Hier_Ethernet_CoS for Hierarchical port scheduling ● Family Type—switching ● Port Mode—trunk ● Power over Ethernet—disabled ● Auto Negotiation—disabled ● Flow Control—disabled ● Maximum Size—no default provided ● Speed—no default provided ● Link Mode—no default provided ● Port Security—enabled ● MAC Limit—no default provided ● MAC Limit Action—no default provided ● Allowed MAC List—no default provided
	<p>Switched Downlink service type has the following default settings:</p> <ul style="list-style-type: none"> ● CoS Profile—juniper_CoS_template for Non-Hierarchical port scheduling; juniper_CS_Hier_Ethernet_CoS for Hierarchical port scheduling ● Family Type—switching ● Port Mode—trunk ● Power over Ethernet—disabled ● Auto Negotiation—disabled ● Flow Control—disabled ● Maximum Size—no default provided ● Speed—no default provided ● Link Mode—no default provided ● MAC Limit—no default provided ● MAC Limit Action—no default provided ● Allowed MAC list—no default provided

Table 95: Port Profile Quick Setup Settings for Campus Switching ELS (continued)

Field	Action
	<p>Server Port service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—juniper_CoS_template for Non-Hierarchical port scheduling; juniper_CS_Hier_Ethernet_CoS for Hierarchical port scheduling • Family Type—switching • Port Mode—trunk • Power over Ethernet—disabled • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • MAC Limit—no default provided • MAC Limit Action—no default provided • Allowed MAC list—no default provided <hr/> <p>Wireless Access Port service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—juniper_CS_Hier_Ethernet_CoS for Hierarchical port scheduling; no default profile for Non-Hierarchical port scheduling • Family Type—switching • Port Mode—Access • Power over Ethernet—disabled • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • Trust DHCP—disabled • MAC Limit—no default provided • MAC Limit Action—no default provided • Allowed MAC List—no default provided

Port Family Options

The available settings and defaults for these options depend on the Service Type you selected.

Table 95: Port Profile Quick Setup Settings for Campus Switching ELS (continued)

Field	Action
Family Type: Switching or Routing	<p>This setting cannot be changed if any preconfigured Service Type was selected. If you selected the Custom Service Type, indicate whether the interface operates as a Layer 2 (Switching) or a Layer 3 (Routing) interface.</p> <p>TIP: Service Type must be set to Custom to configure a routing interface.</p> <p>If you select routing, you configure an IP address on a per-interface basis when you assign the profile to individual interfaces.</p>
Port Mode for switching interfaces only	<p>This setting cannot be changed if any preconfigured Service Type was selected. If you selected the Custom Service Type, select the port mode for the interface, either Access, Trunk, or Tagged Access.</p> <ul style="list-style-type: none"> ● Access—Use for interfaces that connect to an end device, such as a desktop computer, an IP telephone, a wireless access point, a printer, or a security camera. The interface must belong to a single VLAN. Frames sent and received over the over the interface are untagged Ethernet frames. ● Trunk—Use for interfaces that connect to a switch or router. Trunk interfaces can belong to more than one VLAN, enabling VLAN traffic to be multiplexed on a single physical interface. The Ethernet frames sent and received over the interface are tagged frames, in which IEEE 802.1Q tagging is used to segregate the traffic from each VLAN. ● Tagged Access—Use for access interfaces where VLAN tagging is required, typically when the interface connects to a server running virtual machines using virtual Ethernet port aggregator (VEPA) technology. The traffic generated by the server can contain an aggregation of VLAN packets from different virtual machines on that server, requiring that packets be tagged.

VLAN Options

Available VLAN options depend on the Service Type selected. VLAN association is required for Campus Switching ELS.

Member VLAN (Switched Uplink, Switched Downlink, Server Port)	<p>Click All if you want to assign an interface to all the VLANs.</p> <p>This option is enabled when Port Mode is Trunk or TaggedAccess.</p>
Member VLAN (all Service Types)	<p>This configuration is for one VLAN. Select a VLAN for the interface by clicking Select, selecting one of the listed filters, and then clicking OK.</p> <p>Or</p> <p>Configure a VLAN by clicking Configure VLAN Settings and clicking Create. Enter the VLAN name and ID and click OK.</p>

Table 95: Port Profile Quick Setup Settings for Campus Switching ELS (continued)

Field	Action
Voice VLAN (Desktop Phone Port, Custom Port)	<p>This configuration is for one VLAN. Select a voice VLAN for the interface by clicking Select, selecting one of the listed filters, and then clicking OK.</p> <p>Or</p> <p>Configure a VLAN by clicking Configure VLAN Settings and clicking Create. Enter the VLAN name and ID and click OK.</p>
Native VLAN (Switched Uplink, Switched Downlink)	<p>Select a native VLAN for the interface by clicking Select, selecting one of the listed VLANs, and then clicking OK. The VLAN is added to the Member VLANs list.</p> <p>Or</p> <p>Configure a VLAN by clicking Configure VLAN Settings and clicking Create. Enter the VLAN name and ID and click OK.</p>
Power over Ethernet (PoE)	
Configure Power over Ethernet	<p>Enable to configure PoE settings.</p> <p>If you do not enable this option, Network Director does not send any PoE configuration commands to the device when the profile is deployed on the device. For example, if PoE is enabled on an interface, it remains enabled.</p> <p>If you enable this option, the PoE settings in this profile is deployed on the interfaces that support PoE. If you assign this Port profile to interfaces that do not support PoE, the profile can be deployed successfully on those interfaces, but the PoE settings do not take effect.</p> <p>TIP: EX9200 switches do not support PoE.</p>

Table 95: Port Profile Quick Setup Settings for Campus Switching ELS (continued)

Field	Action
Maximum Power (W)	<p>Use the arrows to adjust the maximum PoE power allocated to a PoE port in watts. The powered device cannot draw more power than the wattage specified. If it does so, PoE power to the port is shut down.</p> <p>The Maximum Power setting has no effect when the PoE management mode for the switch or line card is class mode, which is the default mode. In class mode, the power allocated to a PoE port is determined either by LLDP negotiation with the powered device or by the PoE class of the powered device if LLDP is not supported.</p> <p>You must set the PoE power management mode for the switch or line card to static mode for the Maximum Power setting to take effect. You can do this in the Device Common Settings profile.</p> <p>If you specify a maximum wattage that is greater than the maximum wattage that can be supplied by the port, your configuration is accepted when the Port profile is deployed on the port. However, the maximum wattage is set to the port's maximum supported wattage. The maximum supported wattages for PoE ports are:</p> <ul style="list-style-type: none"> • 15.4W for ports that support IEEE 802.3af only • 18.6W for IEEE 802.3af ports on switches that support enhanced PoE • 30W for ports that support IEEE 802.3at <p>Default: 15.4W</p>
Priority	<p>Select a power priority for the PoE port—either Low or High. If there is a shortage of PoE power on the switch, power to low priority ports is shut down before power to high priority ports. Within ports with the same assigned priority, power priority is determined by port number—ports with a lower port number have a higher power priority.</p> <p>Default: low priority</p>
Disable PoE	Select to disable PoE on the interface.

After providing the information in the fields listed in [Table 95](#), click **Done**.

To use default Port Profile Custom Setup settings, click **Done**. To configure Custom Setup settings, click **Custom Setup** and then provide the information in [Table 96](#) and then click **Done**.

Clicking **Done** in either case displays the dialog Do you want to assign Port Profile to Ports. Click **Yes** to create a profile assignment; else click **No** to navigate to the Manage Port Profile page and to create the Port assignment later on.

Table 96: Port Profile Custom Setup Settings for Campus Switching ELS

Field	Action
<p>Advanced Settings</p>	
<p>Expand Advanced Settings to configure link settings and port security. The Link Setting in Port profile is disabled by default. On enabling Link Settings, autonegotiation and flow control are enabled by default.</p>	
<p>Enable Auto Negotiation</p>	<p>Autonegotiation of link speed and duplex mode is enabled by default; clear to disable autonegotiation.</p> <p>If you disable autonegotiation, you must set link speed and link mode.</p> <p>You cannot disable autonegotiation if a link speed of 1 Gbps is configured. This configuration might be accepted, but autonegotiation will not be disabled.</p> <p>NOTE: This setting is ignored when you assign a Port profile to an Aggregated Ethernet interface.</p>
<p>Enable Flow Control</p>	<p>Select to enable flow control on the interface, which permits the switch suspend packet transmission for a set period of time in response to a PAUSE frame sent by a congested switch.</p> <p>Flow control applies only to links operating at 1 Gbps, full-duplex mode.</p>
<p>MTU</p>	<p>Using the arrows, indicate the maximum transmission unit (MTU), which is the maximum size of Ethernet frames sent by the interface. To calculate the MTU, add 14 bytes overhead to the maximum payload you want sent.</p> <p>Range: 256 through 9216 bytes</p>
<p>Speed</p>	<p>Select the link speed.</p> <p>If you select a link speed when autonegotiation is enabled, autonegotiation remains enabled and the interface will advertise the link speed that you specify as its maximum link speed.</p> <p>NOTE: This setting is ignored when you assign a Port profile to an Aggregated Ethernet interface.</p>
<p>Link Mode</p>	<p>Select the duplex mode, either Automatic, Full Duplex, or Half Duplex. Select Automatic to enable autonegotiation when autonegotiation is disabled.</p> <p>NOTE: This setting is ignored when you assign a Port profile to an Aggregated Ethernet interface.</p> <p>You cannot select Half Duplex with link speed set to Autonegotiation or 1 Gbps.</p>

Table 96: Port Profile Custom Setup Settings for Campus Switching ELS (continued)

Field	Action
<p>Storm Control Settings</p> <p>Enabling storm control on a switching device monitors traffic levels and drops broadcast, multicast, and unknown unicast packets when a specified traffic level—called the storm control level or storm control bandwidth—is exceeded, thus preventing packets from proliferating and degrading the LAN.</p> <p>You can customize the storm control level for a specific interface by explicitly configuring either bandwidth or level.</p> <p>NOTE: You cannot configure both bandwidth and level for the same interface.</p>	<p>Unit</p> <ul style="list-style-type: none"> <p>Percentage—Configures the storm control level as a percentage of the available bandwidth used by the combined applicable traffic streams that are subject to storm control on that interface.</p> <p>The level can be set from 0% to 100%, where 0% indicates that the entire traffic is being suppressed and 100% indicates no traffic is being suppressed, in other words there is no storm control.</p> <p>The default level is 80%.</p> <p>Kbps—Configures the storm control level as the bandwidth in kilobits per second (Kbps) of the applicable traffic streams on that interface.</p> <p>Set the bandwidth from 100 through 10,000,000 in Kbps. When you configure storm control bandwidth, the value you configure is rounded off internally to the closest multiple of 64 Kbps, and the rounded-off value represents the bandwidth that is actually used. For example, if you configure a bandwidth limit of 150 Kbps, storm control uses a bandwidth limit of 128 Kbps.</p> <p>Value</p> <p>Configures the traffic storm control threshold level value as a percentage of bandwidth or bandwidth in kilobits per second depending upon the specified unit.</p>
<p>Power over Ethernet (PoE)</p>	
<p>You can enable PoE and display the configuration options by enabling Configure Power over Ethernet.</p>	
<p>Configure Power over Ethernet</p>	<p>Enable to configure PoE settings.</p> <p>If you do not enable this option, Network Director does not send any PoE configuration commands to the device when the profile is deployed on the device. For example, if PoE is enabled on an interface, it remains enabled. On EX Series switches, the factory-default configuration enables PoE on all interfaces that support PoE.</p> <p>If you enable this option, the PoE settings in this profile is deployed on the interfaces that support PoE. If you assign this Port profile to interfaces that do not support PoE, the profile is deployed successfully on those interfaces, but the PoE settings will not take effect.</p>

Table 96: Port Profile Custom Setup Settings for Campus Switching ELS (continued)

Field	Action
Maximum Power (W)	<p>Use the arrows to adjust the maximum PoE power in watts allocated to a PoE port. The powered device cannot draw more power than the wattage specified. If it does so, PoE power to the port is shut down. Maximum power for PoE is 15.4W, Extended PoE is 18.6W and PoE+ is 30W.</p> <p>The Maximum Power setting has no effect when the PoE management mode for a switch or line card is class mode, which is the default mode. In class mode, the power allocated to a PoE port is determined either by LLDP negotiation with the powered device or by the PoE class of the powered device if LLDP is not supported.</p> <p>You must set the PoE power management mode for the switch or line card to static mode for the Maximum Power setting to take effect. Do this in the Device Common Settings profile.</p> <p>If you specify a maximum wattage that is greater than the maximum wattage that can be supplied by the port, your configuration is accepted when the Port profile is deployed on the port. However, the maximum wattage is set to the port's maximum supported wattage. The maximum supported wattages for PoE ports are:</p> <ul style="list-style-type: none"> • 15.4W for ports that support IEEE 802.3af only • 18.6W for IEEE 802.3af ports on switches that support enhanced PoE • 30W for ports that support IEEE 802.3at <p>Default: 15.4W</p>
Priority	<p>Select a power priority for the PoE port—either Low or High. If there is a shortage of PoE power on the switch, power to low priority ports is shut down before power to high priority ports. Within ports with the same assigned priority, power priority is determined by the port number—ports with a lower port number have a higher power priority.</p> <p>Default: low priority</p>
Disable PoE	<p>Select to disable PoE on the interfaces using this Port profile.</p>

Port Security (Switching Interfaces Only)

Select to enable port security (default); clear to disable port security.

When port security is enabled, you can configure port security options such as learned MAC address limits on an interface. When port security is disabled, no port security is applied to the interface, including the default port security options.

Table 96: Port Profile Custom Setup Settings for Campus Switching ELS (continued)

Field	Action
Trust DHCP	<p>Select to permit messages from a DHCP server to be received on the interface—this is the default. Clear to block all messages from a DHCP server from being received on the interface.</p> <p>TIP: For this port security feature to work, DHCP snooping must be enabled on the VLAN the interface belongs to. You can enable DHCP snooping on the VLAN in the VLAN profile. For directions, see “Creating and Managing VLAN Profiles” on page 502.</p>
MAC Limit	<p>Type the number of MAC address that can be dynamically learned on the interface.</p> <p>Range: 1 through 163839</p> <p>Default: For Desktop Ports, 1. For Desktop Phone Ports, 2. For all others, none.</p>
MAC Limit Action	<p>Select the action to be taken if the MAC address limit is exceeded:</p> <ul style="list-style-type: none"> ● Drop—Drop any packet with a previously unlearned MAC address and generate a system log entry, and SNMP trap, or an alarm. This is the default for a Desktop Port and Desktop Phone Ports. ● Log—Accept packets with new MAC addresses and learn the addresses, but generate a system log entry, and SNMP trap, or an alarm. ● Shutdown—Shut down the interface and generate a system log message, SNMP trap, or an alarm. <p>If an interface is shut down because the MAC address limit has been exceeded, you must use the CLI command clear ethernet-switching port-error interface <i>name</i> to clear the error and bring the interface back into service.</p> <p>TIP: You can use the CLI to configure auto-recovery on an interface that has been shut down by a MAC limit error.</p> <ul style="list-style-type: none"> ● None—No action. This selection effectively disables MAC address limiting on the interface. This is the default for Switched Uplink Ports, Switched Downlink Ports, and Server Ports.

Table 96: Port Profile Custom Setup Settings for Campus Switching ELS (continued)

Field	Action
<p>Allowed MAC List</p>	<p>Indicate the MAC addresses of devices that are allowed access to the interface in the Allowed MAC List. Any device whose MAC address does not match an address in the list will not be allowed access to the interface. A list with no entries means that a client with any MAC address is permitted to access the interface.</p> <p>To enter a MAC address, click Add and then type the MAC addresses in the field provided. Enter MAC addresses as two-character hexadecimal numbers separated by colons. Click Save to save the entry.</p> <p>NOTE: Configuring an allowed MAC address list does not block the switch from receiving Layer 2 control packets—such as Link Layer Discovery Protocol (LLDP) packets—transmitted from MAC addresses that are not specified in the address list. Control packets do not undergo the MAC address check. However, the switch does not forward them to another destination.</p> <p>Default: No entries</p>
<p>RSTP Settings</p> <p>In addition to enabling or disabling the Spanning Tree Protocol (STP) as part of device profiles, this feature enables you to fine-tune STP by setting interfaces into edge, disable, or no-root-port states.</p>	<p>Edge</p> <p>RSTP defines the concept of an edge port, which is a designated port that connects to non-STP-capable devices, such as PCs, servers, routers, or hubs that are not connected to other switches. Because edge ports connect directly to end stations, they cannot create network loops and can transition to the forwarding state immediately. You can manually configure edge ports, and a switch can also detect edge ports by noting the absence of communication from the end stations.</p> <p>Disable</p> <p>Disables the RSTP on interface.</p> <p>NOTE: Configuring interfaces to one of these states is not mandatory for ELS switches. Hence, the option Disable is not applicable for ELS switches and therefore not supported.</p> <p>No Root Port</p> <p>Configures an interface to be a spanning-tree designated port. If the bridge receives superior STP bridge protocol data units (BPDUs) on a root-protected interface, that interface transitions to a root-prevented STP state (inconsistency state) and the interface is blocked. This blocking prevents a bridge from being elected the root bridge. When the bridge stops receiving superior STP BPDUs on the root-protected interface, interface traffic is no longer blocked.</p>

Table 96: Port Profile Custom Setup Settings for Campus Switching ELS (continued)

Field	Action
CoS Settings	<p>Click Select Cos Profile to choose from existing CoS profiles. The CoS configuration contained in the CoS profile is applied to the interfaces that the Port profile is assigned to when you deploy the configuration. Click OK. Some preconfigured Service Types have a default CoS profile—see Service Types for details.</p> <p>Or</p> <p>Click Configure CoS settings to configure CoS profile. See “Creating and Managing Wired CoS Profiles” on page 613 for steps to configure a CoS profile.</p>
Authentication Settings (Desktop Port, Desktop Phone Port, Wireless Access Port, Custom Port)	<p>Select the Authentication profile for the interface from a list of existing profiles by clicking Select, selecting one of the listed profiles, and then clicking OK. By assigning an Authentication profile to the Port profile, you can enable 802.1x and captive portal authentication on interfaces.</p> <p>If you do not specify an Authentication profile, the interface is an open port and no authentication is required to connect.</p> <p>NOTE: You cannot configure 802.1x authentication on aggregated Ethernet interfaces. If you attempt to deploy a Port profile that contains an Authentication profile on an aggregated Ethernet interface, the deployment fails.</p> <p>Or</p> <p>Click Configure Authentication Settings to configure 802.1x and captive portal authentications. See “Creating and Managing Authentication Profiles” on page 383 for steps to configure the Authentication profile.</p>
Filter Settings (available for all Service Types, including Custom for routing)	<ul style="list-style-type: none"> ● Ingress Filter Select an Ingress Filter for the interface by clicking Select, selecting one of the listed filters, and then clicking OK. ● Egress Filters Select an Egress Filter for the interface by clicking Select, selecting one of the listed filters, and then clicking OK.
VRRP Settings (available when Service Type is Custom and Family Type is Routing)	<p>Select the VRRP profile for the interface from a list of existing profiles by clicking Select. Select one of the listed profiles, and then click OK.</p>

Clicking **Done** displays the dialog Do you want to assign Port Profile to Ports. Click **Yes** to create a profile assignment else click **No** to navigate to the Manage Port Profile page and to create the Port assignment later.

Click **Done** to save the Port profile for Campus Switching ELS.

Specifying Settings for the Data Center Switching Non-ELS Port Profile

Use the Create Port Profile page to define a common set of attributes in a Port profile, which you can then apply to a group of interfaces on Data Center Switching non-ELS devices. Data Center Switching Non-ELS Port profiles can apply to Ethernet or Fibre Channel (FC) ports. You can create a new Port profile from scratch, or select an appropriate Service Type and use the default settings that Network Director has defined for that service type to create a Port profile.

TIP: You can reference a VLAN profile, CoS profile, Ingress Filter profile, Egress Filter profile, and an Authentication profile in a Port profile. You can either create these profiles in their respective profile pages before you create Port profiles or you can create these profiles as in-line sub-profiles while configuring Port profiles. You can also enable power over Ethernet (PoE).

After you create a Port profile, you can assign it to either individual interfaces or to members of port groups. During this process, you can also configure interface-specific attributes, such as IP address, and certain device-specific attributes, such as the Access profile to use for all ports on the device. You can assign only one Port profile to either an interface or port group.

[Table 97](#) describes the Quick Setup settings available in the data center Port profile.

[Table 98](#) describes the Custom Setup settings available for an Non-ELS Ethernet ports.

Table 97: Port Profile Quick Setup Settings for Data Center Non-ELS Ports

Field	Action
Profile Name	Type the name of profile by using up to 64 alphanumeric characters and no special characters other than the underscore. The name must be unique among Port profiles.
Description	Type a description of the Port profile, which will appear on Manage Port Profiles page. You can use up to 256 characters.

Table 97: Port Profile Quick Setup Settings for Data Center Non-ELS Ports (continued)

Field	Action
Service Type	<p>Select a service type for the profile. The service type describes the function that the port will serve. Selecting a service type automatically configures some of the other page fields to support that service type. Some of the automatic settings are mandatory for the service type, so you cannot edit those fields.</p> <p>Select one the preconfigured switching options, Desktop Port, Desktop Phone Port, Switched Uplink, Switched Downlink, Server Port, FCoE Gateway, Fibre Channel Port, or FCoE Transit Port. To create your own switching or routing service type, select Custom.</p> <p>TIP: No preconfigured routing Service Types are provided. You must create them by using the Custom option.</p> <hr/> <p>Desktop Port service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—no default provided • Family Type—switching • Port Mode—access • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • Trust DHCP—disabled • MAC Limit—1 • MAC Limit Action—drop • Allowed MAC List—no default provided

Table 97: Port Profile Quick Setup Settings for Data Center Non-ELS Ports (continued)

Field	Action
	<p>Desktop Phone Port service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—juniper_CoS_template • Family Type—switching • Port Mode—access • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • Trust DHCP—disabled • MAC Limit—2 • MAC Limit Action—drop • Allowed MAC List—no default provided
	<p>Switched Uplink service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—juniper_CoS_template • Family Type—switching • Port Mode—trunk • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • MAC Limit—no default provided • MAC Limit Action—no default provided • Allowed MAC List—no default provided

Table 97: Port Profile Quick Setup Settings for Data Center Non-ELS Ports (continued)

Field	Action
	<p>Switched Downlink service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—juniper_CoS_template • Family Type—switching • Port Mode—trunk • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • MAC Limit—no default provided • MAC Limit Action—no default provided • Allowed MAC list—no default provided
	<p>Server Port service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—juniper_CoS_template • Family Type—switching • Port Mode—trunk • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • MAC Limit—no default provided • MAC Limit Action—no default provided • Allowed MAC list—no default provided

Table 97: Port Profile Quick Setup Settings for Data Center Non-ELS Ports (*continued*)

Field	Action
	<p>FCoE Gateway service type has the following default settings:</p> <ul style="list-style-type: none"> ● Port Type—Ethernet Port ● CoS Profile—juniper_DC_Hier_CoS ● Family Type—switching ● Port Mode—trunk ● Filters—no default provided ● VLAN Options—no default provided ● DCBX Version—Auto ● Disable DCBX—disabled ● Disable Priority Flow Control—disabled ● ETS No Auto Negotiation—disabled ● Recommendation TVL—no default provided ● Auto Negotiation—disabled ● Flow Control—disabled ● Maximum Size—2500 ● Speed—no default provided ● Link Mode—no default provided <hr/> <p>Fibre Channel Port service type has the following default settings:</p> <ul style="list-style-type: none"> ● Port Type—Fibre Channel Port ● Speed—4Gbps ● Buffer to Buffer State Change Number—no default provided ● Loopback Setting—no default provided

Table 97: Port Profile Quick Setup Settings for Data Center Non-ELS Ports (continued)

Field	Action
	<p>FCoE Transit Port service type has the following default settings:</p> <ul style="list-style-type: none"> • Port Type—Ethernet Port • CoS Profile—juniper_DC_Hier_CoS • Family Type—switching • Port Mode—trunk • Filters—no default provided • VLAN Options—no default provided • DCBX Version—Auto • Disable DCBX—disabled • Disable Priority Flow Control—disabled • ETS No Auto Negotiation—disabled • Recommendation TVL—no default provided • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—2500 • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • Trust DHCP—disabled • FCoE Trusted—enabled • MAC Limit—no default provided • MAC Limit Action—no default provided • Allowed MAC List—no default provided
<p>Family Type The available settings and defaults for these options depend on the Service Type you selected.</p>	<p>Switching—Select whether the interface operates as a Layer 2 (switching) or a Layer 3 (routing) interface.</p> <p>TIP: Routing—If you select routing, you configure an IP address on a per-interface basis when you assign the profile to individual interfaces.</p>

Table 97: Port Profile Quick Setup Settings for Data Center Non-ELS Ports (continued)

Field	Action
Port Mode	<p>(Switching interfaces only) Select the port mode for the interface:</p> <ul style="list-style-type: none"> • Access—Use for interfaces that connect to an end device, such as a desktop computer, an IP telephone, a wireless access point, a printer, or a security camera. The interface must belong to a single VLAN. Frames sent and received over the interface are untagged Ethernet frames. • Trunk—Use for interfaces that connect to a switch or router. Trunk interfaces can belong to more than one VLAN, enabling VLAN traffic to be multiplexed on a single physical interface. The Ethernet frames sent and received over the interface are tagged frames, in which IEEE 802.1Q tagging is used to segregate the traffic from each VLAN. • Tagged Access—Use for access interfaces where VLAN tagging is required, typically when the interface connects to a server running virtual machines using VEPA technology. The traffic generated by the server can contain an aggregation of VLAN packets from different virtual machines on that server, requiring that packets be tagged.
Port Type	<p>If you selected the service type Custom, select the port type. If you selected a service type other than Custom, you cannot edit this field. Each port type has a different set of fields to configure. The options are Ethernet Port or Fibre Channel Port.</p>

VLAN Options

Available VLAN options depend on the Service Type selected.

Member VLAN (available for Switched Uplink, Switched Downlink, Server Port, FCoE Gateway, FCoE Transit port)	<p>Click All if you want to assign an interface to all the VLANs.</p> <p>This option is enabled when Port Mode is Trunk or TaggedAccess.</p>
Member VLANs (All except Fibre Channel Port)	<p>Select a VLAN for the interface by clicking Select, selecting one of the listed filters, and then clicking OK. The VLAN is added to the Member VLANs list.</p> <p>Or</p> <p>Configure a VLAN by clicking Configure VLAN Settings and clicking Create. Enter the VLAN name and ID and click OK.</p>

Table 97: Port Profile Quick Setup Settings for Data Center Non-ELS Ports (*continued*)

Field	Action
Voice VLAN (available for Desktop Phone Port, Custom Port)	<p>Select a voice VLAN for the interface by clicking Select, selecting one of the listed filters, and then clicking OK. The VLAN is added to the Member VLANs list.</p> <p>Or</p> <p>Configure a VLAN by clicking Configure VLAN Settings and clicking Create. Enter the VLAN name and ID and click OK.</p>
Native VLAN (available for Switched Uplink, Switched Downlink)	<p>Select a native VLAN for the interface by clicking Select, selecting one of the listed VLANs, and then clicking OK. The VLAN is added to the Member VLANs list.</p> <p>Or</p> <p>Configure a VLAN by clicking Configure VLAN Settings and clicking Create. Enter the VLAN name and ID and click OK.</p>
Fibre Channel Settings (Fibre Channel Port)	<p>Speed—Select the FC speed: Auto (default), 2Gbps, 4Gbps, or 8Gbps.</p> <p>Buffer to Buffer State Change Number—Configure the buffer-to-buffer credit state change number to prevent the permanent loss of Fibre Channel credits over time (buffer-to-buffer credit recovery). Select a number from 0 through 15.</p> <p>Loopback Setting—Select Loopback to configure an FC loopback interface.</p>

DCBX Settings

Data Center Bridging Capability Exchange protocol is a discovery and exchange protocol for conveying configuration and capabilities among network neighbors to ensure consistent configuration across the network. It is an extension of the Link Layer Data Protocol (LLDP, described in IEEE 802.1AB).

The defaults for these settings depend on the Service Type you selected.

DCBX Version	<p>Select one of the following versions of the Data Center Bridging Capability Exchange protocol:</p> <ul style="list-style-type: none"> • Auto—automatic configuration • DCBX v1.01—The Converged Enhanced Ethernet (CEE) version of DCBX. It has a subtype of 2 and an IEEE DCBX Organizationally Unique Identifier (OUI) of 0x001b21. • IEEE DCBX—The newest DCBX version. Different TLVs have different subtypes (for example, the subtype for the ETS configuration TLV is 9); the OUI is 0x0080c2.
--------------	--

Table 97: Port Profile Quick Setup Settings for Data Center Non-ELS Ports (continued)

Field	Action
Disable DCBX	Select this option to turn off Data Center Bridging Capability Exchange protocol.
Disable Priority Flow Control	Select this option to turn off priority flow control. Priority-based flow control (PFC) is a link-level flow control mechanism defined by IEEE 802.1Qbb that enables independent flow control for each class of service (as defined in the 3-bit CoS field of the Ethernet header by IEEE 802.1Q tags) to ensure that no frame loss from congestion occurs in DCB networks.
ETS No Auto Negotiation	Select this option to turn off ETS autonegotiation. Enhanced transmission selection (ETS) is a mechanism that provides finer granularity of bandwidth management within a link.
Recommendation TLV	Select either Enable TLV or Disable TLV. The enhanced transmission selection (ETS) Recommendation TLV communicates the ETS settings that the switch wants the connected peer interface to use. If the peer interface is willing, the peer interface changes its configuration to match the configuration in the ETS Recommendation TLV. By default, the switch interfaces send the ETS Recommendation TLV to the peer. The settings communicated are the egress ETS settings defined by configuring hierarchical scheduling on the interface.

Table 98: Port Profile Custom Setup Settings for Data Center Non-ELS Ethernet Ports

Field	Action
Advanced Settings	
Expand Advanced Settings to configure link settings and port security. The Link Setting in Port profile is disabled by default. On enabling Link Settings, autonegotiation and flow control are enabled by default.	
Enable Auto Negotiation	Autonegotiation of link speed and duplex mode is enabled by default; clear to disable autonegotiation. If you disable autonegotiation, you must set link speed and link mode. You cannot disable autonegotiation if a link speed of 1 Gbps is configured. This configuration might be accepted, but autonegotiation is not disabled. NOTE: This setting is ignored when you assign a Port profile to an Aggregated Ethernet interface.

Table 98: Port Profile Custom Setup Settings for Data Center Non-ELS Ethernet Ports (continued)

Field	Action
Enable Flow Control	<p>Select to enable flow control on the interface, which permits the switch suspend packet transmission for a set period of time in response to a PAUSE frame sent by a congested switch.</p> <p>Flow control applies only to links operating at 1 Gbps, full-duplex mode.</p>
MTU	<p>Using the arrows, indicate the maximum transmission unit (MTU), which is the maximum size of Ethernet frames sent by the interface. To calculate the MTU, add 14 bytes overhead to the maximum payload you want sent.</p> <p>Range: 256 through 9216 bytes</p>
Speed	<p>Select the link speed.</p> <p>If you select a link speed when autonegotiation is enabled, autonegotiation remains enabled and the interface will advertise the link speed that you specify as its maximum link speed.</p> <p>NOTE: This setting is ignored when you assign a Port profile to an Aggregated Ethernet interface.</p>
Link Mode	<p>Select the duplex mode, either Automatic, Full Duplex, or Half Duplex. Select Automatic to enable autonegotiation when autonegotiation is disabled.</p> <p>NOTE: This setting is ignored when you assign a Port profile to an Aggregated Ethernet interface.</p> <p>You cannot select Half Duplex with link speed set to Autonegotiation or 1 Gbps.</p>

Table 98: Port Profile Custom Setup Settings for Data Center Non-ELS Ethernet Ports (continued)

Field	Action
<p>Storm Control Settings</p> <p>Enabling storm control on a switching device monitors traffic levels and drops broadcast, multicast, and unknown unicast packets when a specified traffic level—called the storm control level or storm control bandwidth—is exceeded, thus preventing packets from proliferating and degrading the LAN.</p> <p>You can customize the storm control level for a specific interface by explicitly configuring either bandwidth or level.</p> <p>NOTE: You cannot configure both bandwidth and level for the same interface.</p>	<p>Unit</p> <ul style="list-style-type: none"> <p>Percentage—Configures the storm control level as a percentage of the available bandwidth used by the combined applicable traffic streams that are subject to storm control on that interface.</p> <p>The level can be set from 0% to 100%, where 0% indicates that the entire traffic is being suppressed and 100% indicates no traffic is being suppressed, in other words there is no storm control.</p> <p>The default level is 80%.</p> <p>Kbps—Configures the storm control level as the bandwidth in kilobits per second (Kbps) of the applicable traffic streams on that interface.</p> <p>Set the bandwidth from 100 through 10,000,000 in Kbps. When you configure storm control bandwidth, the value you configure is rounded off internally to the closest multiple of 64 Kbps, and the rounded-off value represents the bandwidth that is actually used. For example, if you configure a bandwidth limit of 150 Kbps, storm control uses a bandwidth limit of 128 Kbps.</p> <p>Value</p> <p>Configures the traffic storm control threshold level value as a percentage of bandwidth or bandwidth in kilobits per second depending upon the specified unit.</p>
<p>Power over Ethernet (PoE)</p>	
<p>You can enable PoE and display the configuration options by enabling Configure Power over Ethernet.</p>	
<p>Configure Power over Ethernet</p>	<p>Enable to configure PoE settings.</p> <p>If you do not enable this option, Network Director does not send any PoE configuration commands to the device when the profile is deployed on the device. For example, if PoE is enabled on an interface, it remains enabled. On EX Series switches, the factory-default configuration enables PoE on all interfaces that support PoE.</p> <p>If you enable this option, the PoE settings in this profile is deployed on the interfaces that support PoE. If you assign this Port profile to interfaces that do not support PoE, the profile is deployed successfully on those interfaces, but the PoE settings will not take effect.</p>

Table 98: Port Profile Custom Setup Settings for Data Center Non-ELS Ethernet Ports (continued)

Field	Action
Maximum Power (W)	<p>Use the arrows to adjust the maximum PoE power in watts allocated to a PoE port. The powered device cannot draw more power than the wattage specified. If it does so, PoE power to the port is shut down. Maximum power for PoE is 15.4W, Extended PoE is 18.6W and PoE+ is 30W.</p> <p>The Maximum Power setting has no effect when the PoE management mode for a switch or line card is class mode, which is the default mode. In class mode, the power allocated to a PoE port is determined either by LLDP negotiation with the powered device or by the PoE class of the powered device if LLDP is not supported.</p> <p>You must set the PoE power management mode for the switch or line card to static mode for the Maximum Power setting to take effect. Do this in the Device Common Settings profile.</p> <p>If you specify a maximum wattage that is greater than the maximum wattage that can be supplied by the port, your configuration is accepted when the Port profile is deployed on the port. However, the maximum wattage is set to the port's maximum supported wattage. The maximum supported wattages for PoE ports are:</p> <ul style="list-style-type: none"> • 15.4W for ports that support IEEE 802.3af only • 18.6W for IEEE 802.3af ports on switches that support enhanced PoE • 30W for ports that support IEEE 802.3at <p>Default: 15.4W</p>
Priority	<p>Select a power priority for the PoE port—either Low or High. If there is a shortage of PoE power on the switch, power to low priority ports is shut down before power to high priority ports. Within ports with the same assigned priority, power priority is determined by the port number—ports with a lower port number have a higher power priority.</p> <p>Default: low priority</p>
Disable PoE	<p>Select to disable PoE on the interfaces using this Port profile.</p>

Port Security (Switching Interfaces Only)

Select to enable port security (default); clear to disable port security.

When port security is enabled, you can configure port security options such as learned MAC address limits on an interface. When port security is disabled, no port security is applied to the interface, including the default port security options.

Table 98: Port Profile Custom Setup Settings for Data Center Non-ELS Ethernet Ports (continued)

Field	Action
Trust DHCP	<p>Select to permit messages from a DHCP server to be received on the interface—this is the default. Clear to block all messages from a DHCP server from being received on the interface.</p> <p>TIP: For this port security feature to work, DHCP snooping must be enabled on the VLAN the interface belongs to. You can enable DHCP snooping on the VLAN in the VLAN profile. For directions, see “Creating and Managing VLAN Profiles” on page 502.</p>
FCoE Trusted	<p>Select to configure the interface to trust Fibre Channel over Ethernet (FCoE) traffic. If an interface is connected to another switch such as an FCoE forwarder (FCF) or a transit switch, you can configure the interface as trusted so that the interface forwards FCoE traffic from the switch to the FCoE devices without installing FIP snooping filters.</p>
MAC Limit	<p>Type the number of MAC address that can be dynamically learned on the interface.</p> <p>Range: 1 through 163839</p> <p>Default: For Desktop Ports, 1. For Desktop Phone Ports, 2. For all others, none.</p>
MAC Limit Action	<p>Select the action to be taken if the MAC address limit is exceeded:</p> <ul style="list-style-type: none"> ● Drop—Drop any packet with a previously unlearned MAC address and generate a system log entry, and SNMP trap, or an alarm. This is the default for a Desktop Port and Desktop Phone Ports. ● Log—Accept packets with new MAC addresses and learn the addresses, but generate a system log entry, and SNMP trap, or an alarm. ● Shutdown—Shut down the interface and generate a system log message, SNMP trap, or an alarm. <p>If an interface is shut down because the MAC address limit has been exceeded, you must use the CLI command clear ethernet-switching port-error interface <i>name</i> to clear the error and bring the interface back into service.</p> <p>TIP: You can use the CLI to configure auto-recovery on an interface that has been shut down by a MAC limit error.</p> <ul style="list-style-type: none"> ● None—No action. This selection effectively disables MAC address limiting on the interface. This is the default for Switched Uplink Ports, Switched Downlink Ports, and Server Ports.

Table 98: Port Profile Custom Setup Settings for Data Center Non-ELS Ethernet Ports (continued)

Field	Action
<p>Allowed MAC List</p>	<p>Indicate the MAC addresses of devices that are allowed access to the interface in the Allowed MAC List. Any device whose MAC address does not match an address in the list will not be allowed access to the interface. A list with no entries means that a client with any MAC address is permitted to access the interface.</p> <p>To enter a MAC address, click Add and then type the MAC addresses in the field provided. Enter MAC addresses as two-character hexadecimal numbers separated by colons. Click Save to save the entry.</p> <p>NOTE: Configuring an allowed MAC address list does not block the switch from receiving Layer 2 control packets—such as Link Layer Discovery Protocol (LLDP) packets—transmitted from MAC addresses that are not specified in the address list. Control packets do not undergo the MAC address check. However, the switch does not forward them to another destination.</p> <p>Default: No entries</p>
<p>RSTP Settings</p> <p>In addition to enabling or disabling the Spanning Tree Protocol (STP) as part of device profiles, this feature enables you to fine-tune STP by setting interfaces into edge, disable, or no-root-port states.</p>	<p>Edge</p> <p>RSTP defines the concept of an edge port, which is a designated port that connects to non-STP-capable devices, such as PCs, servers, routers, or hubs that are not connected to other switches. Because edge ports connect directly to end stations, they cannot create network loops and can transition to the forwarding state immediately. You can manually configure edge ports, and a switch can also detect edge ports by noting the absence of communication from the end stations.</p> <p>Disable</p> <p>Disables the RSTP on interface.</p> <p>NOTE: Configuring interfaces to one of these states is not mandatory for ELS switches. Hence, the option Disable is not applicable for ELS switches and therefore not supported.</p> <p>No Root Port</p> <p>Configures an interface to be a spanning-tree designated port. If the bridge receives superior STP bridge protocol data units (BPDUs) on a root-protected interface, that interface transitions to a root-prevented STP state (inconsistency state) and the interface is blocked. This blocking prevents a bridge from being elected the root bridge. When the bridge stops receiving superior STP BPDUs on the root-protected interface, interface traffic is no longer blocked.</p>

Table 98: Port Profile Custom Setup Settings for Data Center Non-ELS Ethernet Ports (continued)

Field	Action
CoS Settings (All except Fibre Channel Type)	<p>Click Select Cos Profile to choose from existing CoS profiles. The CoS configuration contained in the CoS profile is applied to the interfaces that the Port profile is assigned to when you deploy the configuration. Click OK. Some preconfigured Service Types have a default CoS profile—see Service Types for details.</p> <p>Or</p> <p>Click Configure CoS settings to configure CoS profile. See “Creating and Managing Wired CoS Profiles” on page 613 for steps to configure a CoS profile.</p>
Authentication Settings (Desktop Port, Desktop Phone Port, Custom Port)	<p>Select the Authentication profile for the interface from a list of existing profiles by clicking Select, selecting one of the listed profiles, and then clicking OK. By assigning an Authentication profile to the Port profile, you can enable 802.1x and captive portal authentication on interfaces.</p> <p>If you do not specify an Authentication profile, the interface is an open port and no authentication is required to connect.</p> <p>NOTE: You cannot configure 802.1x authentication on aggregated Ethernet interfaces. If you attempt to deploy a Port profile that contains an Authentication profile on an aggregated Ethernet interface, the deployment fails.</p> <p>Or</p> <p>Click Configure Authentication Settings to configure 802.1x and captive portal authentications. See “Creating and Managing Authentication Profiles” on page 383 for steps to configure the authentication profile.</p>
Filter Settings (available for all Service Types, including Custom for routing)	<ul style="list-style-type: none"> ● Ingress Filter Select an Ingress Filter for the interface by clicking Select, selecting one of the listed filters, and then clicking OK. ● Egress Filters Select an Egress Filter for the interface by clicking Select, selecting one of the listed filters, and then clicking OK.
VRRP Settings (available when Service Type is Custom and Family Type is Routing)	<p>Select the VRRP profile for the interface from a list of existing profiles by clicking Select. Select one of the listed profiles, and then click OK.</p>

Specifying Settings for a Data Center Switching ELS Port Profile

Use the Create Port Profile page to define a common set of port attributes in a Port profile. You can create a new Port profile from scratch, or select an appropriate Service Type and use the default settings that

Network Director has defined for that service type to create a Port profile. You can then apply the Port profile to interfaces on a group of Data Center Switching ELS devices.

TIP: You can reference a VLAN profile, CoS profile, Ingress Filter profile, Egress Filter profile, and an Authentication profile in a Port profile. You can either create these profiles in their respective profile pages before you create Port profiles or you can create these profiles as in-line sub-profiles while configuring Port profiles.

After you create a Port profile, you can assign it to individual interfaces or to members of a Port group. During this assignment process, you can also configure interface-specific attributes, such as IP address, and certain device-specific attributes, such as the Access profile to use for all ports on the device. You can assign only one Port profile to an interface.

[Table 99](#) describes the Quick Setup settings available in a Port profile. [Table 100](#) describes the Custom Setup settings. The defaults for these options depend on the Service Type you select.

Table 99: Port Profile Quick Setup Settings for Data Center Switching ELS

Field	Action
Profile Name	Type the name of profile by using up to 64 alphanumeric characters and no special characters other than the underscore. The name must be unique among Port profiles.
Description	Type a description of the Port profile, which will appear on Manage Port Profiles page. You can use up to 256 characters.

Table 99: Port Profile Quick Setup Settings for Data Center Switching ELS (continued)

Field	Action
Service Type	<p>Select one the preconfigured options Desktop Port, Switched Uplink, Switched Downlink, Server Port, or FCoE Transit Port. To create your own service type, select Custom.</p> <hr/> <p>Desktop Port service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—no default provided • Family Type—switching • Port Mode—access • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • Trust DHCP—disabled • MAC Limit—1 • MAC Limit Action—drop • Allowed MAC List—no default provided <hr/> <p>Switched Uplink service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—juniper_DC_Hier_Ethernet_CoS (for Hierarchical port scheduling), juniper_DC_NonHier_CoS_Fusion (for Non-Hierarchical (Fusion) port scheduling), and juniper_DC_Hier_Fusion_CoS (for Hierarchical (Fusion) port scheduling) • Family Type—switching • Port Mode—trunk • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • MAC Limit—no default provided • MAC Limit Action—no default provided • Allowed MAC List—no default provided

Table 99: Port Profile Quick Setup Settings for Data Center Switching ELS (continued)

Field	Action
	<p>Switched Downlink service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—juniper_DC_Hier_Ethernet_CoS (for Hierarchical port scheduling), juniper_DC_NonHier_CoS_Fusion (for Non-Hierarchical (Fusion) port scheduling), and juniper_DC_Hier_Fusion_CoS (for Hierarchical (Fusion) port scheduling) • Family Type—switching • Port Mode—trunk • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • MAC Limit—no default provided • MAC Limit Action—no default provided • Allowed MAC list—no default provided <hr/> <p>Server Port service type has the following default settings:</p> <ul style="list-style-type: none"> • CoS Profile—juniper_DC_Hier_Ethernet_CoS (for Hierarchical port scheduling), juniper_DC_NonHier_CoS_Fusion (for Non-Hierarchical (Fusion) port scheduling), and juniper_DC_Hier_Fusion_CoS (for Hierarchical (Fusion) port scheduling) • Family Type—switching • Port Mode—trunk • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—no default provided • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • MAC Limit—no default provided • MAC Limit Action—no default provided • Allowed MAC list—no default provided

Table 99: Port Profile Quick Setup Settings for Data Center Switching ELS (continued)

Field	Action
	<p>FCoE Transit Port service type has the following default settings:</p> <ul style="list-style-type: none"> • Port Type—Ethernet Port • CoS Profile—juniper_DC_Hier_CoS (for Hierarchical port scheduling), juniper_DC_NonHier_CoS_ELS (for Non-Hierarchical port scheduling), juniper_DC_NonHier_CoS_Fusion (for Non-Hierarchical (Fusion) port scheduling), and juniper_DC_Hier_Fusion_CoS (for Hierarchical (Fusion) port scheduling) • Family Type—switching • Port Mode—trunk • Filters—no default provided • VLAN Options—no default provided • DCBX Version—Auto • Disable DCBX—disabled • Disable Priority Flow Control—disabled • ETS No Auto Negotiation—disabled • Recommendation TVL—no default provided • Auto Negotiation—disabled • Flow Control—disabled • Maximum Size—2500 • Speed—no default provided • Link Mode—no default provided • Port Security—enabled • FCoE Trusted—enabled • MAC Limit—no default provided • MAC Limit Action—no default provided • Allowed MAC List—no default provided

Table 99: Port Profile Quick Setup Settings for Data Center Switching ELS (continued)

Field	Action
<p>Family Type: Switching or Routing</p> <p>The available settings and defaults for these options depend on the Service Type you selected.</p>	<p>This setting cannot be changed if any preconfigured Service Type was selected. If you selected the Custom Service Type, indicate whether the interface operates as a Layer 2 (Switching) or a Layer 3 (Routing) interface.</p> <p>TIP: Service Type must be set to Custom to configure a routing interface.</p> <p>If you select routing, you configure an IP address on a per-interface basis when you assign the profile to individual interfaces.</p>
<p>Port Mode for switching interfaces only</p>	<p>This setting cannot be changed if any preconfigured Service Type was selected. If you selected the Custom Service Type, select the port mode for the interface, either Access or Trunk.</p> <ul style="list-style-type: none"> • Access—Use for interfaces that connect to an end device, such as a desktop computer, an IP telephone, a wireless access point, a printer, or a security camera. The interface must belong to a single VLAN. Frames sent and received over the over the interface are untagged Ethernet frames. • Trunk—Use for interfaces that connect to a switch or router. Trunk interfaces can belong to more than one VLAN, enabling VLAN traffic to be multiplexed on a single physical interface. The Ethernet frames sent and received over the interface are tagged frames, in which IEEE 802.1Q tagging is used to segregate the traffic from each VLAN.
<p>Port Type</p>	<p>For Data Center ELS profiles, the port type is always Ethernet Port.</p>
<p>VLAN Options</p> <p>Available VLAN options depend on the Service Type selected.</p>	
<p>Member VAN (available for Switched Uplink, Switched Downlink, Server Port, FCoE Transit Port, Custom)</p>	<p>Click All if you want to assign an interface to all the VLANs.</p> <p>This option is enabled when Port Mode is Trunk.</p>

Table 99: Port Profile Quick Setup Settings for Data Center Switching ELS (continued)

Field	Action
Member VLANs (available for Desktop Port, Desktop Phone Port, Switched Uplink, Switched Downlink, Server Port, Wireless Access Port, Custom Port)	<p>Select a VLAN for the interface by clicking Select, selecting one of the listed filters, and then clicking OK. The VLAN is added to the Member VLANs list.</p> <p>Or</p> <p>Configure a VLAN by clicking Configure VLAN Settings and clicking Create. Enter the VLAN name and ID and click OK.</p>
Voice VLAN (available for Desktop Phone Port, Custom Port)	<p>Select a voice VLAN for the interface by clicking Select, selecting one of the listed filters, and then clicking OK. The VLAN is added to the Member VLANs list.</p> <p>Or</p> <p>Configure a VLAN by clicking Configure VLAN Settings and clicking Create. Enter the VLAN name and ID and click OK.</p>
Native VLAN (available for Switched Uplink, Switched Downlink)	<p>Select a native VLAN for the interface by clicking Select, selecting one of the listed VLANs, and then clicking OK. The VLAN is added to the Member VLANs list.</p> <p>Or</p> <p>Configure a VLAN by clicking Configure VLAN Settings and clicking Create. Enter the VLAN name and ID and click OK.</p>
Member VLAN	(Access ports only) Select a VLAN profile for the interface from a list of existing profiles by clicking Select .
Member VLANs	(Trunk ports only) Select a set of VLAN profiles for the interface from a list of existing profiles by using the Add and Remove functions.
Native VLAN	(Trunk ports only) Select a native VLAN profile for the interface from a list of existing profiles by clicking Select .

Table 99: Port Profile Quick Setup Settings for Data Center Switching ELS (continued)

Field	Action
<p>DCBX Settings</p> <p>Data Center Bridging Capability Exchange protocol is a discovery and exchange protocol for conveying configuration and capabilities among network neighbors to ensure consistent configuration across the network. It is an extension of the Link Layer Data Protocol (LLDP, described in IEEE 802.1AB).</p> <p>The defaults for these settings depend on the Service Type you selected.</p>	
DCBX Version	<p>Select one of the following versions of the Data Center Bridging Capability Exchange protocol:</p> <ul style="list-style-type: none"> ● Auto—automatic configuration ● DCBX v1.01—The Converged Enhanced Ethernet (CEE) version of DCBX. It has a subtype of 2 and an IEEE DCBX Organizationally Unique Identifier (OUI) of 0x001b21. ● IEEE DCBX—The newest DCBX version. Different TLVs have different subtypes (for example, the subtype for the ETS configuration TLV is 9); the OUI is 0x0080c2.
Disable DCBX	<p>Select this option to turn off Data Center Bridging Capability Exchange protocol.</p>
Disable Priority Flow Control	<p>Select this option to turn off priority flow control.</p> <p>Priority-based flow control (PFC) is a link-level flow control mechanism defined by IEEE 802.1Qbb that enables independent flow control for each class of service (as defined in the 3-bit CoS field of the Ethernet header by IEEE 802.1Q tags) to ensure that no frame loss from congestion occurs in DCB networks.</p>
ETS No Auto Negotiation	<p>Select this option to turn off ETS autonegotiation.</p> <p>Enhanced transmission selection (ETS) is a mechanism that provides finer granularity of bandwidth management within a link.</p>

Table 99: Port Profile Quick Setup Settings for Data Center Switching ELS (continued)

Field	Action
Recommendation TLV	<p>Select either Enable TLV or Disable TLV.</p> <p>The enhanced transmission selection (ETS) Recommendation TLV communicates the ETS settings that the switch wants the connected peer interface to use. If the peer interface is willing, the peer interface changes its configuration to match the configuration in the ETS Recommendation TLV. By default, the switch interfaces send the ETS Recommendation TLV to the peer. The settings communicated are the egress ETS settings defined by configuring hierarchical scheduling on the interface.</p>

After providing the information in the fields listed in [Table 96](#), click **Done**.

To use default Port profile Custom Setup settings, click **Done**. To configure Custom Setup settings, click **Custom Setup** and then provide the information in [Table 100](#) and then click **Done**.

Clicking **Done** in either case displays the dialog Do you want to assign Port Profile to Ports. Click **Yes** to create a profile assignment; else click **No** to navigate to the Manage Port Profile page and to create the Port assignment later on.

Table 100: Port Profile Custom Setup Settings for Data Center Switching ELS

Field	Action
Advanced Settings	
Expand Advanced Settings to configure link settings and port security. The Link Setting in Port profile is disabled by default. On enabling Link Settings, autonegotiation and flow control are enabled by default.	
Enable Auto Negotiation	<p>Autonegotiation of link speed and duplex mode is enabled by default; clear to disable autonegotiation.</p> <p>If you disable autonegotiation, you must set link speed and link mode.</p> <p>You cannot disable autonegotiation if a link speed of 1 Gbps is configured. This configuration might be accepted, but autonegotiation will not be disabled.</p> <p>NOTE: This setting is ignored when you assign a Port profile to an Aggregated Ethernet interface.</p>
Enable Flow Control	<p>Select to enable flow control on the interface, which permits the switch suspend packet transmission for a set period of time in response to a PAUSE frame sent by a congested switch.</p> <p>Flow control applies only to links operating at 1 Gbps, full-duplex mode.</p>

Table 100: Port Profile Custom Setup Settings for Data Center Switching ELS (continued)

Field	Action
MTU	<p>Using the arrows, indicate the maximum transmission unit (MTU), which is the maximum size of Ethernet frames sent by the interface. To calculate the MTU, add 14 bytes overhead to the maximum payload you want sent.</p> <p>Range: 256 through 9216 bytes</p>
Speed	<p>Select the link speed.</p> <p>If you select a link speed when autonegotiation is enabled, autonegotiation remains enabled and the interface will advertise the link speed that you specify as its maximum link speed.</p> <p>NOTE: This setting is ignored when you assign a Port profile to an Aggregated Ethernet interface.</p>
Link Mode	<p>Select the duplex mode, either Automatic, Full Duplex, or Half Duplex. Select Automatic to enable autonegotiation when autonegotiation is disabled.</p> <p>NOTE: This setting is ignored when you assign a Port profile to an Aggregated Ethernet interface.</p> <p>You cannot select Half Duplex with link speed set to Autonegotiation or 1 Gbps.</p>

Table 100: Port Profile Custom Setup Settings for Data Center Switching ELS (continued)

Field	Action
<p>Storm Control Settings</p> <p>Enabling storm control on a switching device monitors traffic levels and drops broadcast, multicast, and unknown unicast packets when a specified traffic level—called the storm control level or storm control bandwidth—is exceeded, thus preventing packets from proliferating and degrading the LAN.</p> <p>You can customize the storm control level for a specific interface by explicitly configuring either bandwidth or level.</p> <p>NOTE: You cannot configure both bandwidth and level for the same interface.</p>	<p>Unit</p> <ul style="list-style-type: none"> • Percentage—Configures the storm control level as a percentage of the available bandwidth used by the combined applicable traffic streams that are subject to storm control on that interface. <ul style="list-style-type: none"> The level can be set from 0% to 100%, where 0% indicates that the entire traffic is being suppressed and 100% indicates no traffic is being suppressed, in other words there is no storm control. The default level is 80%. • Kbps—Configures the storm control level as the bandwidth in kilobits per second (Kbps) of the applicable traffic streams on that interface. <ul style="list-style-type: none"> Set the bandwidth from 100 through 10,000,000 in Kbps. When you configure storm control bandwidth, the value you configure is rounded off internally to the closest multiple of 64 Kbps, and the rounded-off value represents the bandwidth that is actually used. For example, if you configure a bandwidth limit of 150 Kbps, storm control uses a bandwidth limit of 128 Kbps. <p>Value</p> <p>Configures the traffic storm control threshold level value as a percentage of bandwidth or bandwidth in kilobits per second depending upon the specified unit.</p>
<p>Power over Ethernet (PoE)</p>	
<p>You can enable PoE and display the configuration options by enabling Configure Power over Ethernet.</p>	
<p>Configure Power over Ethernet</p>	<p>Enable to configure PoE settings.</p> <p>If you do not enable this option, Network Director does not send any PoE configuration commands to the device when the profile is deployed on the device. For example, if PoE is enabled on an interface, it remains enabled. On EX Series switches, the factory-default configuration enables PoE on all interfaces that support PoE.</p> <p>If you enable this option, the PoE settings in this profile is deployed on the interfaces that support PoE. If you assign this Port profile to interfaces that do not support PoE, the profile is deployed successfully on those interfaces, but the PoE settings will not take effect.</p>

Table 100: Port Profile Custom Setup Settings for Data Center Switching ELS (continued)

Field	Action
Maximum Power (W)	<p>Use the arrows to adjust the maximum PoE power in watts allocated to a PoE port. The powered device cannot draw more power than the wattage specified. If it does so, PoE power to the port is shut down. Maximum power for PoE is 15.4W, Extended PoE is 18.6W and PoE+ is 30W.</p> <p>The Maximum Power setting has no effect when the PoE management mode for a switch or line card is class mode, which is the default mode. In class mode, the power allocated to a PoE port is determined either by LLDP negotiation with the powered device or by the PoE class of the powered device if LLDP is not supported.</p> <p>You must set the PoE power management mode for the switch or line card to static mode for the Maximum Power setting to take effect. Do this in the Device Common Settings profile.</p> <p>If you specify a maximum wattage that is greater than the maximum wattage that can be supplied by the port, your configuration is accepted when the Port profile is deployed on the port. However, the maximum wattage is set to the port's maximum supported wattage. The maximum supported wattages for PoE ports are:</p> <ul style="list-style-type: none"> • 15.4W for ports that support IEEE 802.3af only • 18.6W for IEEE 802.3af ports on switches that support enhanced PoE • 30W for ports that support IEEE 802.3at <p>Default: 15.4W</p>
Priority	<p>Select a power priority for the PoE port—either Low or High. If there is a shortage of PoE power on the switch, power to low priority ports is shut down before power to high priority ports. Within ports with the same assigned priority, power priority is determined by port number—ports with a lower port number have a higher power priority.</p> <p>Default: low priority</p>
Disable PoE	<p>Select to disable PoE on the interfaces using this Port profile.</p>

Port Security (Switching Interfaces Only)

Select to enable port security (default); clear to disable port security.

When port security is enabled, you can configure port security options such as learned MAC address limits on an interface. When port security is disabled, no port security is applied to the interface, including the default port security options.

Table 100: Port Profile Custom Setup Settings for Data Center Switching ELS (continued)

Field	Action
Trust DHCP	<p>Select to permit messages from a DHCP server to be received on the interface—this is the default. Clear to block all messages from a DHCP server from being received on the interface.</p> <p>TIP: For this port security feature to work, DHCP snooping must be enabled on the VLAN the interface belongs to. You can enable DHCP snooping on the VLAN in the VLAN profile. For directions, see “Creating and Managing VLAN Profiles” on page 502.</p>
FCoE Trusted	<p>Select to configure the interface to trust Fibre Channel over Ethernet (FCoE) traffic. If an interface is connected to another switch such as an FCoE forwarder (FCF) or a transit switch, you can configure the interface as trusted so that the interface forwards FCoE traffic from the switch to the FCoE devices without installing FIP snooping filters.</p>
MAC Limit	<p>Type the number of MAC address that can be dynamically learned on the interface.</p> <p>Range: 1 through 163839</p> <p>Default: For Desktop Ports, 1. For Desktop Phone Ports, 2. For all others, none.</p>
MAC Limit Action	<p>Select the action to be taken if the MAC address limit is exceeded:</p> <ul style="list-style-type: none"> ● Drop—Drop any packet with a previously unlearned MAC address and generate a system log entry, and SNMP trap, or an alarm. This is the default for a Desktop Port and Desktop Phone Ports. ● Log—Accept packets with new MAC addresses and learn the addresses, but generate a system log entry, and SNMP trap, or an alarm. ● Shutdown—Shut down the interface and generate a system log message, SNMP trap, or an alarm. <p>If an interface is shut down because the MAC address limit has been exceeded, you must use the CLI command clear ethernet-switching port-error interface <i>name</i> to clear the error and bring the interface back into service.</p> <p>TIP: You can use the CLI to configure auto-recovery on an interface that has been shut down by a MAC limit error.</p> <ul style="list-style-type: none"> ● None—No action. This selection effectively disables MAC address limiting on the interface. This is the default for Switched Uplink Ports, Switched Downlink Ports, and Server Ports.

Table 100: Port Profile Custom Setup Settings for Data Center Switching ELS (continued)

Field	Action
<p>Allowed MAC List</p>	<p>Indicate the MAC addresses of devices that are allowed access to the interface in the Allowed MAC List. Any device whose MAC address does not match an address in the list will not be allowed access to the interface. A list with no entries means that a client with any MAC address is permitted to access the interface.</p> <p>To enter a MAC address, click Add and then type the MAC addresses in the field provided. Enter MAC addresses as two-character hexadecimal numbers separated by colons. Click Save to save the entry.</p> <p>NOTE: Configuring an allowed MAC address list does not block the switch from receiving Layer 2 control packets—such as Link Layer Discovery Protocol (LLDP) packets—transmitted from MAC addresses that are not specified in the address list. Control packets do not undergo the MAC address check. However, the switch does not forward them to another destination.</p> <p>Default: No entries</p>
<p>RSTP Settings</p> <p>In addition to enabling or disabling the Spanning Tree Protocol (STP) as part of device profiles, this feature enables you to fine-tune STP by setting interfaces into edge, disable, or no-root-port states.</p>	<p>Edge</p> <p>RSTP defines the concept of an edge port, which is a designated port that connects to non-STP-capable devices, such as PCs, servers, routers, or hubs that are not connected to other switches. Because edge ports connect directly to end stations, they cannot create network loops and can transition to the forwarding state immediately. You can manually configure edge ports, and a switch can also detect edge ports by noting the absence of communication from the end stations.</p> <p>Disable</p> <p>Disables the RSTP on interface.</p> <p>NOTE: Configuring interfaces to one of these states is not mandatory for ELS switches. Hence, the option Disable is not applicable for ELS switches and therefore not supported.</p> <p>No Root Port</p> <p>Configures an interface to be a spanning-tree designated port. If the bridge receives superior STP bridge protocol data units (BPDUs) on a root-protected interface, that interface transitions to a root-prevented STP state (inconsistency state) and the interface is blocked. This blocking prevents a bridge from being elected the root bridge. When the bridge stops receiving superior STP BPDUs on the root-protected interface, interface traffic is no longer blocked.</p>

Table 100: Port Profile Custom Setup Settings for Data Center Switching ELS (continued)

Field	Action
CoS Settings (All except Fibre Channel Type)	<p>Click Select Cos Profile to choose from existing CoS profiles. The CoS configuration contained in the CoS profile is applied to the interfaces that the Port profile is assigned to when you deploy the configuration. Select the type of port scheduling for the CoS profile. Port scheduling depends on the device model. When you select a port scheduling type, Network Director displays the devices that support the selected port scheduling type. Click OK. Some preconfigured Service Types have a default CoS profile—see Service Types for details.</p> <p>Or</p> <p>Click Configure CoS settings to configure CoS profile. Select the type of port scheduling for the CoS profile. Port scheduling depends on the device model. When you select a port scheduling type, Network Director displays the devices that support the selected port scheduling type. See “Creating and Managing Wired CoS Profiles” on page 613 for steps to configure a CoS profile.</p>
Authentication Settings (Desktop Port, Desktop Phone Port, Custom Port)	<p>Select the Authentication profile for the interface from a list of existing profiles by clicking Select, selecting one of the listed profiles, and then clicking OK. By assigning an Authentication profile to the Port profile, you can enable 802.1x and captive portal authentication on interfaces.</p> <p>If you do not specify an Authentication profile, the interface is an open port and no authentication is required to connect.</p> <p>NOTE: You cannot configure 802.1x authentication on aggregated Ethernet interfaces. If you attempt to deploy a Port profile that contains an Authentication profile on an aggregated Ethernet interface, the deployment fails.</p> <p>Or</p> <p>Click Configure Authentication Settings to configure 802.1x and captive portal authentications. See “Creating and Managing Authentication Profiles” on page 383 for steps to configure the authentication profile.</p>
Filter Settings (available for all Service Types, including Custom for routing)	<ul style="list-style-type: none"> ● Ingress Filter Select an Ingress Filter for the interface by clicking Select, selecting one of the listed filters, and then clicking OK. ● Egress Filters Select an Egress Filter for the interface by clicking Select, selecting one of the listed filters, and then clicking OK.

Table 100: Port Profile Custom Setup Settings for Data Center Switching ELS (continued)

Field	Action
VRRP Settings (available when Service Type is Custom and Family Type is Routing)	Select the VRRP profile for the interface from a list of existing profiles by clicking Select . Select one of the listed profiles, and then click OK .

Clicking **Done** displays the dialog Do you want to assign Port Profile to Ports. click **Yes** to create a profile assignment else click **No** to navigate to the Manage Port Profile page and to create the Port assignment later.

What to Do Next

After you create a Port profile, you can assign it to interfaces or members of port groups. During this process, you can also configure interface-specific attributes, such as IP address, and certain device-specific attributes, such as which Access profile to use for all ports on the device. For more information, see [“Assigning and Unassigning Port Profiles from Interfaces” on page 476](#).

RELATED DOCUMENTATION

[Understanding Port Profiles | 408](#)

[Assigning and Unassigning Port Profiles from Interfaces | 476](#)

[Creating and Managing Port Groups | 495](#)

[Creating and Managing VLAN Profiles | 502](#)

[Creating and Managing Authentication Profiles | 383](#)

[Creating and Managing Wired CoS Profiles | 613](#)

[Assigning and Unassigning Port Profiles from Interfaces | 476](#)

[Understanding VRRP Profiles | 845](#)

[Creating and Managing VRRP Profiles | 846](#)

[Network Director Documentation home page](#)

Assigning and Unassigning Port Profiles from Interfaces

IN THIS SECTION

- [Selecting Devices for Assignment | 477](#)
- [Selecting Interfaces for Assignment | 478](#)
- [Reviewing and Accepting the Assignments | 481](#)
- [Editing Profile Assignments | 482](#)
- [Unassigning a Port Profile from an Interface | 483](#)

You can assign an existing user-created or system-created Port profile to network interfaces (including aggregated Ethernet interfaces), or Port Group member interfaces on one or more devices.


During the process of assigning a Port profile to interfaces, you can also:

- Configure IPv4 or IPv6 addresses on interfaces to which you have assigned a routing Port profile.

TIP: IPv4 filters are separate from IPv6 filters.

- Configure certain authentication attributes—such as the RADIUS server or servers to use—for all 802.1X interfaces on the device. Because configuring these attributes involves assigning an Access profile to the device, you must have previously created an Access profile.

To assign a Port profile to interfaces:

1. Click  in the Network Director banner.
2. Under Select View, select one of the following views: **Logical View**, **Location View**, **Device View** or **Custom Group**.

TIP: Do not select **Datacenter View** or **Topology View**.

3. In the Tasks pane, select **Wired > Profiles > Port**.

The Manage Port Profile page is displayed.

4. Select the Port profile you want to assign and then click **Assign**.

The Assign Port Profile wizard appears. It has three parts—Device Selection, Profile Assignment, and Review.

5. Complete device selection for assignment by following the directions [“Selecting Devices for Assignment” on page 477](#).
6. Assign the port profile to one or more objects by following the directions [“Selecting Interfaces for Assignment” on page 478](#).
7. Review your configuration by following the directions [“Reviewing and Accepting the Assignments” on page 481](#).
8. Click **Finish**.

After you assign a Port profile to ports, you can modify your assignments by selecting the Port profile from the Manage Port Profiles page and clicking **Edit Assignments**.

The following sections describe how to use the Assign Port Profile wizard and the Edit Assignments page.

Selecting Devices for Assignment

Use the Device Selection page in the Assign Port Profile wizard to select one or more devices that have ports. You can select container nodes, individual devices, or port groups. For more information about Port Groups, see [“Creating and Managing Port Groups” on page 495](#).

To select devices for Port profile assignment:

1. Enable either **Select Devices** or **Select Port Groups**.
2. If you enabled **Select Devices**, expand the list of objects and select the objects that contain the devices and interfaces you want to assign by clicking the check box next to the them. If you select a container node, all devices under that node are selected.

TIP: The list of objects is filtered to include only devices that match the profile’s family type. If you do not see a device that you expected to see, verify that the device matches the profile’s family type. For example, a profile with the Data Center Switching ELS family type cannot be assigned to a Data Center Non-ELS device.

3. If you enabled **Select Port Groups**, select one or more port groups from the Select Port Group list.
4. Click either **Next** or **Profile Assignment** to proceed to the next step in the wizard, Profile Assignment.

For directions to complete Port Profile Assignment, see [“Selecting Interfaces for Assignment” on page 478](#).

Selecting Interfaces for Assignment

Use Profile Assignment in the Assign Port Profile wizard to select the interfaces to which you want to assign the Port profile. After you have selected the interfaces, you can configure specific attributes on the interfaces or on the devices to which the interfaces belong.

TIP: Before you start the procedure below, you might want to select a device and click **View Assignments** to view what profiles and attributes are already configured on the device. Any profile assignments or attributes you define during this procedure replace the existing ones.

One optional attribute you can configure for switching interfaces is the Access profile that defines RADIUS server authentication for 802.1X ports. If you will be configuring this optional attribute, make sure that an Access profile has been created.

If you enabled **Select Port Groups** during Object Selection, you can assign the Port profile to any or all existing port groups.

If you enabled **Select Devices** during Object Selection, assign the Port profile to interfaces and configure the port-specific or device-specific attributes:

1. Select one or more container nodes or devices from the Assignments list:

- To assign the profile to nonconsecutive interfaces or to aggregated Ethernet interfaces, select a single device.
- To assign the profile to interfaces in the same consecutive interface range (for example, ge-0/0/0 through ge-0/0/15) on one or more devices, select one or more devices. To make multiple selections, press Shift or Ctrl while making the selections.
- To assign a profile to ports within a QFabric system select the member node group or groups that contain the ports.
- To assign a profile to ports within a Virtual Chassis Fabric (VCF), you can select any container nodes or member devices within the VCF, including the VCF container node.
- To assign a profile to aggregated Ethernet ports within a Virtual Chassis or VCF, select the Virtual Chassis or VCF container node. To assign a profile to physical device ports within a Virtual Chassis or VCF, select one or more member devices.
- Channelized ports are only applicable for Data Center Switching ELS devices and only XE interfaces can be used as channelized ports.

NOTE: If Network Director fails to read the configuration of one or more devices after the device discovery, such devices are not displayed in the Assignments list. You will not be able to assign profiles to such devices. The Manage Jobs page in System mode displays details of the device discovery jobs. Use the information displayed on this page to take appropriate corrective steps to enable Network Director to reread the configuration of the failed device. For more information, see [“Discovering Devices in a Physical Network” on page 204](#).

2. Click **Assign to Port**.

The Assign Profile to Ports window opens.

3. Select either **Ports** (default) or **Port Range**. If you selected multiple devices in the previous step, you cannot choose the Port option.

4. If you selected the Port option, select the ports from the list of ports.

By default, aggregated Ethernet interfaces are listed after the ge- and xe- interfaces in the list of ports. Members of aggregated Ethernet interfaces are not included in the port list.

5. If you selected the Port Range option, enter the port range:
 - a. In the Normal Ports section, enter a first and last port name in the text boxes, then click **Add**. The port range appears in the Selected Port Range section.
 - b. Repeat the add process to add any additional port ranges.
 - c. To delete a port range, select its check box, then click **Delete**.

At least one port within the port range must be available on each selected device for the port range to succeed. Channelized ports are supported in a port range. Assignments are created for the ports within the port range that are available. You can assign the profile to the same interface on multiple devices by entering the interface name in both fields of the port range.

6. Click **Assign** to complete the port assignments and close the window.

The port assignment appears in the list of Assignments, with the Device, Type, Assigned To, and Attributes columns completed. In the Attributes column, you see a triangle and the link **Define**.

7. Configure the following port-specific or device-specific attributes:

- If the Port profile is a switching profile that contains an Authentication profile—in other words, the profile is enabling 802.1X authentication on ports—click the **Define** link in the Attributes column for a device to define additional authentication attributes.

The Configure attributes window opens. Fill in the fields described in [Table 101](#).

Table 101: Configure Device Attributes for Port Profile Assignments

Field	Action
Access Profile	Select an Access profile. The RADIUS server attributes defined in the Access profile is configured on the device when you deploy the configuration.
Radius Server Source IP Address	Type an IP address to be used as the source IP address for RADIUS server requests sent by the switch. The source address must a valid IPv4 or IPv6 (either format) address configured on one of the switch interfaces.
Post authentication URL	Type a URL to be used for the captive portal post-authentication website.

TIP: If you see the message *Port profile does not have an associated Authentication profile.* Please configure the Authentication profile.; then click **OK**, and edit the Port profile by selecting **Port** under Profile and Configuration Management, selecting the Port profile from the list and clicking **Edit**. The Authentication profile association is located in the Port Family Options section.

NOTE: The attributes you define for the device apply to all 802.1X authenticator interfaces on the switch. Different sets of interfaces on the switch cannot have different attributes.

- If the Port profile is a routing profile, click the **Define** link in the port's Attributes field to configure an IPv4 or IPv6 address on the interface.

Repeat this step for all the ports on which you want to configure IPv4 or IPv6 interfaces.

8. Repeat the previous steps as needed to complete the port assignments and then click either **Next** or **Review**.

For review directions, see [“Reviewing and Accepting the Assignments” on page 481](#).

Reviewing and Accepting the Assignments

Use the Review step of the Assign Port Profile wizard to review and accept your assignments:

- Click **Edit** to return to the Profile Assignment step and make changes to your assignments.
- Click **Finish** to accept the assignments.

After you click Finish, the Create Profile Assignments Job Details window opens, which reports on the status of the profile assignment job. If you have assigned the profile to a large number of objects, the profile assignment job can take some time to complete. Instead of waiting for the Job Details window to report job completion status, you can close it and check the details of the profile assignment job at a later time by using the Manage Job task in System mode.

NOTE: If any assignment fails, the profile assignment job fails and none of the assignments are created. Check the details for the profile assignment job for information about why the assignment failed.

After the profile assignment job completes, you can deploy the configuration defined in the Port profile and in the port-specific and device-specific attributes on the affected devices. See [“Deploying Configuration to Devices” on page 1180](#).

Editing Profile Assignments

Use the Edit Assignments page to change Port profile assignments. You can:

- Delete a port from the profile assignments.

If the profile has been already deployed on the port, then the configuration is removed from the port when you next deploy the configuration on the device. The configuration removed includes any port configuration that was defined in associated profiles, such as the CoS, Authentication, and IPv4 or IPv6 Filter profiles.

- Change the IPv4 or IPv6 address for ports associated with a routing Port profile.
- Change the device-specific authentication attributes, such as the Access profile associated with the device. For more information about these attributes, see [Table 101](#).

NOTE: You cannot assign the Port profile to additional ports by using the Edit Assignment page. To add port assignments, use the Assign Port Profile wizard.

[Table 102](#) describes the fields in the Edit Assignments page and how to use them to change the profile assignments. When you are finished with your modifications, click **Apply**. You can then deploy your modifications on the affected devices.

Table 102: Edit Assignment for Port Profile Fields

Field	Description
Objects	Expand the device nodes to see the ports or port group the profile is assigned to.
Assigned State	Indicates the current state of profile assignment on the port: <ul style="list-style-type: none"> • Deployed—The profile configuration has already been deployed on the port. • Pending—The profile configuration has not yet been deployed on the port. • Pending Removal—The profile configuration was deployed on this port, but will be removed from the port the next time the device configuration is deployed.
Attributes	If the attributes for a port or device are currently undefined, you can click the Define link to define them. If attributes have been defined and you want to view them or change them, click the Change link.
Operation	Click the Delete link to delete the profile assignment from the port.

Table 102: Edit Assignment for Port Profile Fields (continued)

Field	Description
Record Status	<p>Shows the current assignment status:</p> <ul style="list-style-type: none"> • An X indicates that you have marked a port for deletion. • A pencil indicates that you have modified the associated attribute. <p>After you apply your assignment changes, these indicators disappear.</p>

Unassigning a Port Profile from an Interface

Starting Network Director Release 3.5, you can unassign multiple port profiles that are assigned to multiple ports, at the same time.

To unassign port profiles:

1. On the Network Director banner, under **Views**, select one of the following views—Logical View, Location View, Device View, or Custom Group.

2. On the Tasks pane, click **Wired > Profiles > Port**.

The Manage Port Profile page appears.

3. Select one or more port profiles that you want to unassign from the ports and click **Unassign**.

A confirmation message indicating the profiles were successfully unassigned appears and the status of the profiles change to Pending Deployment.

RELATED DOCUMENTATION

[Creating and Managing Port Profiles | 414](#)

[Creating and Managing Access Profiles | 352](#)

[Creating and Managing Port Groups | 495](#)

[Network Director Documentation home page](#)

Managing Auto Assignment Policies

To support rapid network deployment, Network Director enables you to define your network configuration in a set of profiles that you can apply to multiple objects in your network. Auto assignment policies go one step ahead and further automate profile assignment. When Network Director detects the devices included in a policy, in a campus network, the Port profiles that you have created in Network Director are automatically assigned to various switch ports on supported devices.

Network Director uses LLDP to detect the type of network device. When the devices such as Desktop, Desktop Phone, Server Port, Wireless Access Port, and Printer are LLDP enabled, Network Director triggers auto assignment of Port profiles for these devices.


NOTE: Network Director detects most of the printers by using organizationally unique identifier (OUI). For more information, see [“Adding and Managing OUI Data in Network Director” on page 271](#).

To create an auto assignment policy, you specify one or more Port profiles, devices, ports or port ranges on the devices to which the Port profile is to be deployed, and a few additional parameters.

After you create an auto assignment policy, when any of the device ports that you specified in the auto assignment policy are connected to a Desktop, Desktop Phone, Server Port, Wireless Access Port, or Printer, Network Director performs the following tasks:

- a. Initiates a job to update the Port profile configuration on the connected ports. If you enable a policy, Network Director overwrites the configuration that is already deployed on the ports and deploys the configuration from the profile that you specified in the auto assignment policy. You can view details of this job in the Job Management page and also from the Policy Assignment Log window.
- b. Updates the port associations and displays the results in the Manage Port Profiles page. In the Manage Port Profiles page, the Port profiles that are assigned through an auto assignment policy are highlighted with a ★ next to the profile name. You can view more details about the auto assignment in the Port Profile Details window.

To manage auto assignment policies:

1. Click  in the Network Director banner.
2. Under Views, select one of the following views: **Logical View**, **Location View**, **Device View**, or **Custom Group**.

TIP: Auto assignment is not available in **Datacenter View** or **Topology View**.

3. Click **Wired > Tasks > Manage Auto Assignments** under the Tasks pane.

The Manage Auto Assignment Policy page opens.

4. From the Manage Auto Assignment Policy page, you can:

- Create a new policy by clicking **Create**. For details, see [“Creating Auto Assignments” on page 486](#).
- Modify an existing policy by selecting the policy and clicking **Edit**.
- View information about a policy by selecting the policy and clicking **Details**. Network Director opens the Auto Assignment Policy Details page.
- Delete a policy by selecting the policy and clicking **Delete**.
- Deploy a policy by selecting the policy and clicking **Run Now**.

[Table 103](#) describes the information provided about the policies on the Manage Auto Assignment Policy page. This page lists all auto assignment policies defined for your network, regardless of your current selected scope in the network view.

Table 103: Manage Auto Assignment Policy Page

Column	Description
Name	Unique name given to the auto assignment policy when the policy was created.
Description	Description of the policy that was entered when the policy was created.
Device Family Mode	Displays one of the following: <ul style="list-style-type: none"> • EX—for EX Series switches • ELS—for Campus Switching ELS
Log	Displays link details corresponding to each policy. Click Details corresponding to each policy. The Policy Assignment Log window opens listing the internal log of a policy.
Creation Time	Date and time when the policy was created.
Update Time	Date and time when the policy was last modified.

TIP: All columns might not be currently displayed. To show or hide fields in the table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

RELATED DOCUMENTATION

| [Creating Auto Assignments](#) | 486

Creating Auto Assignments

IN THIS SECTION

- [Adding Port Profiles using the Select Port Profiles Page](#) | 487
- [Adding Devices and Ports for Auto Assignment](#) | 487
- [Viewing the Auto Assignment Policy Summary](#) | 488

Auto assignments help in automating Port profile assignments. Auto assignments automatically assign Port profiles that you have defined, to various switch ports based on the endpoints detected on the port on supported devices when Network Director detects the devices included in the auto-profile in the network. Auto assignments are supported for desktop ports, desktop phone ports, server ports, wireless access ports, and printer ports.

You can create auto assignment policies using the Create Auto Assignment wizard. The Create Auto Assignment wizard consists of three pages—Select Port Profile, Select Devices, and Summary.

To open the Create Auto Assignment wizard, click **Create** in the Manage Auto Assignment Policy page. Perform the following tasks to create an auto assignment:

Adding Port Profiles using the Select Port Profiles Page

To create an auto assignment policy and add Port profiles:

1. Enter a name and a description for the auto assignment.
2. Select the **Enable Policy** check box to enable the policy for the auto assignment.
3. Select a device and specify a port range for this auto-assignment in the next wizard page.
4. Select a device family. Auto assignment supports devices that belong to *Switching (EX)* or *Campus Switching ELS*.
5. Click **Select** in the Port Profiles table to add Port profiles for auto assignment. Network Director opens the Select Profiles window and displays only those Port profiles that are created for the device family that you selected. Select one or more Port profiles that you want to add to the auto assignment and click **Add**.

NOTE: If there are multiple profiles of the same service type for a device you can select only one of the profiles for auto assignment. This ensures that you can assign only one service type profile across all auto assignment policies. For example, if the Select Profiles window lists two Port profiles with Server as the service type, then you can select only one of these two profiles for auto assignment.

Network Director adds the selected Port profiles to the Port Profiles table. To remove a Port profile from the list, select a profile and click **Remove**.

6. Click **Next** to add devices for auto assignment.

The Select Device page opens.

Adding Devices and Ports for Auto Assignment

To add devices and specify ports for auto assignment:

1. Click **Select Devices**. The Select Devices window opens.
2. Expand the list of objects and select the objects that contain the devices you want to add by clicking the check box next to them. If you select a container node, all devices under that node are selected.

NOTE: The list of objects is filtered to include only devices that match the Port profile family type. If you do not see a device that you expected to see, verify that the device matches the profile's family type. For example, a Port profile created for Switching (EX) family type cannot be assigned to an auto assignment policy that you are creating for Campus Switching ELS.

3. Click **OK** to add the selected devices to the Select Device(s) table.
4. Select a device from the Select Device(s) table and click **Configure Range** to specify the ports or port ranges to which you want to auto assign the Port profiles.

The Configure Port Inclusion window opens.

NOTE: Network Director deploys the Port profiles only if you specify a port range for the auto assignment policy and when the end points match.

5. Select the starting and ending port numbers and click **Add** to add the port range to the Selected Port Range table. You can add multiple port ranges that do not overlap with each other, for each device across policies.
6. When you have added all the required port ranges, click **OK**.
7. Repeat steps 4 through 6 to add port ranges for all the devices that you have added to the Select Devices table.
8. To remove a device from the Select Devices table, select the device and click **Remove**.
9. Click **Next** to view a summary of the auto assignment policy.

Viewing the Auto Assignment Policy Summary

The Summary page displays the details of the auto assignment policy. You can review and make modifications to the auto assignment policy. To modify the configuration details, click the appropriate buttons in the Auto Assignment Policy workflow at the top of the page that corresponds to the configuration you want to change.

When you are finished with your modifications, click **Summary** to return to this page.

To complete the auto assignment policy creation, click **Finish**.

Network Director displays the policy that you created in the Manage Auto Assignment Policy page.

When any of the device ports that you specified in the auto assignment policy are connected to a Desktop, Desktop Phone, Server Port, Wireless Access Port, and Printer, Network Director automatically assigns Port profiles to these device ports. You can view the profile deployed for a port in the Policy Assignment Log window for each policy.

RELATED DOCUMENTATION

| [Managing Auto Assignment Policies | 484](#)

Configuring Easy Config Setup

IN THIS SECTION

- [Configuring Interface Settings | 489](#)

In addition to the Port profile configuration, Network Director enables users to quickly configure interfaces on devices by using the Easy Config Setup task. You can perform configurations by directly selecting the device, instead of creating a new profile and assigning a profile to the device port. You can also deploy the configuration changes without creating additional profiles which results in growing number of profiles in Port profile configurations. Easy Config Setup is supported only for the configurations that are automatically approved; for configurations that require manual approvals, this task is disabled.

Configuring Interface Settings


This section describes the steps to configure the interface settings by using Easy Config Setup.

You can configure the following interface settings in the EX Switching, Campus Switching ELS (MX series devices are supported in L2NG mode only and are not supported in native MX series mode), Data Center Non-ELS, and Data Center ELS devices..

- Port Settings
- VLAN Settings
- PoE Settings

- 802.1x Settings
- Access Settings

To configure an interface in a device by using the Easy Config Setup:

1. Select a switching device from the left navigation pane.
2. Click  in the Network Director banner.
3. Under Select View, select either **Logical View**, **Location View**, or **Device View**.
4. In the Tasks pane, click **Wired > Tasks > Easy Config Setup**.

NOTE: This task is not visible for a fabric device.

5. Enter the settings for the interface described in [Table 104](#).

Table 104: Easy Config Setup Settings

Field	Action
Port Settings	
Ports/Interfaces	Select an Ethernet switching interface, an IPv4 routing interface, or an IPv6 routing interface. All the ports associated with the device (except Layer 3 interfaces) are available in the list.
Description(optional)	Provide a description of the device configuration or port details of the device. You can use up to 256 characters.
Port Mode	Configure a switching interface port to be an access, trunk, or tagged-access port for EX Series switches. Campus Switching ELS, Data Center Switching ELS, and Data Center Switching non-ELS series devices supports access mode and trunk mode. For more information about port modes, see Creating and Managing Port Profiles .
Disable Port	Disables the port. You can still configure and deploy all the settings but these settings become active only when you enable the port by clearing this selection.

Table 104: Easy Config Setup Settings (continued)

Field	Action
<p>Member VLAN Settings</p> <p>You can enable VLAN settings and display the configuration options by enabling Member VLAN Settings.</p>	<p>You can either select an existing VLAN profile or create a new VLAN profile that you want to assign to the port.</p> <p>To select an existing profile:</p> <ol style="list-style-type: none"> Select the option Select VLAN Profile. Select the option Select. The Choose VLAN profile window opens. Select the VLAN profile name and click OK. <p>To create a new profile:</p> <ol style="list-style-type: none"> Select Configure VLAN Settings. Click Create. The Create VLAN Profile window opens. Enter the VLAN name. Under VLAN ID, select Single and enter a VLAN ID from 1 to 4094 if you want to configure a single VLAN. or Under VLAN ID, select Range and enter a range of VLAN IDs that you want to assign to the VLAN profile. TIP: Single VLAN IDs can be configured for all products. VLAN lists or VLAN ID ranges are available for some products, depending on the technology used for implementation. Click OK.
<p>PoE Settings</p> <p>You can enable PoE and display the configuration options by enabling PoE Settings.</p>	
<p>Maximum Power(W)</p>	<p>Use the arrows to adjust the maximum PoE power in watts allocated to a PoE port. The powered device cannot draw more power than the wattage specified. If it does, PoE power to the port is shut down.</p>

Table 104: Easy Config Setup Settings (continued)

Field	Action
Priority	<p>Select a power priority for the PoE port—either Low or High. If there is a shortage of PoE power on the switch, power to low priority ports is shut down before power to high priority ports. Within ports with the same assigned priority, power priority is determined by the port number—ports with a lower port number have a higher power priority.</p> <p>Default: low priority</p>
Disable PoE	Select to disable PoE on the interfaces that use this Port profile.

802.1x Settings (Authentication)

You can configure 802.1x and display the configuration options by enabling **802.1x Settings (Authentication)**.

Enable 802.1x	<p>802.1x authentication is enabled by default for a switching profile. 802.1x authentication works by using an Authenticator Port Access Entity (the switch) to block all traffic to and from a supplicant (end device) at the port until the supplicant's credentials are presented and matched on the Authentication server (a RADIUS server). When authenticated, the switch stops blocking traffic and opens the port to the supplicant. Network access can be further defined using VLANs.</p>
Enable MAC-RADIUS	<p>Select to enable MAC-RADIUS based authentication for this profile. MAC RADIUS authentication enables LAN access to permitted MAC addresses. When a new MAC address appears on an interface, the switch consults the RADIUS server to check whether the MAC address is a permitted address. If the MAC address is configured on the RADIUS server, the device is allowed access to the LAN.</p>
Supplicant Mode	<p>Specify the mode authentication supplicants use, either Single, Multiple, or Single-Secure.</p> <ul style="list-style-type: none"> • Single—Allows only one host for authentication. This is the default mode. • Single-Secure—Allows only one end device to connect to the port. No other end device is enabled to connect until the first logs out. • Multiple—Allows multiple hosts for authentication. Each host is checked before being admitted to the network.
Guest VLAN	<p>Click Select and then select the VLAN to which an interface is moved when no 802.1x supplicants are connected on the interface. The VLAN specified must already exist on the switch.</p>

Table 104: Easy Config Setup Settings (continued)

Field	Action
Reject VLAN	Click Select and then select the VLAN to which an interface is moved when the switch receives an Extensible Authentication Protocol over LAN (EAPoL) Access-Reject message during the authentication process between the switch and the RADIUS authentication server.
Server Fail Type	<p>Specify the server fail fallback action the switch takes when all RADIUS authentication servers are unreachable; one of None, Deny, Permit, Use cache, or VLAN Name.</p> <ul style="list-style-type: none"> • None—No server fallback action is used. This option is selected by default. • Deny—Force fails supplicant authentication. No traffic will flow through the interface. • Permit—Force succeeds the supplicant authentication. Traffic will flow through the interface as if it were successfully authenticated by the RADIUS server. • Use cache—Force succeeds the supplicant authentication only if it was previously authenticated successfully. This action ensures that already authenticated supplicants are not affected. • VLAN Name—Move supplicant on the interface to the VLAN specified by this name. This action is allowed only for the first supplicant connecting to an interface. If an authenticated supplicant is already connected, then the supplicant is not moved to the VLAN and is not authenticated. If you select this option, you must provide a Fail VLAN name.

Access Settings

You can configure authentication parameters and accounting parameters on the network and display the configuration options by enabling **Access Settings**.

Server Address	Enter the IP address of the RADIUS server.
Authentication Port	The default RADIUS authentication port is 1812. You can change the port number by using the up and down arrows.
Secret	Provide a password. If the password contains spaces, enclose it in quotation marks. The secret password used by the switch must match the one used by the server.
Retry Count	Specify the number of times that a device attempts to contact the LDAP authentication server. The default retry count is 3. You can change this value by using the up and down arrows to 1 through 100 times.

Table 104: Easy Config Setup Settings (continued)

Field	Action
Timeout (seconds)	Specify the number of seconds the switch waits to receive a response from a RADIUS server. The default timeout is 5 seconds. You can change this value, using the up and down arrows, to 1 through 65535 seconds.

6. Click **Preview** to view the configuration changes that will be deployed to a device when a job runs. Use the Configuration window to see changes that were deployed to a device when a completed job ran. The configuration changes are shown in these formats:

- Select the CLI View tab to view the configuration changes in CLI format. This view shows the Junos configuration statements that will be deployed to the device.
- Select the XML View tab to view the configuration changes in XML format. This view shows the XML-formatted configuration that will be deployed to the device by using the Device Management Interface (DMI), which is used to remotely manage devices.

7. Click **Deploy** to deploy configuration to a device.

After you deploy the configuration, the device goes out of sync and Network Director triggers auto-resynchronization of the device. If there is any conflicting configuration, the new port profile (created during easy config setup) is prompted during the assignment.

The Deploy EasyPortal Configuration window opens. If you chose to deploy the changes immediately, the Deployment Status column shows the status as **INPROGRESS** and changes to **SUCCESS** after the deployment is successfully completed.

Click **Cancel** if you want to cancel your changes.

8. Click **Close** to close the deployment page.

NOTE: Clicking either **Close** or **Cancel** takes you to the Device Inventory My Network page, which displays the details of the device you selected in the first step.

Understanding Port Groups

Ports on virtual and physical devices regulate data-packet traffic to both ensure security and a guaranteed rate of packet flow, and prevent unsolicited traffic. Since configuring each port individually would be tedious, especially if the ports are configured with the same settings, port groups enable configuration of multiple ports simultaneously. First you create the port group (see [“Creating and Managing Port Groups” on page 495](#), and then you can assign a Port profile to the Port Group—see [“Assigning and Unassigning Port Profiles from Interfaces” on page 476](#).

In the case of wireless ports in a group, you can assign a VLAN profile to the port group—see [“Assigning a VLAN Profile to Devices or Ports” on page 531](#).

The ports in a port group can be located on any of your switches or controllers and can include ports from different devices and from different series. Group port configuration has precedence over any individual port configuration.

RELATED DOCUMENTATION

[Creating and Managing Port Groups | 495](#)

[Assigning and Unassigning Port Profiles from Interfaces | 476](#)

[Assigning a VLAN Profile to Devices or Ports | 531](#)

[Network Director Documentation home page](#)

Creating and Managing Port Groups

IN THIS SECTION

- [Managing Port Groups | 496](#)
- [Creating Port Groups | 497](#)
- [Specifying Settings for a Port Group | 498](#)
- [What to Do Next | 499](#)

From Network Director, you can group ports and then name that port group. The ports can be located on any of your switches or controllers and can include ports from different devices and from different series. Creating Port groups enables simultaneous multiple port configuration. For example, when a Port profile

(see [“Creating and Managing Port Profiles” on page 414](#)) is assigned to the members of a Port group (see [“Assigning and Unassigning Port Profiles from Interfaces” on page 476](#)), all ports in the group are configured with the Port profile.

In the case of wireless ports in a group, you can assign a VLAN profile to the port group—see [“Assigning a VLAN Profile to Devices or Ports” on page 531](#).

NOTE: Configuration applied to members of a port group has precedence over any individual port configuration.

This topic describes:

Managing Port Groups

Use the Manage Port Groups page to manage existing Port groups and to create new ones. Port groups enable simultaneous multiple port configuration.

From the Manage Port Groups page, you can:

- Create a new Port Group by clicking **Add**. For directions, see [“Creating Port Groups” on page 497](#).
- Modify an existing Port Group by selecting it and clicking **Edit**.
- View information about a Port Group by selecting the group and clicking **Details** or by clicking the group name.
- Delete a Port Group by selecting a group and then clicking **Delete**.

TIP: To see the current assignments for a group, click the group name.

[Table 105](#) describes the information provided about Port Groups on the Manage Port Groups page. This page lists all Port Groups defined for your network, regardless of your current selected scope in the network view.

Table 105: Manage Port Groups Information

Field	Description
Name	Name given to the group when the group was created. Click the group name to view group details.
Last Updated	Date the port group was last altered.


Table 105: Manage Port Groups Information (*continued*)

Field	Description
User	User name of the person who last altered the Port Group.

Creating Port Groups

TIP: The required configurations for a Port Group are a Port Group name and the configuration of at least one device port.

To create a Port group for switches:

1. Select one of the following device views in the Network Director banner:
 - **Logical View**—Displays devices in hierarchal groupings based on logical relationships.
 - **Location View**—Displays devices in hierarchal groupings based on physical locations.
 - **Device View**—Displays devices in hierarchal groupings based on device type.
 - **Custom Group View**—Displays devices in hierarchal groupings based on custom group.
2. Click  in the Network Director banner.
3. Under Select View, select either **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Datacenter View** or **Topology View**.

4. Click **Manage Port Groups** under Device Management in the Tasks pane.
The Manage Port Groups page appears.
5. Click **Add**.
The Create Port Group page appears.
6. Enter settings for the Port Group as described in [“Specifying Settings for a Port Group” on page 498](#).
7. Click **Done**.

The message *Port group successfully created* is displayed.

8. Click **OK**.

The new port group appears on the list of managed port groups.

Specifying Settings for a Port Group

Use the Create Port Group page to define the members of a Port Group.

[Table 106](#) describes the settings available in the Port Group.

Table 106: Port Group Settings

Field	Action
Port Group Name	Type the name of the Port Group, using up to 64 alphanumeric characters and no special characters other than the underscore. The name must be unique among Port groups.
Port Group Description	Type a description of the Port Group, which will appear on Manage Port Groups page. You can type up to 256 characters.

Select Devices and Ports

Add Ports to the Port Group (task)	<p>To add ports to this Port Group:</p> <ol style="list-style-type: none"> 1. Select a device with ports from the tree displayed in the column labelled Select Device and Ports. 2. For the current selection (highlighted device), click either All Ports or Selected Ports. If you click All Ports, the ports or port range is immediately listed under Selected Ports and Port Ranges. If you click Selected Ports, you must then select individual ports from the Port Selector list, and click Done. 3. Click Done. An information window displays the message <i>Port group created successfully</i>. 4. Click OK to close the information window. The new Port Group is now listed under Manage Port Groups. <p>TIP: To edit a port group, select it from the Manage Port Groups list and then click Edit. To remove a port group, select it and then click Delete.</p>
------------------------------------	--

What to Do Next

After you create a Port Group, you can treat it like an individual port—see [“Assigning and Unassigning Port Profiles from Interfaces” on page 476](#). For wireless ports in a group, you can assign a VLAN profile to all members of the group—see [“Assigning a VLAN Profile to Devices or Ports” on page 531](#).

RELATED DOCUMENTATION

[Creating and Managing Port Profiles | 414](#)

[Assigning and Unassigning Port Profiles from Interfaces | 476](#)

[Assigning a VLAN Profile to Devices or Ports | 531](#)

[Understanding Port Groups | 495](#)

[Network Director Documentation home page](#)

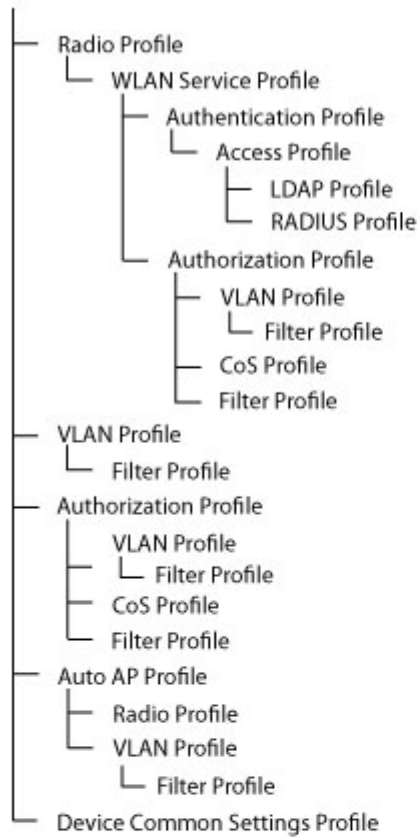
Understanding VLAN Profiles

A virtual LAN (VLAN) is a Layer 2 broadcast domain that can span multiple wired or wireless LAN segments. Each VLAN is a separate logical network, grouping hosts with common requirements, regardless of their physical location.

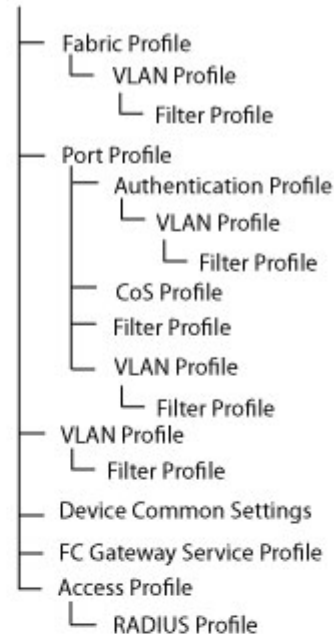
VLAN profiles in Network Director enable multiple VLAN configuration from a single profile. Each VLAN profile is specific to a device family: EX Switching, Wireless, Campus Switching ELS, or Data Center Switching Non-ELS. In addition, the VLANs are created for different purposes at different levels, as shown in [Figure 22](#).

Figure 22: VLANs Are Specific to Device Families and Function Levels

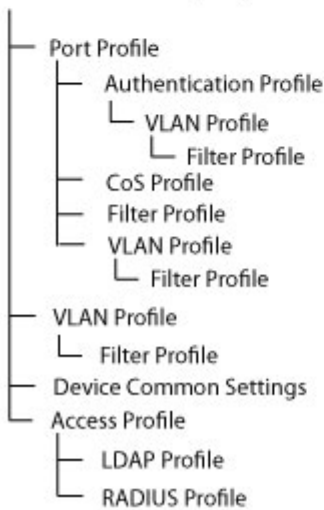
Wireless LAN Device



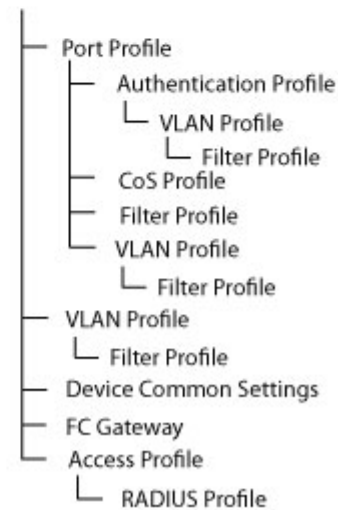
QFX Series Switch



EX Series Switch (ELS)



EX Series Switch



Note: An Access profile is assigned to a switch when you assign a Port profile to the switch interfaces.

After a VLAN profile is created, you assign it to either wireless ports, switches, controllers, custom groups, port groups with wireless ports, or access points that are managed by controllers or clusters. You can also assign VLAN profiles to controller-managed access points and cluster-managed access points for local switching on the access points.

For EX Series Switches and Campus Switching ELS, apart from the basic settings, you can specify:

- MAC parameters
- Switching and routing parameters
- L2 and L3 Filters
- VLAN security DHCP, ARP inspection, and MAC movement.

For controllers, apart from the basic settings, you can specify:

- Filter details
- Routing
- DHCP enabled
- IGMP settings
- Spanning tree protocol settings

For data center switching devices, apart from the basic settings, you can specify:

- MAC parameters
- Switching and routing parameters
- Filters
- VLAN security settings
- Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) settings

TIP: Single VLAN IDs can be configured for all products and circumstances. VLAN lists or VLAN ranges of IDs are available for some products, depending on the technology used for implementation. For example, EX Switching does not support a VLAN list. Campus Switching ELS supports a VLAN ID range only as part of a VLAN ID list. You will only see the available configurations for the selected device family.

RELATED DOCUMENTATION

[Creating and Managing VLAN Profiles | 502](#)

[Assigning a VLAN Profile to Devices or Ports | 531](#)

[Understanding Network Configuration Profiles | 197](#)

[Network Director Documentation home page](#)

Creating and Managing VLAN Profiles

IN THIS SECTION

- [Managing VLAN Profiles | 503](#)
- [Creating a VLAN Profile | 504](#)
- [Specifying Basic EX Switching VLAN Settings | 506](#)
- [Specifying Basic Wireless VLAN Settings | 507](#)
- [Specifying Basic Campus Switching ELS VLAN Settings | 508](#)
- [Specifying Basic VLAN Settings for Data Center Switching Non-ELS | 510](#)
- [Specifying Basic VLAN Settings for Data Center Switching ELS | 511](#)
- [Specifying Advanced VLAN Profile Settings for EX Series Switches | 513](#)
- [Specifying Advanced VLAN Profile Settings for Wireless VLANs | 515](#)
- [Specifying Advanced VLAN Settings for Campus Switching ELS | 521](#)
- [Specifying Advanced VLAN Profile Settings for Data Center Switching Non-ELS | 523](#)
- [Specifying Advanced VLAN Settings for Data Center Switching ELS | 528](#)
- [Reviewing and Saving the VLAN Profile Configuration | 530](#)
- [What to Do Next | 531](#)

You can create and manage VLAN profiles on switches, wireless LAN controllers, and QFX Series devices by using the Manage VLAN Profiles window. Each VLAN profile is specific to a device family. After you create a VLAN profile, you can assign the profile at port level, switch level, or controller level. You can also assign VLAN profiles to controller managed access points and cluster managed access points.

Use the Manage VLAN Profiles page to create new VLAN profiles and to manage existing VLAN profiles.

This topic describes:

Managing VLAN Profiles

From the Manage VLAN Profiles page, you can:

- Create a new profile by clicking **Add**. For directions, see [“Creating a VLAN Profile” on page 504](#).
- Modify an existing profile by selecting the profile and clicking **Edit**.
- Assign a profile to a port, a switch, access point, or a controller by selecting the profile and clicking **Assign**. For directions, see [“Assigning a VLAN Profile to Devices or Ports” on page 531](#).
- Modify an existing assignment of a profile by selecting the profile and clicking **Edit Assignment**.
- View information about a VLAN profile, including the interfaces it is associated with, by either clicking the profile name or by selecting the profile and clicking **Details**.
- Delete profiles by selecting the profiles and clicking **Delete**.

TIP: You cannot delete profiles that are in use—that is, assigned to objects or being used by other profiles. To see the current assignments for a profile, select the profile, click **Details**, and then click the Assigned Objects Tab in the Details window.

- Clone a VLAN profile by selecting the profile and clicking **Clone**.

[Table 107](#) describes the fields in the Manage VLAN Profiles page. This page lists all VLAN profiles defined for your network.

Table 107: Manage VLAN Profile Fields

Field	Description
Profile Name	Name given to the profile when the profile was created.
VLAN Name	Name given to the VLAN when the VLAN profile was created.
Family Type	The device family; an EX Series switch, wireless LAN controller (WLC), Campus Switching ELS, or Data Center Switching.
VLAN ID	VLAN ID assigned when the profile was created.
VLAN Range	Range of VLAN IDs assigned when the profile was created. TIP: If a VLAN ID is displayed, VLAN range will be null. Also, Campus Switching ELS supports a VLAN ID range only as part of a VLAN ID list.

Table 107: Manage VLAN Profile Fields (continued)

Field	Description
VLAN ID List	VLAN IDs can be either individually listed (with a space to separate each ID), an inclusive list separating the starting VLAN ID and ending VLAN ID with a hyphen, or a combination of both. TIP: If a VLAN ID is displayed, VLAN range will be null. Also, this column will never have a value for EX Switching because it is not available.
Description	Description of the VLAN profile entered when the profile was created.
Assignment State	Displays the assignment state of the profile. A profile can be: <ul style="list-style-type: none"> • Unassigned—When the profile is not assigned to any object. • Deployed—When the profile is assigned and is deployed from Deploy mode. • Pending Deployment—When the profile is assigned, but not yet deployed in the network.
Creation Time	Date and time when the profile was created.
Last Updated Time	Date and time when the profile was last modified.
User Name	The username of the person who created or modified the profile.

TIP: All columns might not be displayed. To show or hide fields listed in the Manage Authorization Profiles table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

Creating a VLAN Profile

To create a VLAN profile, at minimum, you must specify the VLAN name and the IEEE 802.1Q VLAN tag for the profile. You also must indicate a device family for the VLAN: EX Series Switches, Wireless (WLC), Campus Switching ELS, or Data Center Switching.


In the VLAN, you can specify additional VLAN profile configuration such as:

- Ingress or egress filters to be used on the VLAN
- Parameters for handling the MAC forwarding table

To create a VLAN profile:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.
3. From the Tasks pane, select the type of network (Wired or Wireless), the appropriate functional area, and then select the name of the profile that you want to create. For example, to create a RADIUS profile for a wireless device, click **Wireless** > **Profiles** > **RADIUS**. The appropriate Manage Profile page opens.

4. Click **Add** to add a new profile.

If you chose to create a profile for the wired network, Network Director opens the Device Family Chooser window.

- a. From the Device Family Chooser, select the device family for which you want to create a profile. The available device families are **Switching (EX)**, **Campus Switching ELS** (Enhanced Layer 2 Software), **Data Center Switching Non-ELS** and **Data Center Switching ELS**.
- b. Click **OK**.

The Create VLAN Profile page for the selected device family is displayed. It consists of three sections, Basic Settings, Advanced Settings, and Review.

If you chose to create a profile for the wireless network, Network Director opens the Create VLAN Profile for Wireless page.

5. Specify the basic VLAN settings by using the appropriate directions:
 - [Specifying Basic EX Switching VLAN Settings on page 506](#)
 - [Specifying Basic Wireless VLAN Settings on page 507](#)
 - [Specifying Basic Campus Switching ELS VLAN Settings on page 508](#)
 - [Specifying Basic VLAN Settings for Data Center Switching Non-ELS on page 510](#)
 - [Specifying Basic VLAN Settings for Data Center Switching ELS on page 511](#)
6. When you have completed the basic settings, click **Next** or click **Advanced Settings** at the top of the wizard window.

7. Specify the advanced settings. Complete the Advanced Settings options as described in the online help:
 - [“Specifying Advanced VLAN Profile Settings for EX Series Switches” on page 513](#) for EX Series switches.
 - [“Specifying Advanced VLAN Profile Settings for Wireless VLANs” on page 515](#) for wireless LAN controllers.
 - [“Specifying Advanced VLAN Settings for Campus Switching ELS” on page 521](#) for Campus Switching ELS.
 - [“Specifying Advanced VLAN Profile Settings for Data Center Switching Non-ELS” on page 523](#) for Data Center Switching Non-ELS.
 - [“Specifying Advanced VLAN Settings for Data Center Switching ELS” on page 528](#) for Data Center Switching ELS.
8. When you have completed the advanced settings, click **Next** or click **Review** at the top of the wizard window.
9. You can make changes to your profile from the **Review** page. Click **Save > Finish** to save the profile. For directions, see [“Reviewing and Saving the VLAN Profile Configuration” on page 530](#).
10. Click **Finish**.

The system saves the VLAN profile and displays the Manage VLAN Profiles page. Your new or modified VLAN profile is listed in the table of VLAN profiles.

Specifying Basic EX Switching VLAN Settings

To configure the basic settings for an EX Switching VLAN profile, enter the settings described in [Table 108](#). Required settings are indicated by a red asterisk (*) that appears next to the field label.

Table 108: VLAN Profile Basic Settings for EX Switching

Field	Action
Profile Name	Type a name for the profile. You can use up to 32 characters for profiles created for wireless devices. Profile name must not contain special characters or spaces. Note that profiles that are automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character.
VLAN Name	Type the name of VLAN. The profile name and the VLAN name can be the same or different.
Description	Type a description to identify the group or function the VLAN will be part of. The character limit is 256 characters.

Table 108: VLAN Profile Basic Settings for EX Switching (continued)

Field	Action
VLAN ID	
You can indicate a single VLAN ID or a VLAN Range for EX Switching.	
Single VLAN ID	To specify a single VLAN ID, type the single unique IEEE 802.1Q identifier for the VLAN (VLAN tag). The range for VLAN IDs is 1 through 4094.
Range of VLAN IDs	To indicate a range of VLAN IDs for EX Series switches, follow these steps: <ol style="list-style-type: none"> 1. Select Range instead of Single in the VLAN ID section. 2. Provide the first and last VLAN IDs in the range. <p>TIP: For example, if you enter 10 and 12, when you deploy the profile on a device, three VLANs are created with VLAN IDs 10, 11, and 12. The names of the VLANs are created from the name you specified by adding the VLAN ID as a suffix to the name, for example vlanname_10.</p>

Click **Next** or click **Advanced Settings** at the top of the wizard window to configure advanced VLAN EX Switching profile settings. Advanced Settings are described in [“Specifying Advanced VLAN Profile Settings for EX Series Switches” on page 513](#).

Specifying Basic Wireless VLAN Settings

To configure the basic settings for a wireless VLAN profile, enter the settings described in [Table 109](#). Required settings are indicated by a red asterisk (*) that appears next to the field label.

Table 109: VLAN Wireless Profile Basic Settings

Field	Action
Profile Name	Type a unique name that identifies the profile. Use up to 32 characters for wireless profile names. Profile names must not contain special characters or spaces. Note that profiles automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character.
VLAN Name	Type the name of VLAN. The profile name and the VLAN name can be the same or different.

Table 109: VLAN Wireless Profile Basic Settings (*continued*)

Field	Action
Description	Type a description to identify the group or function the VLAN will be part of. The character limit is 256 characters.
VLAN ID	
Single VLAN ID	Type a single unique IEEE 802.1Q identifier for the VLAN (VLAN tag). The range for VLAN IDs is 1 through 4094.

Click **Next** or click **Advanced Settings** at the top of the wizard window to configure advanced wireless VLAN profile settings. Wireless Advanced Settings are described in [“Specifying Advanced VLAN Profile Settings for Wireless VLANs” on page 515](#).

Specifying Basic Campus Switching ELS VLAN Settings

To configure the basic settings for a Campus Switching ELS VLAN profile, enter the settings described in [Table 110](#). Required settings are indicated by a red asterisk (*) that appears next to the field label.

Table 110: VLAN Profile Basic Settings for Campus Switching ELS

Field	Action
Profile Name	Type a unique name that identifies the profile. You can use up to 32 characters for profiles created for wireless devices. Profile name must not contain special characters or spaces. Note that profiles that are automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character.
VLAN Name	Type the name of VLAN. The profile name and the VLAN name can be the same or be different.
Description	Type a description to identify the group or function of the VLAN. The character limit is 256 characters.
VLAN ID	
NOTE: Campus Switching ELS supports a VLAN ID range only as part of a VLAN ID list. Follow the directions for adding a list of VLAN IDs if you are adding a VLAN range.	
Single VLAN ID	To specify a single VLAN ID (default), type the single unique IEEE 802.1Q identifier for the VLAN—the VLAN tag. The range for VLAN IDs is 1 through 4094.

Table 110: VLAN Profile Basic Settings for Campus Switching ELS (continued)

Field	Action
List of VLAN IDs	<p>To create a list of VLAN IDs for switches, follow these steps:</p> <ol style="list-style-type: none"> 1. Select List instead of Single in the VLAN ID section. 2. Click Add under VLAN IDs. The Add VLAN Details window opens. 3. To add a single VLAN ID to the list, type the VLAN ID and then click either Add which closes this window or Add More which allows you to continue adding to the list. 4. To add a range of VLAN IDs to this list: <ol style="list-style-type: none"> a. In the Add VLAN Details window, select Range to add VLAN IDs in the range format 1 - 3. b. In the Add VLAN Details window, provide the first and last VLAN IDs in the range. TIP: For example, if you enter 10 and 12, when you deploy the profile on a device, three VLANs are created with VLAN IDs 10, 11, and 12. The names of the VLANs are created from the name you specified by adding the VLAN ID as a suffix to the name, for example vlanname_10. c. Click either Add to close this window, or Add More to allow you to continue adding to the list. 5. When you are finished creating the list, close the window (if it is still open). All VLAN IDs you added appear in the VLAN IDs list.
Ethernet VLAN and FCoE VLAN (applies to Data Center Switching only)	Select the type of VLAN, either Ethernet or Fibre Channel Over Ethernet (FCoE).

Click **Next** or click **Advanced Settings** at the top of the wizard window to configure advanced Campus Switching ELS VLAN profile settings. Advanced settings are described in [“Specifying Advanced VLAN Settings for Campus Switching ELS” on page 521](#).

Specifying Basic VLAN Settings for Data Center Switching Non-ELS

To configure the basic settings for a data center switching non-ELS VLAN profile, enter the settings described in [Table 111](#). Required settings are indicated by a red asterisk (*) that appears next to the field label.

Table 111: Data Center Switching Non-ELS VLAN Profile Basic Settings

Field	Action
Profile Name	Type a unique name that identifies the profile. You can use up to 32 characters for profiles created for wireless devices. Profile name must not contain special characters or spaces. Note that profiles that are automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character.
VLAN Name	Type the name of VLAN. The profile name and the VLAN name can be the same or different.
Description	Type a description to identify the group or function the VLAN will be part of. The character limit is 256 characters.
VLAN Type	Select the type of VLAN, either Ethernet or Fibre Channel Over Ethernet (FCoE).

VLAN ID

You can always indicate a single VLAN ID. You can specify a VLAN List or VLAN Range for some products. The VLAN List or VLAN Range options are listed when they apply to the VLAN profile.

Single VLAN ID	To specify a single VLAN ID, type the single unique IEEE 802.1Q identifier for the VLAN (VLAN tag). The range for VLAN IDs is 1 through 4094.
----------------	---

Table 111: Data Center Switching Non-ELS VLAN Profile Basic Settings (*continued*)

Field	Action
Range of VLAN IDs	<p>To indicate a range of VLAN IDs for switches, follow these steps:</p> <ol style="list-style-type: none"> 1. Select List instead of Single in the VLAN ID section. 2. Click Add under VLAN IDs. The Add VLAN Details window opens. 3. In the Add VLAN Details window, select Range instead of Single. 4. In the Add VLAN Details window, provide the first and last VLAN IDs in the range. TIP: For example, if you enter 10 and 12, when you deploy the profile on a device, three VLANs are created with VLAN IDs 10, 11, and 12. The names of the VLANs are created from the name you specified by adding the VLAN ID as a suffix to the name, for example vlanname_10. 5. In the Add VLAN Details window, click Add. The Add VLAN Details window closes and the VLANs are added to the VLAN IDs list.

Click **Next** or click **Advanced Settings** at the top of the wizard window to configure advanced Data Center Non-ELS VLAN profile settings. Advanced Settings are described in [“Specifying Advanced VLAN Profile Settings for Data Center Switching Non-ELS” on page 523](#).

Specifying Basic VLAN Settings for Data Center Switching ELS

To configure the basic settings for a Data Center Switching ELS VLAN profile, enter the settings described in [Table 112](#). Required settings are indicated by a red asterisk (*) that appears next to the field label.

Table 112: VLAN Profile Basic Settings for Data Center Switching ELS

Field	Action
Profile Name	<p>Type a unique name that identifies the profile.</p> <p>You can use up to 32 characters for profiles created for wireless devices. Profile name must not contain special characters or spaces. Note that profiles that are automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character.</p>

Table 112: VLAN Profile Basic Settings for Data Center Switching ELS (continued)

Field	Action
VLAN Name	Type the name of VLAN. The profile name and the VLAN name can be the same or be different.
Description	Type a description to identify the group or function of the VLAN. The character limit is 256 characters.

VLAN ID

NOTE: Data Center Switching ELS supports a VLAN ID range only as part of a VLAN ID list. Follow the directions for adding a list of VLAN IDs if you are adding a VLAN range.

Single VLAN ID	To specify a single VLAN ID (default), type the single unique IEEE 802.1Q identifier for the VLAN—the VLAN tag. The range for VLAN IDs is 1 through 4094.
-----------------------	---

Table 112: VLAN Profile Basic Settings for Data Center Switching ELS (continued)

Field	Action
List of VLAN IDs	<p>To create a list of VLAN IDs for switches, follow these steps:</p> <ol style="list-style-type: none"> 1. Select List instead of Single in the VLAN ID section. 2. Click Add under VLAN IDs. The Add VLAN Details window opens. 3. To add a single VLAN ID to the list, type the VLAN ID and then click either Add which closes this window or Add More which allows you to continue adding to the list. 4. To add a range of VLAN IDs to this list: <ol style="list-style-type: none"> a. In the Add VLAN Details window, select Range to add VLAN IDs in the range format 1 - 3. b. In the Add VLAN Details window, provide the first and last VLAN IDs in the range. TIP: For example, if you enter 10 and 12, when you deploy the profile on a device, three VLANs are created with VLAN IDs 10, 11, and 12. The names of the VLANs are created from the name you specified by adding the VLAN ID as a suffix to the name, for example vlanname_10. c. Click either Add to close this window, or Add More to allow you to continue adding to the list. 5. When you are finished creating the list, close the window (if it is still open). All VLAN IDs you added appear in the VLAN IDs list.

Click **Next** or click **Advanced Settings** at the top of the wizard window to configure advanced Data Center Switching ELS VLAN profile settings. Advanced Settings are described in [“Specifying Advanced VLAN Settings for Data Center Switching ELS” on page 528](#).

Specifying Advanced VLAN Profile Settings for EX Series Switches

To configure the EX Switching advanced settings for the VLAN profile, enter the MAC parameters and Layer 2 filters described in [Table 113](#) for EX Series switching. All settings are optional.

Table 113: VLAN Profile Advanced Settings for an EX Series Switch

EX Switching MAC Parameters	
MAC Limit	<p>Type the number of dynamic MAC addresses that can be learned on the VLAN. If this number is exceeded, packets containing new MAC addresses are dropped and an alarm is raised.</p> <p>Setting a limit on the number of dynamic MAC addresses protects against an Ethernet switching table overflow attack.</p>
MAC Aging Time (ms)	<p>Indicate the number of milliseconds that unused dynamic MAC addresses remain in the MAC forwarding table before being deleted. If you specify the time as unlimited, entries are never removed from the table. Generally, use this setting only if the switch or the VLAN has a fairly static number of end devices—otherwise the table will eventually fill up. You can use this setting to minimize traffic loss and flooding that might occur when traffic arrives for MAC addresses that have been removed from the table.</p> <p>The range is from 60 through 1,000,000.</p>
EX Switching L2 Filters	
L2 Ingress Filter	<p>Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter Profile window and click OK. The filter configuration contained in the profile is applied to ingress traffic on the VLAN.</p> <p>To remove the selected Filter profile, click Clear.</p>
L2 Egress Filter	<p>Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter profile window and click OK. The filter configuration contained in the profile is applied to egress traffic on the VLAN.</p> <p>To remove the selected Filter profile, click Clear.</p>
EX Switching L3 Routing Filters	
<p>If you indicated a single VLAN ID under the Basic Settings, you can specify one or more routing parameters (Layer 3 filters) for the profile.</p>	
L3 Ingress Filter L3 IPv6 Ingress Filter	<p>Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter profile window and click OK. The filter configuration contained in the profile is applied to ingress traffic on the VLAN.</p> <p>To remove the selected Filter profile, click Clear.</p>

Table 113: VLAN Profile Advanced Settings for an EX Series Switch (*continued*)

L3 Egress Filter	Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter profile window and click OK . The filter configuration contained in the profile is applied to egress traffic on the VLAN. To remove the selected Filter profile, click Clear .
L3 IPv6 Egress Filter	

VLAN Security Settings

Optionally, select VLAN Security Settings to display the security options DHCP, ARP inspection, and MAC movement limit for VLAN profiles for EX switching.

Enable DHCP Snooping	Check to apply a series of security techniques to the DHCP infrastructure.
Enable ARP Inspection	The Address Resolution Protocol (ARP), which provides IP communication within a Layer 2 broadcast domain by mapping an IP address to a MAC address, has security issues. Select this option to apply inspection to untrusted interfaces.
MAC Movement Limit	Indicate the number of times a MAC address entry can be moved in the MAC address table without consequences.
MAC Movement Action	When a MAC Movement Limit is specified, select an action to be applied to MAC addresses that exceed the MAC Movement Limit: None , Log , Drop , Shut Down , or Drop and Log .
VRRP Settings	Select the VRRP profile for the interface from a list of existing profiles by clicking Select . Select one of the listed profiles, and then click OK .

Click **Next** or click **Review** to see the Review page of the wizard. For review directions, see [“Reviewing and Saving the VLAN Profile Configuration” on page 530](#).

Specifying Advanced VLAN Profile Settings for Wireless VLANs

To configure the advanced settings for the wireless VLAN profile, enter the settings described in [Table 114](#) for a wireless LAN controller. All fields are optional.

Table 114: VLAN Profile Advanced Settings for a Wireless LAN Controller

Field	Action
Filter	

Table 114: VLAN Profile Advanced Settings for a Wireless LAN Controller (*continued*)

Field	Action
Task: Add a Filter to the VLAN	<p>Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter profile window and click OK. The filter configuration contained in the profile is applied to ingress traffic or egress traffic or both based on the selected direction.</p> <p>To remove the selected Filter profile, click Clear.</p> <hr/> <p>Direction: If you added a filter, select the direction of the traffic towards which the filter will be applied. The options are:</p> <ul style="list-style-type: none"> • In—to apply the filter to ingress traffic only. • Out—to apply the filter to egress traffic only. • Both—to apply the filter for both ingress and egress traffic. • None—if you do not want to apply the filter to either ingress or egress traffic. <p>The default value is <i>Both</i>, which means the filter is applied for both ingress and egress traffic.</p> <p>NOTE: The direction <i>None</i> is applicable only when no Filter profile is selected.</p>
Enable DHCP (Disabled by default)	Enable DHCP client for the VLAN. By selecting the check box, you enable VLAN to obtain the Layer 3 interface IP address from the DHCP server.
IGMP Settings	
<p>Optionally, expand the IGMP Settings section and modify the default settings. Internet Group Management Protocol (IGMP) snooping controls multicast traffic on a controller by forwarding packets for a multicast group only on the ports that are connected to members of the group. The controller listens for multicast packets and maintains a table of multicast groups, as well as their sources and receivers, based on the traffic.</p>	
IGMP Enabled	IGMP is enabled by default. You can disable it by removing the check mark.
Version (default is Version 2)	Select the IGMP version, Version 1 or Version 2 , from the list.

Table 114: VLAN Profile Advanced Settings for a Wireless LAN Controller (*continued*)

Field	Action
<p>Querier Enabled (default is disabled)</p>	<p>Select the check box to enable pseudo-query.</p> <p>TIP: The IGMP querier enables IGMP snooping to operate in a VLAN without a multicast router to send IGMP general queries to clients.</p> <p>Juniper Networks recommends that you enable the querier only when the VLAN contains local multicast traffic sources and no multicast router is supporting the subnet.</p>
<p>Query Interval (default is 125 seconds)</p>	<p>Enter or select the number of seconds that elapse between general queries sent by the controller to advertise multicast groups.</p> <p>You can specify a value from 1 through 65,535. The default is 125 seconds.</p>
<p>Other Querier Present Interval (default is 255 seconds)</p>	<p>Enter or select the number of seconds that the controller waits for a general query to arrive from another querier before becoming the querier.</p> <p>You can specify a value from 1 through 65,535. The default is 255 seconds.</p>
<p>Query Response Interval (default is 1 second)</p>	<p>Enter or select the number of seconds, in tenths, that the controller waits for a receiver to respond to a group-specific query message before removing the receiver from the group receiver list.</p> <p>You can specify a value from 1 through 255 tenths of a second. The default is 100 tenths of a second (10 seconds).</p> <p>TIP: The query interval, other-querier-present interval, and query response interval are applicable only when the controller is the querier for the subnet. For the controller to become the querier, the querier feature must be enabled on the controller and the controller must have the lowest IP address among all the devices eligible to become a querier.</p>

Table 114: VLAN Profile Advanced Settings for a Wireless LAN Controller (*continued*)

Field	Action
Last Member Query Interval (default is 1 second)	<p>Enter or select the number of tenths of a second that the controller waits for a response to a group-specific query after receiving a leave message for that group, before removing the receiver that sent the leave message from the list of receivers for the group. If there are no more receivers for the group, the controller also sends a leave message for the group to multicast routers.</p> <p>You can specify a value from 1 through 255 tenths of a second. The default is 10 tenths of a second (1 second).</p>
Robustness Value (default is 2)	<p>Enter or select a number used as a multiplier to adjust the IGMP timers to the amount of traffic loss that occurs on the network. Set a higher value to adjust for more traffic loss.</p> <p>You can specify a value from 2 through 255. The default is 2.</p>
Proxy Request (default is enabled)	<p>Enable or disable proxy request. Proxy request is enabled by default.</p> <p>Proxy request reduces multicast overhead by sending only one request for each active group to the multicast routers, instead of sending a separate request from each multicast receiver.</p>
Multicast Router Solicitation (default is disabled)	<p>Enable or disable multicast router solicitation. Router solicitation is disabled by default.</p> <p>A controller can search for multicast routers by sending multicast router solicitation messages. This message invites multicast routers receiving the message and support router solicitation to immediately advertise themselves to the controller.</p>
Solicitation Interval (default is 30 seconds)	<p>If you enabled Multicast Router Solicitation, enter or select the multicast router solicitation interval in seconds. You can specify a value from 1 through 65,535 seconds.</p> <p>The default multicast router solicitation interval is 30 seconds.</p>
Spanning Tree Protocol Settings	
Optionally, expand Spanning Tree Protocol Settings to enable and configure STP on this VLAN.	
Enabled (default is disabled)	Select the check box to enable the Spanning Tree Protocol (STP) on the VLAN profile.

Table 114: VLAN Profile Advanced Settings for a Wireless LAN Controller (*continued*)

Field	Action
Instance Number	The VLAN ID you indicated under basic settings is reflected here. You cannot change this number.
Bridge Priority (default is 32768)	<p>Enter or select a bridge priority number from 0 through 65,535. The default bridge priority for all devices is 32,768.</p> <p>The bridge priority determines the controllers eligibility to become the root bridge. You can set this parameter globally or on individual VLANs.</p> <p>The root bridge is elected on the basis of the bridge priority of each device in the spanning tree. The device with the highest bridge priority is elected to be the root bridge for the spanning tree.</p> <p>The bridge priority is a numeric value from 0 through 65,535. Lower numeric values represent higher priorities. The highest priority is 0, and the lowest priority is 65,535.</p> <p>If more than one device has the highest bridge priority (lowest numeric value), the device with the lowest MAC address becomes the root bridge. If the root bridge fails, STP elects a new root bridge on the basis of the bridge priorities of the remaining bridges.</p>
Protocol (default is PVST)	<p>Enter the protocol for the VLAN.</p> <p>Network Director supports 802.1D and Per-VLAN Spanning Tree plus (PVST+).</p>
Max Age (default is 20 seconds)	<p>Enter or select an age from 6 through 40 seconds. The default is 20 seconds.</p> <p>The period of time that a controller acting as a designated bridge waits for a new hello packet from the root bridge before determining that the root bridge is no longer available and is initiating a topology change. You can specify a range from 6 through 40 seconds. The default is 20 seconds.</p>

Table 114: VLAN Profile Advanced Settings for a Wireless LAN Controller (*continued*)

Field	Action
Hello Time (default is 2 seconds)	<p>Enter or select an interval from 1 through 10 seconds. The default is 2 seconds.</p> <p>The interval between configuration messages sent by a controller when the controller is acting as the root bridge. You can specify an interval from 1 through 10 seconds. The default is 2 seconds.</p>
Forward Delay (default is 15 seconds)	<p>Enter or a select a time delay from 4 through 30 seconds. The default is 15 seconds.</p> <p>The period of time a bridge other than the root bridge waits after receiving a topology change notification to begin forwarding data packets. You can specify a delay from 4 through 30 seconds. The default is 15 seconds. (The root bridge always forwards traffic.)</p>

mDNS Settings

mDNS is a simple way to enable Apple TV, Internet Printing, and/or iTunes on this VLAN profile. For more information, see ["Understanding Bonjour" on page 995](#).

TIP: You must have an existing mDNS Profile to add to a VLAN Profile. To create an mDNS Profile, see ["Creating and Managing mDNS Profiles" on page 997](#).

mDNS Profile	Add an existing mDNS Profile to the VLAN Profile by clicking Select , selecting one of the listed mDNS Profiles, and then clicking OK . The profile name is now listed in the mDNS Profile field.
---------------------	---

Table 114: VLAN Profile Advanced Settings for a Wireless LAN Controller (*continued*)

Field	Action
<p>Location Service</p>	<p>To specify the mDNS services for this VLAN Profile (Apple TV, Internet Printing, iTunes):</p> <ol style="list-style-type: none"> 1. Provide a unique name for the service at this VLAN Profile. 2. Add one or more of the available mDNS services to this VLAN by clicking Add under Location Service. The phrase <i>Enter service here</i> appears in the list of services. 3. Select the phrase <i>Enter service here</i>. A list box replaces the phrase. 4. From the list of services, select one of the following: <ul style="list-style-type: none"> • _airplay._tcp—Apple TV • _ipp._tcp—Internet printer • _daap._tcp—Digital Auto Access Protocol (iTunes) • All The selected services are listed under Services.
<p>Restrict L2 Traffic</p> <p>Optionally, check Restrict L2 Traffic to display the list of restricted MAC addresses. You can also add or delete MAC addresses.</p>	
<p>Task: Add a MAC address to the list</p>	<p>To add a MAC address, click Add, double-click <i>Enter MAC address here...</i>, and then type the MAC address in the format 12:ae:53:ef:56:76.</p>
<p>Task: Remove a MAC address from the list</p>	<p>To remove a MAC address, select one of the MAC addresses from the list and then click Delete.</p>

Click **Next** or click **Review** to see the Review page of the wizard. For review directions, see [“Reviewing and Saving the VLAN Profile Configuration” on page 530](#).

Specifying Advanced VLAN Settings for Campus Switching ELS

To configure the advanced settings for a Campus Switching ELS VLAN profile, specify the parameters described in [Table 115](#) for Campus Switching ELS. All settings are optional.

Table 115: VLAN Profile Advanced Settings for Campus Switching ELS

Field	Action
Campus Switching ELS MAC Parameters	
Interface MAC Limit	<p>Indicate the number of dynamic MAC addresses that can be learned on the VLAN. If this number is exceeded, packets containing new MAC addresses are dropped and an alarm is raised.</p> <p>Setting a limit on the number of dynamic MAC addresses protects against an Ethernet switching table overflow attack.</p>
Packet Action	Indicate the packet action for MAC addresses that exceed the Interface MAC Limit, by selecting None , Log , Drop , Shut Down , or Drop and Log .
MAC Table Size	If you indicated an Interface MAC limit, provide a table size here by using the up and down arrows. The MAC table must allow for at least 16 entries—you can increase this limit with the arrow.
L2 Filters	
Ingress Filter	<p>Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter Profile window and then click OK. The filter configuration contained in the profile is applied to ingress traffic on the VLAN.</p> <p>To remove a selected Filter profile, click Clear.</p>
Egress Filter	<p>Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter profile window and then click OK. The filter configuration contained in the profile is applied to egress traffic on the VLAN.</p> <p>To remove a selected Filter profile, click Clear.</p>
Routing	
<p>If you selected a single VLAN ID under Basic Settings, you can specify Layer 3 filter routing parameters for the VLAN profile.</p> <p>NOTE: If an IP address is configured for a VLAN on some devices, then the configured IP address will be retained and a DHCP client will not be enabled on those devices. Also, if you indicated a VLAN range for basic ELS switching configuration, this option is not available.</p>	
Routing L3 Filters	

Table 115: VLAN Profile Advanced Settings for Campus Switching ELS (continued)

Field	Action
Ingress Filter IPv6 Ingress Filter	Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter profile window and then click OK . The filter configuration contained in the profile is applied to ingress traffic on the VLAN. To remove a selected Filter profile, click Clear .
Egress Filter IPv6 Egress Filter	Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter profile window and then click OK . The filter configuration contained in the profile is applied to egress traffic on the VLAN. To remove a selected Filter profile, click Clear .

VLAN Security Settings

Optionally, enable VLAN Security Settings to display the security options DHCP, ARP inspection, and MAC movement limit for VLAN profiles for ELS switching.

Enable DHCP Snooping	When checked (default), this option applies a series of security techniques to the DHCP infrastructure.
Enable ARP Inspection	The Address Resolution Protocol (ARP), which provides IP communication within a Layer 2 broadcast domain by mapping an IP address to a MAC address, has security issues. Select this option to apply inspection to untrusted interfaces.
MAC Movement Limit	Indicate the number of times a MAC address entry can be moved in the MAC address table without consequences.
MAC Movement Action	When a MAC Movement Limit is specified, select an action to be applied to MAC addresses that exceed the MAC Movement Limit: None , Log , Drop , Shut Down , or Drop and Log .

Click **Next** or click **Review** to see the Review page of the wizard. For review directions, see [“Reviewing and Saving the VLAN Profile Configuration” on page 530](#).

Specifying Advanced VLAN Profile Settings for Data Center Switching Non-ELS

To configure the advanced settings for a data center switching non-ELS VLAN profile:

1. Enter advanced settings for the profile on the Advanced Setting page.

The settings that are available depend on which types of VLAN you are configuring—Ethernet or FCoE. Settings for Ethernet VLANs are described in both the online help and in [Table 116](#). Settings for FCoE VLANs are described in both the online help and in [Table 117](#). All settings are optional.

Table 116: VLAN Profile Advanced Settings for a Data Center Switching Non-ELS Ethernet VLAN

Field	Action
Switching	Specify the switching parameters (MAC parameters and Layer 2 filters) for the profile.
MAC Parameters	
MAC Limit	<p>Type the number of dynamic MAC addresses that can be learned on the VLAN. If this number is exceeded, packets containing new MAC addresses are dropped and an alarm is raised.</p> <p>Setting a limit on the number of dynamic MAC addresses protects against an Ethernet switching table overflow attack.</p>
MAC Aging Time (ms)	<p>Indicate the number of milliseconds unused dynamic MAC addresses remain in the MAC forwarding table before being deleted. If you specify the time as unlimited, entries are never removed from the table. Generally, use this setting only if the switch or the VLAN has a fairly static number of end devices—otherwise the table will eventually fill up. You can use this setting to minimize traffic loss and flooding that might occur when traffic arrives for MAC addresses that have been removed from the table.</p> <p>The range is from 60 through 1,000,000 ms.</p>
L2 Filters	
Ingress Filter	<p>Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter Profile window and click OK. The filter configuration contained in the profile is applied to ingress traffic on the VLAN.</p> <p>To remove the selected Filter profile, click Clear.</p>
Egress Filter	<p>Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter profile window and click OK. The filter configuration contained in the profile is applied to egress traffic on the VLAN.</p> <p>To remove the selected Filter profile, click Clear.</p>

Table 116: VLAN Profile Advanced Settings for a Data Center Switching Non-ELS Ethernet VLAN (continued)

Field	Action
Routing	
<p>If you selected a single VLAN ID under Basic Settings, you can specify Layer 3 filter routing parameters for the VLAN profile.</p>	
<p>NOTE: If an IP address is configured for a VLAN on some devices, then the configured IP address will be retained and a DHCP client will not be enabled on those devices. Also, if you indicated a VLAN range for basic ELS switching configuration, this option is not available.</p>	
Routing L3 Filters	
Ingress Filter IPv6 Ingress Filter	<p>Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter profile window and then click OK. The filter configuration contained in the profile is applied to ingress traffic on the VLAN.</p> <p>To remove a selected Filter profile, click Clear.</p>
Egress Filter IPv6 Egress Filter	<p>Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter profile window and then click OK. The filter configuration contained in the profile is applied to egress traffic on the VLAN.</p> <p>To remove a selected Filter profile, click Clear.</p>
VLAN Security Settings	
Enable DHCP Snooping	When checked (default), this option applies a series of security techniques to the DHCP infrastructure.
Enable ARP Inspection	The Address Resolution Protocol (ARP), which provides IP communication within a Layer 2 broadcast domain by mapping an IP address to a MAC address, has security issues. Select this option to apply inspection to untrusted interfaces.
MAC Movement Limit	Indicate the number of times a MAC address entry can be moved in the MAC address table without consequences.
Fabric Limit	Specify the maximum number of times a MAC address can move in a QFabric system. If no fabric limit is specified then the value given for the mac-move-limit applies to the QFabric system.
MAC Movement Action	When a MAC Movement Limit is specified, select an action to be applied to MAC addresses that exceed the MAC Movement Limit. The options are: None , Log , Drop , Shut Down , and Drop and Log .

Table 116: VLAN Profile Advanced Settings for a Data Center Switching Non-ELS Ethernet VLAN (continued)

Field	Action
FIP Snooping Settings	
Enable V2V2N Snooping	Select to enable VN_Port to VN_Port (VN2VN) FIP snooping on the VLAN.
Beacon Period (ms)	Set the interval between periodic beacons, in milliseconds. Beacons perform virtual link maintenance for VN_Ports in a way that is similar to FIP keepalive advertisements. Range: 250 through 90000 milliseconds. Default: 8000 milliseconds.
FC Map	Set the FCoE mapped address prefix (FC-MAP) value for the FCoE VLAN to match the FC switch (or FCoE forwarder) FC-MAP value for the FC fabric. The FC-MAP value is a unique MAC address prefix an FC switch uses to identify FCoE traffic for a given FC fabric (traffic on a particular FCoE VLAN). Range: 0x0EFC00 through 0x0EFCFF. Default: 0xEFC00

[Table 117](#) describes the advanced settings for a Data Center Switching Non-ELS FCoE VLAN.

Table 117: VLAN Profile Advanced Settings for a Data Center Switching Non-ELS FCoE VLAN

Field	Action
Switching	Specify the switching parameters (MAC parameters and Layer 2 filters) for the profile.
MAC Parameters	
MAC Limit	Type the number of dynamic MAC addresses that can be learned on the VLAN. If this number is exceeded, packets containing new MAC addresses are dropped and an alarm is raised. Setting a limit on the number of dynamic MAC addresses protects against an Ethernet switching table overflow attack.

Table 117: VLAN Profile Advanced Settings for a Data Center Switching Non-ELS FCoE VLAN (continued)

Field	Action
MAC Aging Time (ms)	<p>Indicate the number of milliseconds unused dynamic MAC addresses remain in the MAC forwarding table before being deleted. If you specify the time as unlimited, entries are never removed from the table. Generally, use this setting only if the switch or the VLAN has a fairly static number of end devices—otherwise the table will eventually fill up. You can use this setting to minimize traffic loss and flooding that might occur when traffic arrives for MAC addresses that have been removed from the table.</p> <p>The range is from 60 through 1,000,000.</p>
L2 Filters	
Ingress Filter	<p>Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter Profile window and click OK. The filter configuration contained in the profile is applied to ingress traffic on the VLAN.</p> <p>To remove the selected Filter profile, click Clear.</p>
Egress Filter	<p>Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter profile window and click OK. The filter configuration contained in the profile is applied to egress traffic on the VLAN.</p> <p>To remove the selected Filter profile, click Clear.</p>
VLAN Security Settings	
Enable DHCP Snooping	Select the check box to enable DHCP snooping.
Enable ARP Inspection	Select the check box to enable ARP Inspection.
MAC Movement Limit	Enter the MAC movement limit.
Fabric Limit	Specify the maximum number of times a MAC address can move in a QFabric system. If no fabric limit is specified then the value given for the mac-move-limit applies to the QFabric system.
MAC Movement Action	Select one of the options. The options are: None, Log, Drop, and Shutdown.
FIP Snooping Settings	
<p>Selecting the FIP Snooping check box enables FIP snooping with the default FC-Map. This stops access for all traffic other than FCoE traffic.</p>	

Table 117: VLAN Profile Advanced Settings for a Data Center Switching Non-ELS FCoE VLAN (continued)

Field	Action
Enable VN2VN Snooping	Select to enable VN_Port to VN_Port (VN2VN) FIP snooping on the VLAN.
Beacon Period (ms)	Set the interval between periodic beacons, in milliseconds. Beacons perform virtual link maintenance for VN_Ports in a way that is similar to FIP keepalive advertisements. Range: 250 through 90000 milliseconds. Default: 8000 milliseconds.
FC Map	Set the FCoE mapped address prefix (FC-MAP) value for the FCoE VLAN to match the FC switch (or FCoE forwarder) FC-MAP value for the FC fabric. The FC-MAP value is a unique MAC address prefix an FC switch uses to identify FCoE traffic for a given FC fabric (traffic on a particular FCoE VLAN). Range: 0x0EFC00 through 0x0EFCFF. Default: 0xEFC00.

2. Click **Next** or click **Review** to see the Review page of the wizard. For review directions, see [“Reviewing and Saving the VLAN Profile Configuration” on page 530](#).

You can either save your profile or make changes to your profile from the **Review** page.

Specifying Advanced VLAN Settings for Data Center Switching ELS

To configure the advanced settings for a Data Center Switching ELS VLAN profile, specify the parameters described in [Table 118](#) for an Ethernet VLAN profile. All settings are optional.

Table 118: VLAN Profile Advanced Settings for Data Center Switching ELS Ethernet VLAN

Field	Action
Data Center Switching ELS MAC Parameters	
Interface MAC Limit	Indicate the number of dynamic MAC addresses that can be learned on the VLAN. If this number is exceeded, packets containing new MAC addresses are dropped and an alarm is raised. Setting a limit on the number of dynamic MAC addresses protects against an Ethernet switching table overflow attack.
Packet Action	Indicate the packet action for MAC addresses that exceed the Interface MAC Limit. The options are: None , Log , Drop , Shut Down , and Drop and Log .

Table 118: VLAN Profile Advanced Settings for Data Center Switching ELS Ethernet VLAN (continued)

Field	Action
MAC Table Size	If you indicated an Interface MAC limit, provide a table size here by using the up and down arrows. The MAC table must allow for at least 16 entries—you can increase this limit by using the arrow.
L2 Filters	
Ingress Filter	Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter Profile window and then click OK . The filter configuration contained in the profile is applied to ingress traffic on the VLAN. To remove a selected Filter profile, click Clear .
Egress Filter	Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter profile window and then click OK . The filter configuration contained in the profile is applied to egress traffic on the VLAN. To remove a selected Filter profile, click Clear .
Routing	
If you selected a single VLAN ID under Basic Settings, you can specify Layer 3 filter routing parameters for the VLAN profile.	
NOTE: If an IP address is configured for a VLAN on some devices, then the configured IP address will be retained and a DHCP client will not be enabled on those devices. Also, if you indicated a VLAN range for basic ELS switching configuration, this option is not available.	
Routing L3 Filters	
Ingress Filter IPv6 Ingress Filter	Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter profile window and then click OK . The filter configuration contained in the profile is applied to ingress traffic on the VLAN. To remove a selected Filter profile, click Clear .
Egress Filter IPv6 Egress Filter	Click Select to choose from existing Filter profiles. Select a profile from the Choose Filter profile window and then click OK . The filter configuration contained in the profile is applied to egress traffic on the VLAN. To remove a selected Filter profile, click Clear .
VLAN Security Settings	
Optionally, enable VLAN Security Settings to display the security options DHCP, ARP inspection, and MAC movement limit for VLAN profiles for ELS switching.	

Table 118: VLAN Profile Advanced Settings for Data Center Switching ELS Ethernet VLAN (continued)

Field	Action
Enable DHCP Snooping	When checked (default), this option applies a series of security techniques to the DHCP infrastructure.
Enable ARP Inspection	The Address Resolution Protocol (ARP), which provides IP communication within a Layer 2 broadcast domain by mapping an IP address to a MAC address, has security issues. Select this option to apply inspection to untrusted interfaces.
MAC Movement Limit	Indicate the number of times a MAC address entry can be moved in the MAC address table without consequences.
MAC Movement Action	When a MAC Movement Limit is specified, select an action to be applied to MAC addresses that exceed the MAC Movement Limit. The options are: None , Log , Drop , Shut Down , and Drop and Log .
FIP Snooping Settings	
Enable VN2VN Snooping	Select to enable VN_Port to VN_Port (VN2VN) FIP snooping on the VLAN.
Beacon Period (ms)	Set the interval between periodic beacons, in milliseconds. Beacons perform virtual link maintenance for VN_Ports in a way that is similar to FIP keepalive advertisements. Range: 250 through 90000 milliseconds. Default: 8000 milliseconds.
FC Map	Set the FCoE mapped address prefix (FC-MAP) value for the FCoE VLAN to match the FC switch (or FCoE forwarder) FC-MAP value for the FC fabric. The FC-MAP value is a unique MAC address prefix an FC switch uses to identify FCoE traffic for a given FC fabric (traffic on a particular FCoE VLAN). Range: 0x0EFC00 through 0x0EFCFF. Default: 0xEFC00.

Click **Next** or click **Review** to see the Review page of the wizard. For review directions, see [“Reviewing and Saving the VLAN Profile Configuration” on page 530](#).

Reviewing and Saving the VLAN Profile Configuration

From this page, you can either save the VLAN profile or make changes to the VLAN profile:

- To make changes to the profile, click the **Edit** associated with the configuration you want to change. Alternatively, you can click **Basic Settings** or **Advanced Settings** from the wizard workflow at the top of the page and make changes there.

When you are finished with your modifications, click **Review** to return to this page.

- To save a new profile or to save modified settings to an existing profile, click **Finish**.

The Manage VLAN Profiles page is displayed and your new or modified VLAN profile is listed in the table of VLAN profiles.

What to Do Next

Once the VLAN profile is created, you must assign the VLAN profile from the Assign VLAN Profile page to the required ports, switches, or controllers. You can also assign VLAN profiles to controller managed access points and cluster managed access points. To assign a VLAN profile, see “[Assigning a VLAN Profile to Devices or Ports](#)” on page 531. After you assign a VLAN profile to a port, switch, access point, or controller, you must deploy the profile configuration from the Deploy mode. For directions on deploying your configurations, see “[Deploying Configuration to Devices](#)” on page 1180.

FCoE VLANs are assigned to Fabric profiles, where they define the FCoE VLAN for a gateway FC fabric.

RELATED DOCUMENTATION

[Assigning a VLAN Profile to Devices or Ports](#) | 531

[Deploying Configuration to Devices](#) | 1180

[Understanding VLAN Profiles](#) | 499

[Understanding VRRP Profiles](#) | 845

[Creating and Managing VRRP Profiles](#) | 846

[Network Director Documentation home page](#)

Assigning a VLAN Profile to Devices or Ports

IN THIS SECTION

- [Assigning a VLAN Profile](#) | 532
- [Editing Profile Assignments](#) | 535


After a VLAN profile is created, assign it to wireless ports, access points, switches, controllers, aggregation devices in a Junos Fusion fabric, Virtual Chassis Fabric, members of custom groups, members of port

groups with wireless ports, or access points managed by controllers or clusters. You can also assign VLAN profiles to controller- managed access points and cluster- managed access points for local switching on the access points.

You must have one or more existing VLAN profiles, either user-configured or system-created, before you can assign a VLAN profile to a switch, wireless port, or controller, access point, or member of a custom group or port group. For further directions, see [“Creating and Managing VLAN Profiles” on page 502](#), [“Creating Custom Device Groups” on page 276](#), and [“Creating and Managing Port Groups” on page 495](#).

Assigning a VLAN Profile

To assign a VLAN profile:

1. Click  in the Network Director banner.
2. Select **VLAN** from the Profile and Configuration Management menu in the Tasks pane.
The Manage VLAN profiles page is displayed. The page displays all user-configured and system-created VLAN profiles for discovered devices.
3. Select a VLAN profile from the list of VLAN profiles and then click **Assign**.

The Assign VLAN Profile page for the selected VLAN appears.

NOTE: If Network Director fails to read the configuration of one or more devices after the device discovery, those devices are not displayed in the Assign Profile page. You will not be able to assign profiles to those devices. The Manage Jobs page in System mode displays details of the device discovery jobs. Use the information displayed on this page to take appropriate corrective steps to enable Network Director to reread the configuration of the failed device. For more information, see [“Discovering Devices in a Physical Network” on page 204](#).

4. If you are assigning a VLAN profile to a device, a tree is displayed. Choose at least one device or cluster for VLAN assignment. Be sure that a check mark appears in front of the device - just highlighting the name does not select it.

If you are assigning a VLAN profile to members of a Custom Group or Port Group, selecting the group selects all members of that group.

5. Click either **Next** or **Profile Assignment**.

The Profile Assignment page displays the list of existing assignments in the Assignments table.

6. Select a device from the Assignments table and click **Assign to Device**.

NOTE: When you assign a VLAN profile to a wireless port, only those ports, with unit value 0 or none, and the family Ethernet-switching or none, will be available.

7. If you are assigning the profile to an EX9200 and a routing instance has been created on this network, you have a choice between assigning the profile to the EX9200 device or assigning the profile to the routing instance. If a routing instance has been created with the Junos CLI, select either the EX9200 device or the routing instance in the window that opens. If no routing instance exists, the profile is automatically assigned to the device.

NOTE: A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. The set of interfaces belongs to the routing tables, and the routing protocol parameters control the information in the routing tables. There can be multiple routing tables for a single routing instance—for example, unicast IPv4, unicast IPv6, and multicast IPv4 routing tables can exist in a single routing instance. Routing protocol parameters and options control the information in the routing tables.

8. Click **Define** from the **Attributes** column in the Assignments table to modify the attributes. The Configure Attributes page is displayed:

For a wireless LAN controller, specify the following details:

- Enter the Layer 3 interface IP address and logical interface number (between 0 through 32).
- Select the tunnel affinity using the spinner. The available range is from 0 through 10.
- Select the check box to enable the DHCP server.
- Enter the starting IP address for the DHCP server.
- Enter the ending (stop) IP address for the DHCP server.
- Enter the default gateway IP address for the DHCP server.
- Enter the DNS name of the DHCP server.
- Enter the primary DNS address for the DHCP server.
- Enter the secondary DNS server for the DHCP server.
- Click **Save** once you are done with specifying the attributes.

For the wireless port level attributes of a VLAN profile for a controller, specify the following:

- Select the Tag check box to tag the port.
- Select the check box to enable the STP port.
- Select the port priority using the spinner. The available range is from 0 through 255.
- Select the port path cost using the spinner. The available range is from 0 through 65,535

While assigning to an access point specify the following the attributes:

- Select the Tag check box to tag the port.
- Enter the tag value. The range is from 1 through 4093.

For switches, specify the following details:

- Enter the Layer 3 interface IP address and port number (between 0 through 32).
- Enter the unit number. By default, the VLAN ID is populated in this field. You can optionally change the value.

9. You can view the assignment details for the selected device or delete any assignments.

- To view assignment details, select a device and click **View Assignments**.

The Profile Details page for selected device appears. Expand the **Device** name to see the details of the assignment. The assignment status displays the status whether the device is deployed or is pending device update, etc.

- To delete a VLAN profile assignment for a device, select the device from the Assignments table and click **Remove**.

10. Click **Next** or click **Review** from the top wizard workflow to review the assignments. Alternately, click **Edit** to edit the profile assignment.

11. Click **Finish** once you are done reviewing the profile assignment.

After you click Finish, the Create Profile Assignments Job Details dialog box appears, which reports on the status of the profile assignment job. If you have assigned the profile to a large number of objects, the profile assignment job can take some time to complete. Instead of waiting for the Job Details dialog box to report job completion status, you can close it and check the details of the profile assignment job at a later time using the Manage Job task in System mode.

NOTE: If any assignment fails, the profile assignment job fails and none of the assignments are created. Check the details for the profile assignment job for information about why the assignment failed.

After you assign a VLAN profile to a device or port, you can deploy the VLAN profile from the **Deploy** mode. For details, see [“Deploying Configuration to Devices” on page 1180](#).

Editing Profile Assignments

You can edit VLAN assignments from the Manage VLAN Profiles page. To edit an existing assignment:

1. Select a profile from the Manage VLAN profiles page and then click **Edit Assignment**.

The Edit Assignments page for the selected device appears.

2. Expand the **Devices** cabinet select a device.

3. Make the required changes in the **Operation** column of the table.

To change the attributes, see step 8, from the Assign VLAN profiles tasks.

4. Click **Apply**.

The Manage VLAN Profiles page reappears.

RELATED DOCUMENTATION

[Creating and Managing VLAN Profiles | 502](#)

[Deploying Configuration to Devices | 1180](#)

[Creating Custom Device Groups | 276](#)

[Creating and Managing Port Groups | 495](#)

[Understanding VLAN Profiles | 499](#)

[Network Director Documentation home page](#)

Creating and Managing VLAN Pools

IN THIS SECTION

- [Managing VLAN Pools | 536](#)
- [Creating a VLAN Pool | 538](#)
- [What To Do Next | 539](#)

VLAN pooling is a feature that enables you to group multiple wireless controller VLANs to form a VLAN pool. Configure a VLAN pool to load-balance sessions evenly across multiple VLANs. Individual VLANs are then assigned dynamically from the pool, using a round robin algorithm, when a wireless client accesses the network.

Use the Manage VLAN Pools page to create new VLAN pools and manage existing VLAN pools.

This topic describes:

Managing VLAN Pools

From the Manage VLAN Pools page, you can:

- Create a new wireless controller VLAN pool by clicking **Add**. For directions, see [“Creating a VLAN Pool” on page 538](#).
- Modify an existing VLAN pool by selecting it and clicking **Edit**.
- View information about a wireless controller VLAN pool by either clicking the profile name or by selecting the profile and clicking **Details**.
- Delete a VLAN pool by selecting pool name and clicking **Delete**.

TIP: You cannot delete pools that are in use. To see the current state of a pool, select the pool name and click **Details**.

- Clone a VLAN pool by selecting a pool and clicking **Clone**.

[Table 119](#) describes the information provided about VLAN pools. This page lists all VLAN pools defined for your network, regardless of the scope you selected in the network view.

Table 119: Managing VLAN Pools

Field	Description
Profile Name	Unique name, assigned during creation of the profile, that identifies the profile.
Description	Any details that were added when the pool was created

Table 119: Managing VLAN Pools (continued)

Field	Description
Selection Method	<p>Method used to select a VLAN from a VLAN pool—either Client-MAC-hash or Load-balancing.</p> <ul style="list-style-type: none"> • None—The other selection methods are supported only by controller versions 9 and newer. Use None for controllers with older versions—for older controllers, only VLAN pools that are using the selection method None is supported. • Client-MAC-hash—Assign sessions to a VLAN based on the hash computed from the MAC address. Using this method guarantees that a client gets the same VLAN every time it connects, if the VLAN configuration does not change. • Load-balancing—Every controller keeps track of the number of sessions per VLAN in the Mobility Domain. When a session is assigned a VLAN pool, the session is assigned to the VLAN with the fewest sessions. <p>TIP: With client-MAC-hash and load-balancing selection methods, a VLAN is selected only if it has not reached the configured cap of sessions per VLAN. The default value is 0 (unlimited sessions). No more sessions are assigned to a VLAN once it has reached capacity.</p>
Members	List of VLANs in each VLAN pool.
Creation Time	Date and time when the profile was created.
Last Updated Time	Date and time when the profile was last modified.
User Name	The username of the person who created or modified the profile.

TIP: All columns might not be displayed. To show or hide fields listed in the table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

Creating a VLAN Pool

To create a VLAN pool, follow these directions:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View**, or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.

3. In the Tasks pane, expand **Wireless**, expand **Profiles**, and then click **VLAN Pool**.

The Manage VLAN Pool page appears, displaying the list of currently configured VLAN pools.

4. Click **Add**.

The Create VLAN Pool for Wireless page opens.

5. Provide a VLAN Pool Name up to 32 characters with no spaces or special characters. Provide an optional Description, and then choose one of these VLAN selection methods:

- **None** (default)—The other two selection methods are supported only by controller versions 9 and later. Use None for earlier-version controllers—these controllers support only VLAN pools that use the selection method None.
- **Client MAC Hash**—Assign sessions to a VLAN based on the hash computed from the MAC address. Using this method guarantees that a client gets the same VLAN every time it connects, if the VLAN configuration does not change.
- **Load Balancing**—Every controller keeps track of the number of sessions per VLAN in the Mobility Domain. When a session is assigned a VLAN pool, the session is assigned to the VLAN with the fewest sessions.

TIP: With Client MAC Hash and Load Balancing selection methods, a VLAN is selected only if it has not reached the configured cap of sessions per VLAN. The default value is 0 (unlimited sessions). No more sessions are assigned to a VLAN once it has reached capacity.

6. Add existing VLANs to the Members List of the VLAN pool by following these steps:
 - a. Click **Select** under **VLAN Members**.

The Select VLAN Profile for VLAN Pool window opens with a list of existing VLAN Pools..

- b. Select any VLANs from the list by adding a check mark, and then click **OK**.

The selected VLANs are added to the pool and appear in the VLAN Members list.

7. Optionally, add VLANs that are not yet controlled by Network Director. Add them to the VLAN pool by following these steps:

- a. Under Members, click **Add New**.

The Add New Member window opens.

- b. Type a name for a VLAN that does not yet belong to Network Director, type a Capacity, and then click **Add**.

The VLAN name is added to the pool and appears in the VLAN Members list.

8. Click **Done**.

The VLAN pool is created and added to the list of VLAN pools on the Manage VLAN Pool page.

What To Do Next

VLAN pools are assigned to local users, user groups, MAC users and MAC user groups for a given controller by associating them with an Authorization profile—see [“Creating and Managing Wireless Authorization Profiles” on page 395](#) and [“Assigning Wireless Authorization Profiles to Controllers” on page 404](#).

RELATED DOCUMENTATION

[Creating and Managing VLAN Profiles | 502](#)

[Creating and Managing Wireless Authorization Profiles | 395](#)

[Assigning Wireless Authorization Profiles to Controllers | 404](#)

[Network Director Documentation home page](#)

Configuring Firewall Filters (ACLs)

IN THIS CHAPTER

- [Understanding Filter Profiles | 540](#)
- [Creating and Managing Wired Filter Profiles | 542](#)
- [Creating and Managing Wireless Filter Profiles | 598](#)
- [Assigning a Wireless Filter Profile to Controllers | 607](#)

Understanding Filter Profiles

Filter profiles are a set of rules that define whether to accept or discard packets that are transiting on an interface on a Juniper Networks EX Series Ethernet Switch or on a radio connected to a Wireless LAN Controller. You configure Filter profiles to determine whether to accept or decline traffic before it enters or exits a port or a radio to which the Filter profile is applied to.

A Filter profile must contain at least one term. Each term consists of the following components:

- **Match conditions**—Specify the values or fields that the packet must contain. You can define various match conditions, depending on the device for which you are defining these conditions. For example, for EX Series switches, you can specify a match condition based on the IP source address field, IP destination address field, Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) source port field, IP protocol field, Internet Control Message Protocol (ICMP) packet type, TCP flags, interfaces, and so on. For a wireless device, the match conditions can be based on the source and destination IP or MAC address, and EtherTypes.
- **Action**—Specifies what to do if a packet matches the match conditions. Possible actions are to accept or discard the packet or to send the packet to a specific virtual routing interface. In addition, packets can be counted to collect statistical information. If no action is specified for a term, the default action is to accept the packet.
- **Action modifier**—Specifies one or more actions for the switch if a packet matches the match conditions. You can specify action modifiers such as the loss priority, policer details, and forwarding class, depending on the type of device on which you are creating the Filter profile.

The maximum number of terms allowed per Filter profile for EX Series switches is:

- 512 for EX2200 switches
- 1,436 for EX3300 switches

NOTE: On EX3300 switches, if you add and delete filters with a large number of terms (on the order of 1000 or more) in the same commit operation, not all the filters are installed. You must add filters in one commit operation, and delete filters in a separate commit operation.

- 7,042 for EX3200 and EX4200 switches—as allocated by the dynamic allocation of ternary content addressable memory (TCAM) for firewall filters.
- 1,200 for EX4500 and EX4550 switches
- 1,400 for EX6200 switches
- 32,768 for EX8200 switches

NOTE: The on-demand dynamic allocation of the shared space TCAM in EX8200 switches is achieved by assigning free space blocks to firewall filters. Firewall filters are categorized into two different pools. Port and VLAN filters are pooled together (the memory threshold for this pool is 22K) while router firewall filters are pooled separately (the threshold for this pool is 32K). The assignment happens based on the filter pool type. You can share free space blocks only among the firewall filters belonging to the same filter pool type. An error message is generated if you attempt to configure a firewall filter beyond the TCAM threshold.

The Manage Filter Profiles page enables you create, modify, view, and delete Filter profiles.

RELATED DOCUMENTATION

[Creating and Managing Wired Filter Profiles | 542](#)

[Network Director Documentation home page](#)

Creating and Managing Wired Filter Profiles

IN THIS SECTION

- [Managing Wired Filter Profiles | 542](#)
- [Creating a Wired Filter Profile | 543](#)
- [Specifying Settings for an EX Series Switch Filter Profile | 544](#)
- [Specifying Settings for a Campus Switching ELS Switch Filter Profile | 556](#)
- [Specifying Settings for Creating a Data Center Switching Non-ELS Filter Profile | 570](#)
- [Specifying Settings for a Data Center Switching ELS Filter Profile | 583](#)
- [What to Do Next | 597](#)

Filter profiles are sets of rules that determine whether to accept or discard packets transiting on either a switch or wireless radio interface.

Use the Manage Filter Profiles page to create new wired Filter profiles and manage existing Filter profiles.

This topic describes:

Managing Wired Filter Profiles

From the Manage Filter Profiles page, you can:

- Create a new wired Filter profile by clicking **Add**. For directions, see [“Creating a Wired Filter Profile” on page 543](#).
- Modify an existing wired Filter profile by selecting it and clicking **Edit**.
- View information about a wired Filter profile, including the associated interfaces, by either clicking the profile name or by selecting the profile and clicking **Details**.
- Delete a wired Filter profile by selecting the profile and clicking **Delete**.

TIP: You cannot delete profiles that are in use—that is, profiles assigned to objects or used by other profiles. To see the current assignments for a profile, select the profile and click **Details**.

- Clone a wired Filter profile by selecting a profile and clicking **Clone**.

Table 120 describes the information provided about wired Filter profiles on the Manage Filter Profiles page. This page lists all Filter profiles defined for your network, regardless of the scope you selected in the network view.

Table 120: Manage Wired Filter Profile Fields


Field	Description
Profile Name	Name given to the profile when the profile was created.
Family Type	The device family on which the profile was created: Switching (EX) , Campus Switching ELS , or Data Center Switching Non-ELS .
Description	Description of the profile entered when the profile was created. TIP: To display the entire description, you might need to resize the Description column by clicking the column border in the heading and dragging it.
Creation Time	Date and time when the profile was created.
Last Updated Time	Date and time when the profile was last modified.
User Name	The username of the user who created or modified the profile.

TIP: All columns might not be displayed. To show or hide fields in the table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

Creating a Wired Filter Profile

To create a wired Filter profile, you must provide a filter name and configure at least one term. A term is a collection of one or more match conditions, and actions that the system takes when match conditions are met. A term must have at least one match condition.

To create a wired Filter profile:

1. Click  in the Network Director banner.
2. Under Views, select one of these options: **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

3. From the Tasks pane, expand **Wired**, expand **System**, and then select **Filter**.

4. Click **Add** to add a new profile.

Network Director opens the Device Family Chooser window.

5. From the Device Family Chooser, select the wired device family for which you want to create a profile. The available device families are **Switching (EX)**, **Campus Switching ELS** (Enhanced Layer 2 Software), and **Data Center Switching**.

6. Click **OK**.

The Create Filter Profile wizard for the selected device family is displayed.

7. Specify the filter settings by following these directions:

- For EX Series switches, specify the settings as described in both the online help and in [“Specifying Settings for an EX Series Switch Filter Profile” on page 544](#).
- For Campus Switching ELS, specify the settings as described in both the online help and in [“Specifying Settings for a Campus Switching ELS Switch Filter Profile” on page 556](#).
- For Data Center Switching Non-ELS, specify the settings as described in both the online help and in [“Specifying Settings for Creating a Data Center Switching Non-ELS Filter Profile” on page 570](#).
- For Data Center Switching ELS, specify the settings as described in both the online help and in [“Specifying Settings for a Data Center Switching ELS Filter Profile” on page 583](#).

8. Click **Done** to save the Filter profile.

The system saves the Filter profile and displays the Manage Filter Profiles page. Your new or modified Filter profile is listed in the table of Filter profiles.

Specifying Settings for an EX Series Switch Filter Profile

A Filter profile must have at least one term in it. Each term has one filtering function. For example, if a term is evaluating the source of packets, then that term cannot also evaluate the protocols used by the packets. Some switch models do accommodate multiple terms in one filter. When you have more than one term in a filter, the ordering of the terms is important. The system evaluates multiple filter terms as follows:

- The packet is evaluated against the first term's conditions. If the packet matches all of the conditions in that term, the action specified for that condition is taken and evaluation ends. Subsequent terms in the filter are not evaluated.
- If a packet does not match all conditions in the first term, the packet is then evaluated against the conditions in the second term. This process continues until either the packet matches all conditions in a term or there are no more terms in the filter. Whenever a match occurs, the term's corresponding action is taken and evaluation ends—subsequent terms in the filter are not evaluated.
- If a packet passes through all the terms in the filter without a match, the packet is discarded.

To configure a Filter profile for EX Series switches:

1. Specify a filter name and description for the Filter profile.
2. Select the switch filter family for which you want to create the profile:
 - If you want to create a Layer 2 based filter, select **Ethernet switching**.
 - If you want to create a Layer 3 based filter for IPv4, select **INET**.
 - If you want to create a Layer 3 based filter for IPv6, select **INET6**.
3. Under Terms, click **Add** to add one or more terms with match condition(s) for this filter.

The Create Term window opens, displaying a section for each type of term you can create, Source and Destination Parameters, Protocols, DSCP Settings, TCP Settings, and ICMP Settings. The Action section applies to all of those types.

NOTE: The order of the terms within a Filter profile configuration is important. Packets are tested against each term in the order in which terms are listed.

4. Enter a name for the filter term.
5. Specify the match condition(s) for the filter term as described in [Table 121](#). Required settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 121: Create Term Fields for EX Switching

Task	Description
------	-------------

Source and Destination Parameters

You can specify match conditions for either packets' origin (source) or packets' destination, or both. You are indicating the location of the filtering here—either specifying that packets that originate at a specific place (source) will be filtered or packets destined for a specific location (destination) will be filtered. You can have multiple sources and destinations for one filter term.

Table 121: Create Term Fields for EX Switching (continued)

Task	Description
Add Source Parameters and Destination Parameters	<p>To add source and destination parameters to the named filter term:</p> <ol style="list-style-type: none"> Click Add to the right of the Destination Parameters list. The Add Source/Destination Parameter window appears. Select either Source (default) or Destination from the Add Source/Destination Parameter page. Select one of following available Parameter Types from the Add Source/Destination Parameter page and provide the corresponding information: TIP: Available parameter types vary. <ul style="list-style-type: none"> ● IP Address—also provide the IP address of the source or destination device ● MAC Address—also provide the MAC address of the source or destination device ● Port—also provide the port type of the source or destination port. Select either AFS (Andrew File System), BGP (Border Gateway Protocol), BIFF (UNIX mail notification), Bootpc (bootstrap protocol client), Bootps, Cmd, CVS pserver, DHCP, Domain, EK login, EK shell, EXEC, Finger (protocol), or FTP. <p>NOTE: If you selected Port as the parameter and do not find the type of port that you want to add from the Port list, then select Other and enter a port number.</p> Click OK The parameter term is added to the appropriate list, either Source Parameters or Destination Parameters.

Protocols and EtherTypes

For either INET family, you can apply a filter term based on protocols being used by packets. For the Ethernet-switching family, you can apply a filter term based on either the protocols being used by packets or on the EtherTypes being used by packets. EtherType indicates a protocol that is encapsulated in the payload of an Ethernet Frame. Expand the Protocols section to see the configuration.

Table 121: Create Term Fields for EX Switching (continued)

Task	Description
Add a Protocol Match Condition (Ethernet-switching family or INET family)	<p>To add a protocol match condition to the named filter term:</p> <ol style="list-style-type: none"> Expand the Protocols and EtherTypes section. Click Add under Protocols. The Select Protocols window opens, displaying a list of protocols. From the list of protocols, select one or more. The options are AH, DSTOPTS, EGP, ESP, Fragment, GRE, Hop-by-hop, ICMP, ICMP6, IPIP, IPv6, No-text-header, OSPF, PIM, Routing, RSVP, SCTP, TCP, UDP, and VRRP. Click OK. The protocols are added to the Protocols list.
Add an EtherType Match Condition (Ethernet-switching family)	<p>To add an EtherType match condition to the named filter Ethernet-switching family term:</p> <ol style="list-style-type: none"> Expand the Protocols and EtherTypes section. Click Add under EtherTypes. The Select EtherTypes window opens, displaying a list of protocols. From the list of EtherTypes, select one or more. The options are AARP, AppleTalk, ARP, IPv4, IPv6, MPLS multicast, MPLS unicast, OAM, PPP, PPPOE discovery, PPPOE session, and SNA. Click OK. The EtherTypes are added to the EtherTypes list.

DSCP Settings

Expand this section to see the DSCP term settings. DiffServ is a simple mechanism for classifying and managing network traffic and providing quality-of-service (QoS) on IP networks. DiffServ can, for example, be used to apply low-latency to critical network traffic such as voice or streaming media while providing simple best-effort service to non-critical services such as Web traffic. Here, you can apply a filter term based on the Differentiated Services code point (DSCP) which is a field in IPv4 and IPv6 headers.

NOTE: With IPv6 packets, the DS field and ECN field replace the IPv4 TOS field.

Table 121: Create Term Fields for EX Switching (continued)

Task	Description
Add a DSCP Match Condition	<p>To add a DSCP match condition to the named filter term:</p> <p>NOTE: A DSCP IP match condition and a precedence match condition cannot be both specified for the same term.</p> <p>a. Click Add in the DSCP section to see a list of match conditions.</p> <p>The Select DSCP list appears.</p> <p>b. Select one or more of the following DSCP types from the list:</p> <ul style="list-style-type: none"> ● AF11—Assured forwarding class 1, low drop precedence ● AF12—Assured forwarding class 1, medium drop precedence ● AF21—Assured forwarding class 2, low drop precedence ● AF22—Assured forwarding class 2, medium drop precedence ● AF23—Assured forwarding class 2, high drop precedence ● AF31—Assured forwarding class 3, low drop precedence ● AF32—Assured forwarding class 3, medium drop precedence ● AF33—Assured forwarding class 3, high drop precedence ● AF41—Assured forwarding class 4, low drop precedence ● AF42—Assured forwarding class 4, medium drop precedence ● AF43—Assured forwarding class 4, high drop precedence ● BE—Best effort (default) ● EF—Expedited forwarding ● CS0—Class selector 0 ● CS1—Class selector 1 ● CS2—Class selector 2 ● CS3—Class selector 3 ● CS4—Class selector 4 ● CS5—Class selector 5 ● CS6—Class selector 6 ● CS7—Class selector 7 <p>c. Click OK.</p> <p>The DSCP code term for the named filter is added to the DSCP list.</p>

Table 121: Create Term Fields for EX Switching (continued)

Task	Description
Add a Precedence match condition	<p>You can apply an IP precedence match condition to the named term. With IP precedence, a device prioritizes traffic by class first. Then it differentiates and prioritizes same-class traffic.</p> <p>NOTE: The two match conditions IP Precedence and DSCP cannot be simultaneously applied to a term.</p> <p>To apply an IP precedence value match condition to the named term:</p> <ol style="list-style-type: none"> a. Click Add in the Precedence section. The Select Precedence list appears. b. Select one of the following precedence settings from the list: Routine (0 or lowest, also called Best Effort), Priority (1), Immediate (2), Flash (3, mainly used for voice signaling or for video), Flash-override (4), Critical-ECP (5, mainly used for voice RTP), Internet-control (6, used for IP routing protocols), or Net-control (7 or highest, used for link layer and routing protocol keep alive). c. Click OK. The precedence match condition is added to the named term, and the condition is listed in the Precedence list.
<h3>TCP Settings</h3>	
<p>Expand this section to see the TCP term settings. The Transmission Control Protocol (TCP) is the most common core protocol of the Internet protocol suite (IP). TCP provides reliable, ordered, error-checked delivery of a stream of octets between programs running on computers connected to the Internet or an intranet. You can use the TCP initial flag for a match condition.</p>	
Enable TCP Initial flag match condition	<p>Select to use the TCP initial flag for a match condition. The TCP flags option becomes unavailable as a result.</p>

Table 121: Create Term Fields for EX Switching (continued)

Task	Description
Enable other TCP flag match conditions	<p>If you are not using the TCP initial flag for a match condition, select one of the TCP flags from the list—RST, ACK, SYN, Urgent, Push, FIN, None. These flags have the following meaning:</p> <ul style="list-style-type: none"> • RST—Reset flag indicates that the TCP connection will be reset. • ACK—Third step in TCP three-way handshake for connection. In response to a server's SYN-ACK, the client replies with an ACK. • SYN—First step in TCP three-way handshake for connection. The active open is performed by the client sending a SYN to the server. • Urgent—If the URG flag is set, then the 16-bit field is an offset from the sequence number indicating the last urgent data byte. • Push—Push flags request that buffered data to the receiving application be sent now. • FIN—The final flag indicates that no more data will be sent.

ICMP Settings

Expand the ICMP Settings section to select an ICMP code value for the filter item's match condition. The Internet Control Message Protocol (ICMP) is one of the core IP protocols used by operating systems of networked computers to send error messages. ICMP can also be used to relay query messages.

Add an ICMP Code match condition	<p>To apply an ICMP code match condition to the named term:</p> <ol style="list-style-type: none"> a. Click Add in the ICMP Codes section. The Select ICMP Code list appears. b. Select one or more ICMP codes from the list. These codes vary, depending on the Filter Family you selected. c. Click OK. The ICMP code match condition is listed in the ICMP Code list and added to the named term. <p>NOTE: ICMP code specifies more specific information than ICMP type. Because the value's meaning depends upon the associated ICMP type, you must specify an ICMP type along with an ICMP code. The keywords are grouped by the ICMP type with which they are associated.</p>
----------------------------------	---

Table 121: Create Term Fields for EX Switching (continued)

Task	Description
Add an ICMP Type match condition	<p>NOTE: ICMP type specifies the ICMP packet type field. Typically, you specify this match condition in conjunction with the protocol match condition to determine which protocol is being used on the port.</p> <p>ICMP code specifies more specific information than ICMP type. Because the value's meaning depends upon the associated ICMP type, you must specify an ICMP type along with the ICMP code. The keywords are grouped by the ICMP type with which they are associated.</p> <p>To apply an ICMP type match condition to the named term:</p> <ol style="list-style-type: none"> Click Add in the ICMP Type section. The Select ICMP Type list appears. Select one or more ICMP types from the list. Options vary, depending on which Filter Family you selected. Click OK. The ICMP type match condition is listed in the ICMP Type list and is added to the named term.
Action	
Select the action that the system performs on an IP packet if all match conditions that you specified above are met. Possible actions are Discard and Accept. The default action is to discard a packet that matches the filter term's conditions.	
Action	<p>Select either Discard or Accept to indicate what the filter term does with a packet when a match is made.</p> <p>NOTE: The remaining fields in this section are enabled only if you select Accept as the action.</p>
Counter Name	When Accept is the action, specify a counter name.
Loss Priority	<p>When Accept is the action, specify the packet loss priority, Low, High, or None.</p> <p>NOTE: Forwarding class and loss priority must be specified together for the same term.</p>

Table 121: Create Term Fields for EX Switching (continued)

Task	Description
Policer	<p>When you create a Filter profile, you can specify a policer action for any term or terms within the filter. Policing, or rate limiting, enables you to limit the amount of traffic that passes into or out of an interface. All traffic that matches a term that contains a policer action goes through the policer that the term references.</p> <p>You have two options with a policer. You can specify that an existing policer be used for the packet that matches the match condition. Or, you can create a new policer for the packet that matches the match condition.</p> <p>To select a policer from an existing list of policers, click Select. The Select Policer page appears. Select the policer that you want to use for the term and click OK. The system displays the selected policer in the Policer field in the Create Term page.</p>

Table 121: Create Term Fields for EX Switching (continued)

Task	Description
	<p>To create a new policer:</p> <ol style="list-style-type: none"> a. Click Create. The Create Policer page appears. b. Type a name for the policer—you can use this policer again in the future. c. Select a policer type from the list, either a single-rate-two-color policer, or a three-color-policer. The type of policer that you select here affects the rest of the configurations available for the policer. If you selected a three-color-policer, then also select a rate for it, either single-rate or two-rate. <ul style="list-style-type: none"> • Single-rate two-color—A two-color policer (sometimes called simply <i>policer</i>) meters the traffic stream and classifies packets into two categories of packet loss priority (PLP) according to a configured bandwidth and burst-size limit. You can mark packets that exceed the bandwidth and burst-size limit or simply discard them. A two-color policer is most useful for metering traffic at the port (physical interface) level. • Single-Rate Three-color—This type of policer is defined in RFC 2697, A Single Rate Three Color Marker, as part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on the configured committed information rate (CIR), committed burst size (CBS), and the excess burst size (EBS). Traffic is marked as belonging to one of three categories (green, yellow, or red) based on whether the packets are arriving at rates that are below the CBS (green), exceed the CBS but not the EBS (yellow), or exceed the EBS (red). A single-rate three-color policer is most useful when a service is structured according to packet size and not according to peak arrival rate. • Two-rate three-color—This type of policer is defined in RFC 2698, A Two Rate Three Color Marker, as part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on the configured CIR and the peak information rate (PIR), along with their associated burst sizes; the CBS, and the peak burst size (PBS). Traffic is marked as belonging to one of three categories (green, yellow, or red) based on packets are arriving at rates that are below the CIR (green), exceed the CIR but not the PIR (yellow), or exceed the PIR (red). A two-rate three-color policer is most useful when a service is structured according to arrival rates and not to packet size. <p>NOTE: The system displays and hides various fields in the Create Policer page depending on the type of policer that you want to create.</p>

Table 121: Create Term Fields for EX Switching (continued)

Task	Description
	<p>d. Configure these fields for a single-rate-two-color policer:</p> <ul style="list-style-type: none"> ● Bandwidth Limit—Specify the traffic rate in bits per second, 1000 through 102,300,000,000 (102.3g) bps. ● Burst Size Limit—Specify the maximum number of bytes allowed for incoming packets to burst above the peak information rate (PIR) and still be marked with medium-high packet loss priority (yellow). Packets that exceed the peak burst size (PBS) are marked with high packet loss priority (red). The range is 1500 through 100,000,000,000 bytes. ● Action—Select either Discard or None. ● Loss Priority—Select either High or None. <p>e. Configure these fields for a single-rate-three-color policer:</p> <ul style="list-style-type: none"> ● Committed Information Rate—Specify the guaranteed bandwidth (in bits per second) under normal line conditions and the average rate up to which packets are marked with low packet loss priority (green). The range is 32,000 through 40,000,000,000 bps. ● Committed Burst Size—Specify the maximum number of bytes allowed for incoming packets to burst above the committed information rate (CIR) and still be marked with low packet loss priority (green). The range is 1500 through 100,000,000,000 bytes. ● Excess Burst Size—Specify the maximum number of bytes allowed for incoming packets to burst above the committed information rate and still be marked with medium-high packet loss priority (yellow). Packets that exceed the excess burst size (EBS) are marked with high packet loss priority (red). The range is 1500 through 100,000,000,000 bytes. ● Color Mode—Select the way the preclassified packets are to be metered: <ul style="list-style-type: none"> ● Color-aware—The local switch can assign a higher packet loss priority but cannot assign a lower packet loss priority. ● Color-blind—The local switch ignores the preclassification of packets and can assign a higher or lower packet loss priority. ● None—The preclassified packets are not metered. ● Action—Options are Discard and None. ● Loss Priority—Options are High and None.

Table 121: Create Term Fields for EX Switching (continued)

Task	Description
	<p>f. Configure these fields for a three-color two-rate policer:</p> <ul style="list-style-type: none"> ● Committed Information Rate—Specify the guaranteed bandwidth (in bits per second) under normal line conditions and the average rate up to which packets are marked with low packet loss priority (green). The range is 32,000 through 40,000,000,000 bps. ● Committed Burst Size—Specify the maximum number of bytes allowed for incoming packets to burst above the committed information rate (CIR) and still be marked with low packet loss priority (green). The range is 1500 through 100,000,000,000 bytes. ● Peak Burst Size—Specify the maximum number of bytes allowed for incoming packets to burst above the peak information rate (PIR) and still be marked with medium-high packet loss priority (yellow). Packets that exceed the peak burst size (PBS) are marked with high packet loss priority (red). The range is 1500 through 100,000,000,000 bytes. ● Peak Information Rate—Specify the maximum achievable rate in bits per second. Packets that exceed the peak information rate (PIR) are marked with high packet loss priority (red). You can configure a discard action for packets that exceed the PIR. The range is 32,000 through 40,000,000,000 bps. ● Color Mode—Select the way the preclassified packets are to be metered: <ul style="list-style-type: none"> ● Color-aware—The local switch can assign a higher packet loss priority but cannot assign a lower packet loss priority. ● Color-blind—The local switch ignores the preclassification of packets and can assign a higher or lower packet loss priority. ● None—The preclassified packets are not metered. ● Action—Options are Discard and None. ● Loss Priority—Options are High and None. <p>g. Click OK.</p> <p>The policer is added to the list of applied policers and the list of available policers.</p>

Table 121: Create Term Fields for EX Switching (continued)

Task	Description
Forwarding Class	<p>When Accept is the action, specify the forwarding class (or output queue) that is to be used for the packet that matches the match condition. You can create a new forwarding class or select from a list of available forwarding classes.</p> <p>To select a forwarding class from an existing list of classes, click Select. The Select Forwarding Class page appears. Select the forwarding class that you want to use for the packet and click OK. The system displays the selected forwarding class in the Forwarding Class field in the Create Term page.</p> <hr/> <p>To create a new forwarding class:</p> <ol style="list-style-type: none"> a. Click Create. The Create Forwarding Class page appears. b. Type a name for the forwarding class—you can use this forwarding class again in the future. c. Select a queue number from the list, and then click OK. The system creates a new forwarding class and displays it in the Forwarding Class field in the Create Term page.

Click **OK** to save the term and return to the Create Filter Profile page.

Specifying Settings for a Campus Switching ELS Switch Filter Profile

A Filter profile must have at least one term in it. Each term has one filtering function. For example, if a term is evaluating the source of packets, then that term cannot also evaluate the protocols used by the packets. Some switch models accommodate multiple terms in one filter. When you have more than one term in a filter, the ordering of the terms is important. The system evaluates multiple filter terms as follows:

- The packet is evaluated against the first term's conditions. If the packet matches all of the conditions in that term, the corresponding action for that condition is taken and evaluation ends. Subsequent terms in the filter are not evaluated.
- If the packet does not match all conditions in the first term, the packet is evaluated against the conditions in the second term. This process continues until either the packet matches all the conditions in one of the subsequent terms or there are no more terms in the filter. If a match is found, the action specified in the Action section of the matched term is taken and the evaluation ends. Subsequent terms in the filter are not evaluated.

- The term conditions for protocol, EtherType, DSCP, precedence, ICMP code and ICMP type must all be either match conditions or except conditions.
- If a packet passes through all the terms in the filter without a match, the packet is discarded.

To configure a Filter profile for Campus switching ELS:

1. Specify a filter name and description for the Filter profile.
2. Select the switch filter family for which you want to create the profile:
 - If you want to create a Layer 2 based filter, select **Ethernet switching**.
 - If you want to create a Layer 3 based filter for IPv4, select **INET**.
 - If you want to create a Layer 3 based filter for IPv6, select **INET6**.
3. Under Terms, click **Add** to add one or more terms with match condition(s) to the named filter. You need at least one term for this filter.

The Create Term window opens.

NOTE: The order of the terms within a Filter profile configuration is important. Packets are tested against each term in the order in which the terms are listed.

4. Enter a name for the filter term.
5. Specify the match condition(s) for the filter term as described in [Table 122](#). Required settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 122: Create Term Fields for Campus Switching ELS

Field	Description
-------	-------------

Source and Destination Parameters

You can specify match conditions based on the packets' origin (source) or the packets' destination, or both. You are indicating the location of the filtering here—either specifying that packets that originate at a specific place (source) will be filtered or packets destined for a specific location (destination) will be filtered. You can have multiple sources and destinations for one filter.

Table 122: Create Term Fields for Campus Switching ELS (continued)

Field	Description
Source Parameters and Destination Parameters	<p>To add source and destination parameters to the named filter term:</p> <ol style="list-style-type: none"> Click Add to the right of the Destination Parameters lists. The Add Source/Destination Parameter window opens. Select either Source (default) or Destination from the Add Source/Destination Parameter window. Select one of following available Parameter Types from the Add Source/Destination Parameter page and provide the corresponding information: <ul style="list-style-type: none"> ● IP Address—Provide the IP address of the source or destination device. ● MAC Address—Provide a MAC address. ● Port—Provide the port type of the source or destination port. Select either AFS (Andrew File System), BGP (Border Gateway Protocol), BIFF (UNIX mail notification), Bootpc (bootstrap protocol client), Bootps, Cmd, CVS pserver, DHCP, Domain, EK login, EK shell, EXEC, Finger protocol, FTP, FTP data, HTTP, HTTPS, Ident protocol, IMAP (Internet Message Access protocol), Kerberos-sec (Kerberos security), Klogin forwarding, Kpasswd command, KRB-prop (Kerberos database propagation), Krbupdate (Kerberos database update), Kshell (Kerberos rsh), LDAP, Login (UNIX rlogin), Mobilip-agent (Mobile IP agent), Mobilip-mn (Mobile IP MN), MSDP (Multicast Source Discovery Protocol), NetBIOS dgm, NetBIOS-ns (NetBIOS name service), NetBIOS-ssn (NetBIOS session service), NFSD, NNTP (Network News Transport Protocol), Ntalk, NTP (Network Time Protocol), POP3 (Post Office Protocol3), PPTP, Printer, RADacct (RADIUS accounting), RADIUS, RIP, RKINIT (Kerberos remote kinit), SMTP, SNMP trap, SNPP, SUNRPC, Syslog, TACACS, TACACS-ds, Talk (UNIX Talk), Telnet, TFTP, Timed (UNIX time daemon), Who (UNIX rwho), XDMCP (X Display Manager Control Protocol), Zephyr-clt (Zephyr serv-hm connection), Zephyr-hm (Zephyr hostmanager), Zephyr-srv (Zephyr server), or Other. <p>NOTE: If you selected Port as the parameter and do not find the type of port that you want to add from the Port list, then select other and enter a port number.</p> To select any other source/destination than the one indicated, enable Except. TIP: You cannot indicate both matching and except for a parameter. Click OK The parameter term is added to the appropriate list, either Source Parameters or Destination Parameters.

Table 122: Create Term Fields for Campus Switching ELS (continued)

Field	Description
<p>Protocols and EtherTypes</p>	
<p>Depending on the Filter Family you selected, you can sometimes apply a filter term based on either protocols being used by packets or on EtherTypes being used by packets. Recognized protocols are listed where applicable. Recognized EtherTypes, which indicate the protocol that is encapsulated in the payload of an Ethernet Frame, are also listed where applicable.</p>	
<p>Protocols (apply to Ethernet and INET filter families)</p>	<p>To add a protocol match condition to the named filter term:</p> <ol style="list-style-type: none"> Expand the Protocols and EtherTypes section. Click Add under Protocols. The Select Protocols window opens, displaying a list of protocols. From the list of protocols, select one or more. The options are AH, DSTOPTS, EGP, ESP, Fragment, GRE, Hop-by-hop, ICMP, IPIP, IPv6, No-text-header, OSPF, PIM, Routing, RSVP, SCTP, TCP, UDP, and VRRP. To make the filter exclude the specified protocol, select Except. NOTE: The term conditions for protocol, EtherType, DSCP, precedence, ICMP code and ICMP type must all be either match conditions or except conditions. Click OK. The protocols are added to the Protocols list.

Table 122: Create Term Fields for Campus Switching ELS (continued)

Field	Description
EtherTypes (apply to Ethernet filter family)	<p>To add an EtherTypes match condition to the named filter term:</p> <ol style="list-style-type: none"> Expand the Protocols and EtherTypes section. Click Add under EtherTypes. The Select EtherTypes window opens, displaying a list of protocols. From the list of EtherTypes, select one or more. The options are AARP, AppleTalk, ARP, IPV4, MPLS multicast, MPLS unicast, OAM, PPP, PPPOE discovery, PPPOE session, and SNA. To make the filter exclude the specified EtherType, select Except. NOTE: Term values must all be either match conditions or all except conditions. Click OK. The EtherTypes are added to the EtherTypes list.

DSCP Settings

Expand the DSCP section to see the DSCP match settings. DiffServ is a simple mechanism for classifying and managing network traffic and providing quality-of-service (QoS) on IP networks. DiffServ can, for example, be used to apply low-latency to critical network traffic such as voice or streaming media while providing simple best-effort service to non-critical services such as Web traffic. Here, you can apply a filter term based on the Differentiated Services code point (DSCP) which is a field in IPv4 and IPv6 headers.

NOTE: With IPv6 packets, the DS field and ECN field replace the IPv4 TOS field.

Table 122: Create Term Fields for Campus Switching ELS (continued)

Field	Description
DSCP (Ethernet and INET filter families)	<p>To add a DSCP match condition to the named filter term:</p> <p>NOTE: A DSCP IP match condition and a precedence match condition cannot be both specified for the same term.</p> <p>a. Click Add in the DSCP section.</p> <p>The Select DSCP Match Condition list appears.</p> <p>b. Select one of the following DSCP types from the list:</p> <ul style="list-style-type: none"> ● AF11—Assured forwarding class 1, low drop precedence ● AF12—Assured forwarding class 1, medium drop precedence ● AF21—Assured forwarding class 2, low drop precedence ● AF22—Assured forwarding class 2, medium drop precedence ● AF23—Assured forwarding class 2, high drop precedence ● AF31—Assured forwarding class 3, low drop precedence ● AF32—Assured forwarding class 3, medium drop precedence ● AF33—Assured forwarding class 3, high drop precedence ● AF41—Assured forwarding class 4, low drop precedence ● AF42—Assured forwarding class 4, medium drop precedence ● AF43—Assured forwarding class 4, high drop precedence ● BE—Best effort (default) ● EF—Expedited forwarding ● CS0—Class selector 0 ● CS1—Class selector 1 ● CS2—Class selector 2 ● CS3—Class selector 3 ● CS4—Class selector 4 ● CS5—Class selector 5 ● CS6—Class selector 6 ● CS7—Class selector 7 <p>c. To make the filter exclude a specified DSCP type, select Except.</p> <p>NOTE: Term values must all be either match conditions or all of them need to be except conditions.</p> <p>d. Click OK.</p> <p>The DSCP code term for the named filter is added to the DSCP list.</p>

Table 122: Create Term Fields for Campus Switching ELS (continued)

Field	Description
Precedence for DSCP (Ethernet and INET filter families)	<p>You can apply an IP precedence match condition to the named term. With IP precedence, a device prioritizes traffic by class first. Then it differentiates and prioritizes same-class traffic.</p> <p>NOTE: The match conditions IP Precedence and DSCP cannot be simultaneously applied to a term.</p> <p>To apply an IP precedence value match condition to the named term:</p> <ol style="list-style-type: none"> Click Add in the Precedence section. The Select Precedence list appears. Select one of the following precedence settings from the list: Routine (0 or lowest, also called Best Effort), Priority (1), Immediate (2), Flash (3, mainly used for voice signaling or for video), Flash-override (4), Critical-ECP (5, mainly used for voice RTP), Internet-control (6, used for IP routing protocols), or Net-control (7 or highest, used for link layer and routing protocol keep alive). To make the filter exclude the specified IP precedence value, select Except. NOTE: Term values must all be either match conditions or all of them need to be except conditions. Click OK. The precedence match condition is added to the named term, and the condition is listed in the Precedence list.

TCP Settings

Expand this section to access the TCP settings. The Transmission Control Protocol (TCP) is the most common core protocol of the Internet protocol suite (IP). TCP provides reliable, ordered, error-checked delivery of a stream of octets between programs running on computers connected to the Internet or an intranet. You can use the TCP initial flag for a match condition.

Enable TCP Initial (all families)	<p>Select to use the TCP initial flag for an Ethernet, INET, or INET6 match condition.</p> <p>TIP: If you use the TCP initial flag for filtering, you cannot use any other TCP flag.</p>
-----------------------------------	---

Table 122: Create Term Fields for Campus Switching ELS (continued)

Field	Description
TCP Flags	<p>If you are not using the TCP initial flag for a match condition, you can select one of the TCP flags from the list for a match condition—RST, ACK, SYN, Urgent, Push, FIN, or None. These flags have the following meaning:</p> <ul style="list-style-type: none"> • RST—Reset flag indicates that the TCP connection will be reset. • ACK—Third step in TCP three-way handshake for connection. In response to a server's SYN-ACK, the client replies with an ACK. • SYN—First step in TCP three-way handshake for connection. The active open is performed by the client sending a SYN to the server. • Urgent—If the URG flag is set, then the 16-bit field is an offset from the sequence number indicating the last urgent data byte. • Push—Push flags request that buffered data to the receiving application be sent now. • FIN—The final flag indicates that no more data will be sent.

ICMP Settings

You can select the ICMP code value for the filter item's match condition—expand this section to access the ICMP settings. The Internet Control Message Protocol (ICMP) is one of the core IP protocols used by operating systems of networked computers to send error messages. ICMP can also be used to relay query messages.

ICMP Code	<p>To apply an ICMP code match condition to the named term:</p> <ol style="list-style-type: none"> a. Click Add in the ICMP Code section. The Select ICMP Code window appears. b. Select one or more ICMP codes from the list. These codes vary, depending on the Filter Family you selected. c. To make the filter exclude the specified ICMP code, select Except. NOTE: Term values must all be either match conditions or all of them need to be except conditions. d. Click OK. The ICMP code match condition is added to the named term, and the condition is listed in the ICMP Code list. You can now enable Except. <p>NOTE: An ICMP code specifies more specific information than an ICMP type. Because the value's meaning depends upon the associated ICMP type, you must specify an ICMP type along with ICMP code. The keywords are grouped by the ICMP type with which they are associated.</p>
-----------	---

Table 122: Create Term Fields for Campus Switching ELS (continued)

Field	Description
ICMP Type	<p>NOTE: ICMP type specifies the ICMP packet type field. Typically, you specify this match condition in conjunction with the protocol match condition to determine which protocol is being used on the port.</p> <p>ICMP code specifies more specific information than ICMP type. Because the value's meaning depends upon the associated ICMP type, you must specify ICMP type along with ICMP code. The keywords are grouped by the ICMP type with which they are associated.</p> <p>To apply an ICMP type match condition to the named term:</p> <ol style="list-style-type: none"> Click Add in the ICMP Type section. The Select ICMP Types window appears. Select one or more ICMP types from the list. These types vary, depending on the Filter Family selected. To make the filter exclude the specified ICMP type, select Except. NOTE: Term values must all be either match conditions or all of them need to be except conditions. Click OK. The ICMP type match condition is added to the named term, and the condition is listed in the ICMP Type list. You can now enable Except.
<p>Action</p>	
<p>Select the action that the system performs on an IP packet if all match conditions that you specified above are met. Possible actions are Discard and Accept. The default action is to discard packet that matches the filter term conditions.</p>	
Action	<p>Select either Discard or Accept to indicate what the filter term does with a packet when a match is made.</p> <p>NOTE: All other fields in this section are enabled only if you select Accept as the action.</p>
Counter Name	<p>When the action selected is accept, specify the maximum packet count for this filter, term, or policer.</p>

Table 122: Create Term Fields for Campus Switching ELS (continued)

Field	Description
Loss Priority	<p>When the action selected is accept, specify the packet loss priority, Low, High, Medium-low, Medium-high, or None.</p> <p>NOTE: Forwarding class and loss priority must be specified together for the same term.</p>

Table 122: Create Term Fields for Campus Switching ELS (continued)

Field	Description
Policer	<p>When you create a Filter profile with the action accept, you can specify a policer action for any term or terms within the filter. Policing, or rate limiting, enables you to limit the amount of traffic that passes into or out of an interface. All traffic that matches a term that contains a policer action goes through the policer that the term references.</p> <p>You have two options with a policer. You can specify that an existing policer be used for the packet that matches the match condition. Or, you can create a new policer for the packet that matches the match condition.</p> <hr/> <p>To select an existing policer:</p> <ol style="list-style-type: none">a. Click Select. The Select Policer page appears.b. Click OK. The policer is added to the list of applied policers. <hr/>

Table 122: Create Term Fields for Campus Switching ELS (continued)

Field	Description
	<p>To create a new policer:</p> <ol style="list-style-type: none"> a. Click Create. The Create Policer page appears. b. Type a name for the policer—you can use this policer again in the future. c. Select a policer type from the list, either a single-rate-two-color policer, or a three-color-policer. The type of policer that you select here affects the rest of the configurations available for the policer. If you selected a three-color-policer, then also select a rate for it, either single-rate or two-rate. <ul style="list-style-type: none"> • Single-rate two-color—A two-color policer (sometimes called simply <i>policer</i>) meters the traffic stream and classifies packets into two categories of packet loss priority (PLP) according to a configured bandwidth and burst-size limit. You can mark packets that exceed the bandwidth and burst-size limit or simply discard them. A two-color policer is most useful for metering traffic at the port (physical interface) level. • Single-Rate Three-color—This type of policer is defined in RFC 2697, A Single Rate Three Color Marker, as part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on the configured committed information rate (CIR), committed burst size (CBS), and the excess burst size (EBS). Traffic is marked as belonging to one of three categories (green, yellow, or red) based on whether the packets are arriving at rates that are below the CBS (green), exceed the CBS but not the EBS (yellow), or exceed the EBS (red). A single-rate three-color policer is most useful when a service is structured according to packet size and not according to peak arrival rate. • Two-rate three-color—This type of policer is defined in RFC 2698, A Two Rate Three Color Marker, as part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on the configured CIR and the peak information rate (PIR), along with their associated burst sizes; the CBS, and the peak burst size (PBS). Traffic is marked as belonging to one of three categories (green, yellow, or red) based on packets are arriving at rates that are below the CIR (green), exceed the CIR but not the PIR (yellow), or exceed the PIR (red). A two-rate three-color policer is most useful when a service is structured according to arrival rates and not to packet size. <p>NOTE: The system displays and hides various fields in the Create Policer page depending on the type of policer that you want to create.</p>

Table 122: Create Term Fields for Campus Switching ELS (continued)

Field	Description
	<p>d. Configure these fields for a single-rate-two-color policer:</p> <ul style="list-style-type: none"> ● Bandwidth Limit—Specify the traffic rate in bits per second, 1000 through 102,300,000,000 (102.3g) bps. ● Burst Size Limit—Specify the maximum number of bytes allowed for incoming packets to burst above the peak information rate (PIR) and still be marked with medium-high packet loss priority (yellow). Packets that exceed the peak burst size (PBS) are marked with high packet loss priority (red). The range is 1500 through 100,000,000,000 bytes. ● Action—Select either Discard or None. ● Loss Priority—Select either High or None. <p>e. Configure these fields for a single-rate-three-color policer:</p> <ul style="list-style-type: none"> ● Committed Information Rate—Specify the guaranteed bandwidth (in bits per second) under normal line conditions and the average rate up to which packets are marked with low packet loss priority (green). The range is 32,000 through 40,000,000,000 bps. ● Committed Burst Size—Specify the maximum number of bytes allowed for incoming packets to burst above the committed information rate (CIR) and still be marked with low packet loss priority (green). The range is 1500 through 100,000,000,000 bytes. ● Excess Burst Size—Specify the maximum number of bytes allowed for incoming packets to burst above the committed information rate and still be marked with medium-high packet loss priority (yellow). Packets that exceed the excess burst size (EBS) are marked with high packet loss priority (red). The range is 1500 through 100,000,000,000 bytes. ● Color Mode—Select the way the preclassified packets are to be metered: <ul style="list-style-type: none"> ● Color-aware—The local switch can assign a higher packet loss priority but cannot assign a lower packet loss priority. ● Color-blind—The local switch ignores the preclassification of packets and can assign a higher or lower packet loss priority. ● None—The preclassified packets are not metered. ● Action—Options are Discard and None. ● Loss Priority—Options are High and None.

Table 122: Create Term Fields for Campus Switching ELS (continued)

Field	Description
	<p>f. Configure these fields for a three-color two-rate policer:</p> <ul style="list-style-type: none"> ● Committed Information Rate—Specify the guaranteed bandwidth (in bits per second) under normal line conditions and the average rate up to which packets are marked with low packet loss priority (green). The range is 32,000 through 40,000,000,000 bps. ● Committed Burst Size—Specify the maximum number of bytes allowed for incoming packets to burst above the committed information rate (CIR) and still be marked with low packet loss priority (green). The range is 1500 through 100,000,000,000 bytes. ● Peak Burst Size—Specify the maximum number of bytes allowed for incoming packets to burst above the peak information rate (PIR) and still be marked with medium-high packet loss priority (yellow). Packets that exceed the peak burst size (PBS) are marked with high packet loss priority (red). The range is 1500 through 100,000,000,000 bytes. ● Peak Information Rate—Specify the maximum achievable rate in bits per second. Packets that exceed the peak information rate (PIR) are marked with high packet loss priority (red). You can configure a discard action for packets that exceed the PIR. The range is 32,000 through 40,000,000,000 bps. ● Color Mode—Select the way the preclassified packets are to be metered: <ul style="list-style-type: none"> ● Color-aware—The local switch can assign a higher packet loss priority but cannot assign a lower packet loss priority. ● Color-blind—The local switch ignores the preclassification of packets and can assign a higher or lower packet loss priority. ● None—The preclassified packets are not metered. ● Action—Options are Discard and None. ● Loss Priority—Options are High and None. <p>g. Click OK.</p> <p>The policer is added to the list of applied policers and the list of available policers.</p>

Table 122: Create Term Fields for Campus Switching ELS (continued)

Field	Description
Forwarding Class	<p>Specify the forwarding class (or output queue) that is to be used for the packet that matches the match condition. You can either select from a list of available forwarding classes or create a new forwarding class.</p> <p>To select a forwarding class from an existing list of classes, click Select. The Select Forwarding Class page appears. Select the forwarding class that you want to use for the packet and click OK. The system displays the selected forwarding class in the Forwarding Class field in the Create Term page.</p> <hr/> <p>To create a new forwarding class:</p> <ol style="list-style-type: none"> a. Click Create. The Create Forwarding Class page appears. b. Type a name for the forwarding class—you can use this forwarding class again in the future. c. Select a queue number from the list, and then click OK. The system creates a new forwarding class and displays it in the Forwarding Class field in the Create Term page.

6. Click **OK** to save the term and return to the Create Filter Profile page.

7. Click **Done**.

The new filter is added to the Manage Filter Profile list.

Specifying Settings for Creating a Data Center Switching Non-ELS Filter Profile

A Filter profile must have at least one term in it. Each term has one filtering function. For example, if a term is evaluating the source of packets, then that term cannot also evaluate the protocols used by the packets. Some switch models accommodate multiple terms in one filter. When you have more than one term in a filter, the ordering of the terms is important. The system evaluates multiple filter terms as follows:

- The packet is evaluated against the first term's conditions. If the packet matches all of the conditions in that term, the action specified for that condition is taken and evaluation ends. Subsequent terms in the filter are not evaluated.
- If the packet does not match all conditions in the first term, the packet is evaluated against the conditions in the second term. This process continues until either the packet matches all the conditions in one of the subsequent terms or there are no more terms in the filter. If a match is found, the action specified in the Action section of the matched term is taken and the evaluation ends. Subsequent terms in the filter are not evaluated.
- Term values (protocol, EtherType, DSCP, precedence, ICMP code and ICMP type) must all be either match conditions or all of them need to be except conditions.
- If a packet passes through all the terms in the filter without a match, the packet is discarded.

To configure a Filter profile:

1. Specify a filter name and description for the Filter profile.
2. Select the switch filter family for which you want to create the profile:
 - If you want to create a Layer 2 based filter, select **Ethernet switching**.
 - If you want to create a Layer 3 based filter for IPv4, select **INET**.
 - If you want to create a Layer 3 based filter for IPv6, select **INET6**.
3. Under Terms, click **Add** to add one or more terms with match condition(s) to the named filter.

The Create Term window opens.

NOTE: The order of the terms within a Filter profile configuration is important. Packets are tested against each term in the order in which the terms are listed.

4. Enter a name for the filter term.
5. Specify the match condition(s) for the filter term as described in [Table 123](#). Required settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 123: Create Term Fields for Data Center Switching Non-ELS

Field	Description
<p>Source Parameters and Destination Parameters</p>	<p>To add source and destination parameters to the named filter term:</p> <ol style="list-style-type: none"> Click Add to the right of the Destination Parameters lists. The Add Source/Destination Parameter page appears. Select either Source (default) or Destination from the Add Source/Destination Parameter page. Select one of following available Parameter Types from the Add Source/Destination Parameter page and provide the corresponding information: <ul style="list-style-type: none"> ● IP Address—Provide the IP address of the source or destination device ● MAC Address—MAC address of the source or destination device ● Port—Port type of the source or destination port. Select either AFS (Andrew File System), BGP (Border Gateway Protocol), BIFF (UNIX mail notification), Bootpc (bootstrap protocol client), Bootps, Cmd, CVS pserver, DHCP, Domain, EK login, EK shell, EXEC, Finger (protocol), or FTP. <p>NOTE: If you selected Port as the parameter and do not find the type of port that you want to add from the Port list, then select other and enter a port number.</p> Click OK The parameter term is added to the appropriate list, either Source Parameters or Destination Parameters.

Protocols and EtherTypes

You can apply a filter term that is based on either the protocols being used by packets or on the EtherTypes being used by packets. Protocols such as AH, DSTOPTS, EGP, ESP, FRAGMENT, GRE, HOP-BY-HOP, ICMP, ICMP6, IPIP, IPv6, no-text-header, OSPF, PIM, ROUTING, RSVP, SCTP, TCP, UDP, and VRRP are recognized. EtherType indicates the protocol that is encapsulated in the payload of an Ethernet Frame.

Table 123: Create Term Fields for Data Center Switching Non-ELS (continued)

Field	Description
Protocols	<p>To add a protocol match condition to the named filter term:</p> <ol style="list-style-type: none"> Expand the Protocols and EtherTypes section. Click Add under Protocols. The Select Protocols window opens, displaying a list of protocols. From the list of protocols, select one or more. The options are AH, DSTOPTS, EGP, ESP, fragment, GRE, Hop-by-hop, ICMP, IPIP, IPv6, No-text-header, OSPF, PIM, routing, RSVP, SCTP, TCP, UDP, and VRRP. Click OK. The protocols are added to the Protocols list.
EtherTypes	<p>To add an EtherTypes match condition to the named filter term:</p> <ol style="list-style-type: none"> Expand the Protocols and EtherTypes section. Click Add under EtherTypes. The Select EtherTypes window opens, displaying a list of protocols. From the list of EtherTypes, select one or more. The options are AARP, AppleTalk, ARP, IPV4, FCOE, FIP, MPLS multicast, MPLS unicast, OAM, PPP, PPPOE discovery, PPPOE session, SNA, and VLAN. Click OK. The protocols are added to the Protocols list.

DSCP Settings

DiffServ is a simple mechanism for classifying and managing network traffic and providing quality-of-service (QoS) on IP networks. DiffServ can, for example, be used to apply low-latency to critical network traffic such as voice or streaming media while providing simple best-effort service to non-critical services such as Web traffic. Here, you can apply a filter term based on the Differentiated Services code point (DSCP) which is a field in IPv4 and IPv6 headers.

NOTE: With IPv6 packets, the DS field and ECN field replace the IPv4 TOS field.

Table 123: Create Term Fields for Data Center Switching Non-ELS (continued)

Field	Description
DSCP	<p>To add a DSCP match condition to the named filter term:</p> <p>NOTE: A DSCP IP match condition and a precedence match condition cannot be both specified for the same term.</p> <ol style="list-style-type: none"> a. Expand the DSCP Settings section. b. Click Add in the DSCP section. The Select DSCP Match Condition list appears. c. Select one of the following DSCP types from the list: <ul style="list-style-type: none"> ● AF11—Assured forwarding class 1, low drop precedence ● AF12—Assured forwarding class 1, medium drop precedence ● AF21—Assured forwarding class 2, low drop precedence ● AF22—Assured forwarding class 2, medium drop precedence ● AF23—Assured forwarding class 2, high drop precedence ● AF31—Assured forwarding class 3, low drop precedence ● AF32—Assured forwarding class 3, medium drop precedence ● AF33—Assured forwarding class 3, high drop precedence ● AF41—Assured forwarding class 4, low drop precedence ● AF42—Assured forwarding class 4, medium drop precedence ● AF43—Assured forwarding class 4, high drop precedence ● BE—Best effort (default) ● EF—Expedited forwarding ● CS0—Class selector 0 ● CS1—Class selector 1 ● CS2—Class selector 2 ● CS3—Class selector 3 ● CS4—Class selector 4 ● CS5—Class selector 5 ● CS6—Class selector 6 ● CS7—Class selector 7 d. Click OK. The DSCP code term for the named filter is added to the DSCP list.

Table 123: Create Term Fields for Data Center Switching Non-ELS (continued)

Field	Description
Precedence	<p>You can apply an IP precedence match condition to the named term. With IP precedence, a device prioritizes traffic by class first. Then it differentiates and prioritizes same-class traffic.</p> <p>NOTE: The match conditions IP Precedence and DSCP cannot be simultaneously applied to a term.</p> <p>To apply an IP precedence value match condition to the named term:</p> <ol style="list-style-type: none"> a. Expand the DSCP Settings section. b. Click Add in the Precedence section. The Select Precedence list appears. c. Select one of the following precedence settings from the list: Routine (0 or lowest, also called Best Effort), Priority (1), Immediate (2), Flash (3, mainly used for voice signaling or for video), Flash-override (4), Critical-ECP (5, mainly used for voice RTP), Internet-control (6, used for IP routing protocols), or Net-control (7 or highest, used for link layer and routing protocol keep alive). d. Click OK. The precedence match condition is added to the named term, and the condition is listed in the Precedence list.

TCP Settings

The Transmission Control Protocol (TCP) is the most common core protocol of the Internet protocol suite (IP). TCP provides reliable, ordered, error-checked delivery of a stream of octets between programs running on computers connected to the Internet or an intranet. You can use the TCP initial flag for a match condition.

Enable TCP Initial	Select to use the TCP initial flag for a match condition.
--------------------	---

Table 123: Create Term Fields for Data Center Switching Non-ELS (continued)

Field	Description
TCP Flags	<p>If you are not using the TCP initial flag for a match condition, select one of the TCP flags from the list—RST, ACK, SYN, Urgent, Push, Fin, None. These flags have the following meaning:</p> <ul style="list-style-type: none"> ● RST—Reset flag indicates that the TCP connection will be reset. ● ACK—Third step in TCP three-way handshake for connection. In response to a server's SYN-ACK, the client replies with an ACK. ● SYN—First step in TCP three-way handshake for connection. The active open is performed by the client sending a SYN to the server. ● Urgent—If the URG flag is set, then the 16-bit field is an offset from the sequence number indicating the last urgent data byte. ● Push—Push flags request that buffered data to the receiving application be sent now. ● FIN—The final flag indicates that no more data will be sent.

ICMP Settings

You can select the ICMP code value for the filter item's match condition. The Internet Control Message Protocol (ICMP) is one of the core IP protocols used by operating systems of networked computers to send error messages. ICMP can also be used to relay query messages.

ICMP Code	<p>To apply an ICMP code match condition to the named term:</p> <ol style="list-style-type: none"> a. Expand the ICMP Settings section. b. Click Add in the ICMP Codes section. The Select ICMP Codes list appears. c. Select one or more ICMP codes from the list. These codes vary, depending on the Filter Family you selected. d. Click OK. The ICMP code match condition is added to the named term, and the condition is listed in the ICMP Code list. <p>NOTE: ICMP code specifies more specific information than ICMP type. Because the value's meaning depends upon the associated ICMP type, you must specify ICMP type along with ICMP code. The keywords are grouped by the ICMP type with which they are associated.</p>
-----------	--

Table 123: Create Term Fields for Data Center Switching Non-ELS (continued)

Field	Description
ICMP Type	<p>NOTE: ICMP type specifies the ICMP packet type field. Typically, you specify this match condition in conjunction with the protocol match condition to determine which protocol is being used on the port.</p> <p>ICMP code specifies more specific information than ICMP type. Because the value's meaning depends upon the associated ICMP type, you must specify an ICMP type along with the ICMP code. The keywords are grouped by the ICMP type with which they are associated.</p> <p>To apply an ICMP type match condition to the named term:</p> <ol style="list-style-type: none"> Expand the ICMP Settings section. Click Add in the ICMP Type section. The Select ICMP Types list appears. Select one or more ICMP types from the list: echo-reply, echo-request, info-reply, info-request, mask-reply, mask-request, parameter-problem, redirect, router-advertisement, router-solicit, source-quench, time-exceeded, timestamp, timestamp-reply, or unreachable. Click OK. The ICMP type match condition is added to the named term, and the condition is listed in the ICMP Type list.

QFabric Settings

These settings apply only when the profile is applied to a QFabric device.

Enable-from-fabric	Select to create a match condition that matches packets coming from the fabric.
Enable-to-fabric	Select to create a match condition that matches packets going to the fabric.
except	Select to create a match condition that matches all packets that are not going to the fabric.

Action

Select the action that the system performs on an IP packet if all match conditions that you specified above are met. Possible actions are Discard and Accept. The default action is to discard packet that match the filter term's conditions.

Table 123: Create Term Fields for Data Center Switching Non-ELS (continued)

Field	Description
Action	Select either Discard or Accept to indicate what the filter term does with a packet when a match is made. NOTE: All other fields in this section are enabled only if you select Accept as the action.
Counter Name	Specify the maximum packet count for this filter, term, or policer.
Loss Priority	Specify the packet loss priority, low , high , or none . NOTE: Forwarding class and loss priority must be specified together for the same term.

Table 123: Create Term Fields for Data Center Switching Non-ELS (continued)

Field	Description
Policer	

Table 123: Create Term Fields for Data Center Switching Non-ELS (continued)

Field	Description
	<p>When you create a Filter profile, you can specify a policer action for any term or terms within the filter. Policing, or rate limiting, enables you to limit the amount of traffic that passes into or out of an interface. All traffic that matches a term that contains a policer action goes through the policer that the term references.</p> <p>You have two options with a policer. You can specify that an existing policer be used for the packet that matches the match condition. Or, you can create a new policer for the packet that matches the match condition.</p> <p>To select a policer from an existing list of policers, click Select. The Select Policer page appears. Select the policer that you want to use for the term and click OK. The system displays the selected policer in the Policer field in the Create Term page.</p> <p>To create a new policer:</p> <ol style="list-style-type: none"> a. Click Create. <ul style="list-style-type: none"> The Create Policer page appears. b. Type a name for the policer—you can use this policer again in the future. c. Select a policer type from the list, either a single-rate-two-color policer, a single-rate-three-color policer, or a three-color-policer. and then click OK. The type of policer that you select here affects the rest of the configurations available for the policer. <p>You can create the following three type of policers:</p> <ul style="list-style-type: none"> ● Single-rate two-color—A two-color policer (sometimes called simply <i>policer</i>) meters the traffic stream and classifies packets into two categories of packet loss priority (PLP) according to a configured bandwidth and burst-size limit. You can mark packets that exceed the bandwidth and burst-size limit or simply discard them. A two-color policer is most useful for metering traffic at the port (physical interface) level. ● Single-rate three-color—This type of policer is defined in RFC 2697, A Single Rate Three Color Marker, as part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on the configured committed information rate (CIR), committed burst size (CBS), and the excess burst size (EBS). Traffic is marked as belonging to one of three categories (green, yellow, or red) based on whether the packets are arriving at rates that are below the CBS (green), exceed the CBS but not the EBS (yellow), or exceed the EBS (red). A single-rate three-color policer is most useful when a service is structured according to packet size and not according to peak arrival rate. ● Two-rate three-color—This type of policer is defined in RFC 2698, A Two Rate Three Color Marker, as part of an assured forwarding (AF) per-hop-behavior (PHB)

Table 123: Create Term Fields for Data Center Switching Non-ELS (continued)

Field	Description
	<p>classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on the configured CIR and the peak information rate (PIR), along with their associated burst sizes; the CBS, and the peak burst size (PBS). Traffic is marked as belonging to one of three categories (green, yellow, or red) based on packets are arriving at rates that are below the CIR (green), exceed the CIR but not the PIR (yellow), or exceed the PIR (red). A two-rate three-color policer is most useful when a service is structured according to arrival rates and not to packet size.</p> <p>NOTE: The system displays and hides various fields in the Create Policer page depending on the type of policer that you want to create.</p>

Table 123: Create Term Fields for Data Center Switching Non-ELS (continued)

Field	Description
	<p>d. Configure the appropriate fields for the policer:</p> <ul style="list-style-type: none"> ● Bandwidth Limit—Specify the traffic rate in bits per second, 1000 through 102,300,000,000 (102.3g) bps. ● Burst Size Limit—Specify the maximum number of bytes allowed for incoming packets to burst above the peak information rate (PIR) and still be marked with medium-high packet loss priority (yellow). Packets that exceed the peak burst size (PBS) are marked with high packet loss priority (red). The range is 1500 through 100,000,000,000 bytes. ● Committed Information Rate—Specify the guaranteed bandwidth (in bits per second) under normal line conditions and the average rate up to which packets are marked with low packet loss priority (green). The range is 32,000 through 40,000,000,000 bps. ● Committed Burst Size—Specify the maximum number of bytes allowed for incoming packets to burst above the committed information rate (CIR) and still be marked with low packet loss priority (green). The range is 1500 through 100,000,000,000 bytes. ● Peak Burst Size—Specify the maximum number of bytes allowed for incoming packets to burst above the peak information rate (PIR) and still be marked with medium-high packet loss priority (yellow). Packets that exceed the peak burst size (PBS) are marked with high packet loss priority (red). The range is 1500 through 100,000,000,000 bytes. ● Excess Burst Size—Specify the maximum number of bytes allowed for incoming packets to burst above the committed information rate and still be marked with medium-high packet loss priority (yellow). Packets that exceed the excess burst size (EBS) are marked with high packet loss priority (red). The range is 1500 through 100,000,000,000 bytes. ● Peak Information Rate—Specify the maximum achievable rate in bits per second. Packets that exceed the peak information rate (PIR) are marked with high packet loss priority (red). You can configure a discard action for packets that exceed the PIR. The range is 32,000 through 40,000,000,000 bps ● Color Mode—Select the way the preclassified packets are to be metered: <ul style="list-style-type: none"> ● Color-aware—The local switch can assign a higher packet loss priority but cannot assign a lower packet loss priority. ● Color-blind—The local switch ignores the preclassification of packets and can assign a higher or lower packet loss priority. ● None—The preclassified packets are not metered. <p>e. Click OK.</p> <p>The policer is added to the list of applied policers and the list of available policers.</p>

Table 123: Create Term Fields for Data Center Switching Non-ELS (continued)

Field	Description
Forwarding Class	<p>Specify the forwarding class (or output queue) that is to be used for the packet that matches the match condition. You can create a new forwarding class or select from a list of available forwarding classes.</p> <p>To select a forwarding class from an existing list of classes, click Select. The Select Forwarding Class page appears. Select the forwarding class that you want to use for the packet and click OK. The system displays the selected forwarding class in the Forwarding Class field in the Create Term page.</p> <hr/> <p>To create a new forwarding class:</p> <ol style="list-style-type: none"> a. Click Create. The Create Forwarding Class page appears. b. Type a name for the forwarding class—you can use this forwarding class again in the future. c. Select a queue number from the list, and then click OK. The system creates a new forwarding class and displays it in the Forwarding Class field in the Create Term page.

Click **OK** to save the term and return to the Create Filter Profile page.

Specifying Settings for a Data Center Switching ELS Filter Profile

A Filter profile must have at least one term in it. Each term has one filtering function. For example, if a term is evaluating the source of packets, then that term cannot also evaluate the protocols used by the packets. Some switch models accommodate multiple terms in one filter. When you have more than one term in a filter, the ordering of the terms is important. The system evaluates multiple filter terms as follows:

- The packet is evaluated against the first term's conditions. If the packet matches all of the conditions in that term, the corresponding action for that condition is taken and evaluation ends. Subsequent terms in the filter are not evaluated.
- If the packet does not match all conditions in the first term, the packet is evaluated against the conditions in the second term. This process continues until either the packet matches all the conditions in one of the subsequent terms or there are no more terms in the filter. If a match is found, the action specified in the Action section of the matched term is taken and the evaluation ends. Subsequent terms in the filter are not evaluated.

- The term conditions for protocol, EtherType, DSCP, precedence, ICMP code and ICMP type must all be either match conditions or except conditions.
- If a packet passes through all the terms in the filter without a match, the packet is discarded.

To configure a Filter profile for Data Center switching ELS:

1. Specify a filter name and description for the Filter profile.
2. Select the switch filter family for which you want to create the profile:
 - If you want to create a Layer 2 based filter, select **Switching**.
 - If you want to create a Layer 3 based filter for IPv4, select **INET**.
 - If you want to create a Layer 3 based filter for IPv6, select **INET6**.
3. Under Terms, click **Add** to add one or more terms with match condition(s) to the named filter. You need at least one term for this filter.

The Create Term window opens.

NOTE: The order of the terms within a Filter profile configuration is important. Packets are tested against each term in the order in which the terms are listed.

4. Enter a name for the filter term.
5. Specify the match condition(s) for the filter term as described in [Table 124](#). Required settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 124: Create Term Fields for Data Center Switching ELS

Field	Description
-------	-------------

Source and Destination Parameters

You can specify match conditions based on the packets' origin (source) or the packets' destination, or both. You are indicating the location of the filtering here—either specifying that packets that originate at a specific place (source) will be filtered or packets destined for a specific location (destination) will be filtered. You can have multiple sources and destinations for one filter.

Table 124: Create Term Fields for Data Center Switching ELS (continued)

Field	Description
Source Parameters and Destination Parameters	<p>To add source and destination parameters to the named filter term:</p> <ol style="list-style-type: none"> a. Click Add to the right of the Destination Parameters lists. The Add Source/Destination Parameter window opens. b. Select either Source (default) or Destination from the Add Source/Destination Parameter window. c. Select one of following available Parameter Types from the Add Source/Destination Parameter page and provide the corresponding information: <ul style="list-style-type: none"> ● IP Address—Provide the IP address of the source or destination device. ● MAC Address—Provide a MAC address. ● Port—Provide the port type of the source or destination port. Select either AFS (Andrew File System), BGP (Border Gateway Protocol), BIFF (UNIX mail notification), Bootpc (bootstrap protocol client), Bootps, Cmd, CVS pserver, DHCP, Domain, EK login, EK shell, EXEC, Finger protocol, FTP, FTP data, HTTP, HTTPS, Ident protocol, IMAP (Internet Message Access protocol), Kerberos-sec (Kerberos security), Klogin forwarding, Kpasswd command, KRB-prop (Kerberos database propagation), Krbupdate (Kerberos database update), Kshell (Kerberos rsh), LDAP, Login (UNIX rlogin), Mobilip-agent (Mobile IP agent), Mobilip-mn (Mobile IP MN), MSDP (Multicast Source Discovery Protocol), NetBIOS dgm, NetBIOS-ns (NetBIOS name service), NetBIOS-ssn (NetBIOS session service), NFSD, NNTP (Network News Transport Protocol), Ntalk, NTP (Network Time Protocol), POP3 (Post Office Protocol3), PPTP, Printer, RADacct (RADIUS accounting), RADIUS, RIP, RKINIT (Kerberos remote kinit), SMTP, SNMP trap, SNPP, SUNRPC, Syslog, TACACS, TACACS-ds, Talk (UNIX Talk), Telnet, TFTP, Timed (UNIX time daemon), Who (UNIX rwho), XDMCP (X Display Manager Control Protocol), Zephyr-clt (Zephyr serv-hm connection), Zephyr-hm (Zephyr hostmanager), Zephyr-srv (Zephyr server), or Other. <p>NOTE: If you selected Port as the parameter and do not find the type of port that you want to add from the Port list, then select other and enter a port number.</p> d. To select any other source/destination than the one indicated, enable Except. TIP: You cannot indicate both matching and except for a parameter. e. Click OK The parameter term is added to the appropriate list, either Source Parameters or Destination Parameters.

Table 124: Create Term Fields for Data Center Switching ELS (continued)

Field	Description
<p>Protocols and EtherTypes</p>	
<p>Depending on the Filter Family you selected, you can sometimes apply a filter term based on either protocols being used by packets or on EtherTypes being used by packets. Recognized protocols are listed where applicable. Recognized EtherTypes, which indicate the protocol that is encapsulated in the payload of an Ethernet Frame, are also listed where applicable.</p>	
<p>Protocols (apply to Ethernet and INET filter families)</p>	<p>To add a protocol match condition to the named filter term:</p> <ol style="list-style-type: none"> Expand the Protocols and EtherTypes section. Click Add under Protocols. The Select Protocols window opens, displaying a list of protocols. From the list of protocols, select one or more. The options are AH, DSTOPTS, EGP, ESP, Fragment, GRE, Hop-by-hop, ICMP, IPIP, IPv6, No-text-header, OSPF, PIM, Routing, RSVP, SCTP, TCP, UDP, and VRRP. To make the filter exclude the specified protocol, select Except. NOTE: The term conditions for protocol, EtherType, DSCP, precedence, ICMP code and ICMP type must all be either match conditions or except conditions. Click OK. The protocols are added to the Protocols list.

Table 124: Create Term Fields for Data Center Switching ELS (continued)

Field	Description
EtherTypes (apply to Ethernet filter family)	<p>To add an EtherTypes match condition to the named filter term:</p> <ol style="list-style-type: none"> a. Expand the Protocols and EtherTypes section. b. Click Add under EtherTypes. The Select EtherTypes window opens, displaying a list of protocols. c. From the list of EtherTypes, select one or more. The options are AARP, AppleTalk, ARP, FCoE, FIP, IPV4, MPLS multicast, MPLS unicast, OAM, PPP, PPPOE discovery, PPPOE session, and SNA. d. To make the filter exclude the specified EtherType, select Except. NOTE: Term values must all be either match conditions or all except conditions. e. Click OK. The EtherTypes are added to the EtherTypes list.

DSCP Settings

Expand the DSCP section to see the DSCP match settings. DiffServ is a simple mechanism for classifying and managing network traffic and providing quality-of-service (QoS) on IP networks. DiffServ can, for example, be used to apply low-latency to critical network traffic such as voice or streaming media while providing simple best-effort service to non-critical services such as Web traffic. Here, you can apply a filter term based on the Differentiated Services code point (DSCP) which is a field in IPv4 and IPv6 headers.

NOTE: With IPv6 packets, the DS field and ECN field replace the IPv4 TOS field.

Table 124: Create Term Fields for Data Center Switching ELS (continued)

Field	Description
DSCP (Ethernet and INET filter families)	<p>To add a DSCP match condition to the named filter term:</p> <p>NOTE: A DSCP IP match condition and a precedence match condition cannot be both specified for the same term.</p> <p>a. Click Add in the DSCP section.</p> <p>The Select DSCP Match Condition list appears.</p> <p>b. Select one of the following DSCP types from the list:</p> <ul style="list-style-type: none"> ● AF11—Assured forwarding class 1, low drop precedence ● AF12—Assured forwarding class 1, medium drop precedence ● AF21—Assured forwarding class 2, low drop precedence ● AF22—Assured forwarding class 2, medium drop precedence ● AF23—Assured forwarding class 2, high drop precedence ● AF31—Assured forwarding class 3, low drop precedence ● AF32—Assured forwarding class 3, medium drop precedence ● AF33—Assured forwarding class 3, high drop precedence ● AF41—Assured forwarding class 4, low drop precedence ● AF42—Assured forwarding class 4, medium drop precedence ● AF43—Assured forwarding class 4, high drop precedence ● BE—Best effort (default) ● EF—Expedited forwarding ● CS0—Class selector 0 ● CS1—Class selector 1 ● CS2—Class selector 2 ● CS3—Class selector 3 ● CS4—Class selector 4 ● CS5—Class selector 5 ● CS6—Class selector 6 ● CS7—Class selector 7 <p>c. To make the filter exclude a specified DSCP type, select Except.</p> <p>NOTE: Term values must all be either match conditions or all of them need to be except conditions.</p> <p>d. Click OK.</p> <p>The DSCP code term for the named filter is added to the DSCP list.</p>

Table 124: Create Term Fields for Data Center Switching ELS (continued)

Field	Description
Precedence for DSCP (Ethernet and INET filter families)	<p>You can apply an IP precedence match condition to the named term. With IP precedence, a device prioritizes traffic by class first. Then it differentiates and prioritizes same-class traffic.</p> <p>NOTE: The match conditions IP Precedence and DSCP cannot be simultaneously applied to a term.</p> <p>To apply an IP precedence value match condition to the named term:</p> <ol style="list-style-type: none"> Click Add in the Precedence section. The Select Precedence list appears. Select one of the following precedence settings from the list: Routine (0 or lowest, also called Best Effort), Priority (1), Immediate (2), Flash (3, mainly used for voice signaling or for video), Flash-override (4), Critical-ECP (5, mainly used for voice RTP), Internet-control (6, used for IP routing protocols), or Net-control (7 or highest, used for link layer and routing protocol keep alive). To make the filter exclude the specified IP precedence value, select Except. NOTE: Term values must all be either match conditions or all of them need to be except conditions. Click OK. The precedence match condition is added to the named term, and the condition is listed in the Precedence list.

TCP Settings

Expand this section to access the TCP settings. The Transmission Control Protocol (TCP) is the most common core protocol of the Internet protocol suite (IP). TCP provides reliable, ordered, error-checked delivery of a stream of octets between programs running on computers connected to the Internet or an intranet. You can use the TCP initial flag for a match condition.

Enable TCP Initial (all families)	<p>Select to use the TCP initial flag for an Ethernet, INET, or INET6 match condition.</p> <p>TIP: If you use the TCP initial flag for filtering, you cannot use any other TCP flag.</p>
-----------------------------------	---

Table 124: Create Term Fields for Data Center Switching ELS (continued)

Field	Description
TCP Flags	<p>If you are not using the TCP initial flag for a match condition, you can select one of the TCP flags from the list for a match condition—RST, ACK, SYN, Urgent, Push, FIN, or None. These flags have the following meaning:</p> <ul style="list-style-type: none"> • RST—Reset flag indicates that the TCP connection will be reset. • ACK—Third step in TCP three-way handshake for connection. In response to a server's SYN-ACK, the client replies with an ACK. • SYN—First step in TCP three-way handshake for connection. The active open is performed by the client sending a SYN to the server. • Urgent—If the URG flag is set, then the 16-bit field is an offset from the sequence number indicating the last urgent data byte. • Push—Push flags request that buffered data to the receiving application be sent now. • FIN—The final flag indicates that no more data will be sent.

ICMP Settings

You can select the ICMP code value for the filter item's match condition—expand this section to access the ICMP settings. The Internet Control Message Protocol (ICMP) is one of the core IP protocols used by operating systems of networked computers to send error messages. ICMP can also be used to relay query messages.

ICMP Code	<p>To apply an ICMP code match condition to the named term:</p> <ol style="list-style-type: none"> a. Click Add in the ICMP Code section. The Select ICMP Code window appears. b. Select one or more ICMP codes from the list. These codes vary, depending on the Filter Family you selected. c. To make the filter exclude the specified ICMP code, select Except. NOTE: Term values must all be either match conditions or all of them need to be except conditions. d. Click OK. The ICMP code match condition is added to the named term, and the condition is listed in the ICMP Code list. <p>NOTE: An ICMP code specifies more specific information than an ICMP type. Because the value's meaning depends upon the associated ICMP type, you must specify an ICMP type along with ICMP code. The keywords are grouped by the ICMP type with which they are associated.</p>
-----------	---

Table 124: Create Term Fields for Data Center Switching ELS (continued)

Field	Description
ICMP Type	<p>NOTE: ICMP type specifies the ICMP packet type field. Typically, you specify this match condition in conjunction with the protocol match condition to determine which protocol is being used on the port.</p> <p>ICMP code specifies more specific information than ICMP type. Because the value's meaning depends upon the associated ICMP type, you must specify ICMP type along with ICMP code. The keywords are grouped by the ICMP type with which they are associated.</p> <p>To apply an ICMP type match condition to the named term:</p> <ol style="list-style-type: none"> Click Add in the ICMP Type section. The Select ICMP Types window appears. Select one or more ICMP types from the list. These types vary, depending on the Filter Family selected. To make the filter exclude the specified ICMP type, select Except. NOTE: Term values must all be either match conditions or all of them need to be except conditions. Click OK. The ICMP type match condition is added to the named term, and the condition is listed in the ICMP Type list.
<p>Action</p>	
<p>Select the action that the system performs on an IP packet if all match conditions that you specified above are met. Possible actions are Discard and Accept. The default action is to discard packet that matches the filter term conditions.</p>	
Action	<p>Select either Discard or Accept to indicate what the filter term does with a packet when a match is made.</p> <p>NOTE: All other fields in this section are enabled only if you select Accept as the action.</p>
Counter Name	<p>When the action selected is accept, specify the maximum packet count for this filter, term, or policer.</p>

Table 124: Create Term Fields for Data Center Switching ELS (continued)

Field	Description
Loss Priority	<p>When the action selected is accept, specify the packet loss priority, Low, High, Medium-low, Medium-high, or None.</p> <p>NOTE: Forwarding class and loss priority must be specified together for the same term.</p>

Table 124: Create Term Fields for Data Center Switching ELS (continued)

Field	Description
Policer	<p>When you create a Filter profile with the action accept, you can specify a policer action for any term or terms within the filter. Policing, or rate limiting, enables you to limit the amount of traffic that passes into or out of an interface. All traffic that matches a term that contains a policer action goes through the policer that the term references.</p> <p>You have two options with a policer. You can specify that an existing policer be used for the packet that matches the match condition. Or, you can create a new policer for the packet that matches the match condition.</p> <hr/> <p>To select an existing policer:</p> <ol style="list-style-type: none"> a. Click Select. The Select Policer page appears. b. Click OK. The policer is added to the list of applied policers. <hr/>

Table 124: Create Term Fields for Data Center Switching ELS (continued)

Field	Description
	<p>To create a new policer:</p> <ol style="list-style-type: none"> a. Click Create. <ul style="list-style-type: none"> The Create Policer page appears. b. Type a name for the policer—you can use this policer again in the future. c. Select a policer type from the list, either a single-rate-two-color policer, or a three-color-policer. The type of policer that you select here affects the rest of the configurations available for the policer. <p>If you selected a three-color-policer, then also select a rate for it, either single-rate or two-rate.</p> <ul style="list-style-type: none"> • Single-rate two-color—A two-color policer (sometimes called simply <i>policer</i>) meters the traffic stream and classifies packets into two categories of packet loss priority (PLP) according to a configured bandwidth and burst-size limit. You can mark packets that exceed the bandwidth and burst-size limit or simply discard them. A two-color policer is most useful for metering traffic at the port (physical interface) level. • Single-Rate Three-color—This type of policer is defined in RFC 2697, A Single Rate Three Color Marker, as part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on the configured committed information rate (CIR), committed burst size (CBS), and the excess burst size (EBS). Traffic is marked as belonging to one of three categories (green, yellow, or red) based on whether the packets are arriving at rates that are below the CBS (green), exceed the CBS but not the EBS (yellow), or exceed the EBS (red). A single-rate three-color policer is most useful when a service is structured according to packet size and not according to peak arrival rate. • Two-rate three-color—This type of policer is defined in RFC 2698, A Two Rate Three Color Marker, as part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on the configured CIR and the peak information rate (PIR), along with their associated burst sizes; the CBS, and the peak burst size (PBS). Traffic is marked as belonging to one of three categories (green, yellow, or red) based on packets are arriving at rates that are below the CIR (green), exceed the CIR but not the PIR (yellow), or exceed the PIR (red). A two-rate three-color policer is most useful when a service is structured according to arrival rates and not to packet size. <p>NOTE: The system displays and hides various fields in the Create Policer page depending on the type of policer that you want to create.</p>

Table 124: Create Term Fields for Data Center Switching ELS (continued)

Field	Description
	<p>d. Configure these fields for a single-rate-two-color policer:</p> <ul style="list-style-type: none"> ● Bandwidth Limit—Specify the traffic rate in bits per second, 8000 through 50,000,000,000 bps. ● Burst Size Limit—Specify the maximum number of bytes allowed for incoming packets to burst above the peak information rate (PIR) and still be marked with medium-high packet loss priority (yellow). Packets that exceed the peak burst size (PBS) are marked with high packet loss priority (red). The range is 1 through 2,147,450,880 bytes. ● Action—The default action is Discard. ● Loss Priority—Not available. <p>e. Configure these fields for a single-rate-three-color policer:</p> <ul style="list-style-type: none"> ● Committed Information Rate—Specify the guaranteed bandwidth (in bits per second) under normal line conditions and the average rate up to which packets are marked with low packet loss priority (green). The range is 1500 through 100,000,000,000 bps. ● Committed Burst Size—Specify the maximum number of bytes allowed for incoming packets to burst above the committed information rate (CIR) and still be marked with low packet loss priority (green). The range is 1500 through 100,000,000,000 bytes. ● Excess Burst Size—Specify the maximum number of bytes allowed for incoming packets to burst above the committed information rate and still be marked with medium-high packet loss priority (yellow). Packets that exceed the excess burst size (EBS) are marked with high packet loss priority (red). The range is 1500 through 100,000,000,000 bytes. ● Color Mode—Select the way the preclassified packets are to be metered: <ul style="list-style-type: none"> ● Color-aware—The local switch can assign a higher packet loss priority but cannot assign a lower packet loss priority. ● Color-blind—The local switch ignores the preclassification of packets and can assign a higher or lower packet loss priority. ● None—The preclassified packets are not metered. ● Action—Options are Discard and None. ● Loss Priority—Options are High and None.

Table 124: Create Term Fields for Data Center Switching ELS (continued)

Field	Description
	<p>f. Configure these fields for a three-color two-rate policer:</p> <ul style="list-style-type: none"> ● Committed Information Rate—Specify the guaranteed bandwidth (in bits per second) under normal line conditions and the average rate up to which packets are marked with low packet loss priority (green). The range is 1500 through 100,000,000,000 bps. ● Committed Burst Size—Specify the maximum number of bytes allowed for incoming packets to burst above the committed information rate (CIR) and still be marked with low packet loss priority (green). The range is 1500 through 100,000,000,000 bytes. ● Peak Burst Size—Specify the maximum number of bytes allowed for incoming packets to burst above the peak information rate (PIR) and still be marked with medium-high packet loss priority (yellow). Packets that exceed the peak burst size (PBS) are marked with high packet loss priority (red). The range is 1500 through 100,000,000,000 bytes. ● Peak Information Rate—Specify the maximum achievable rate in bits per second. Packets that exceed the peak information rate (PIR) are marked with high packet loss priority (red). You can configure a discard action for packets that exceed the PIR. The range is 1500 through 100,000,000,000 bps. ● Color Mode—Select the way the preclassified packets are to be metered: <ul style="list-style-type: none"> ● Color-aware—The local switch can assign a higher packet loss priority but cannot assign a lower packet loss priority. ● Color-blind—The local switch ignores the preclassification of packets and can assign a higher or lower packet loss priority. ● None—The preclassified packets are not metered. ● Action—Options are Discard and None. ● Loss Priority—Options are High and None. <p>g. Click OK.</p> <p>The policer is added to the list of applied policers and the list of available policers.</p>

Table 124: Create Term Fields for Data Center Switching ELS (continued)

Field	Description
Forwarding Class	<p>Specify the forwarding class (or output queue) that is to be used for the packet that matches the match condition. You can either select from a list of available forwarding classes or create a new forwarding class.</p> <p>To select a forwarding class from an existing list of classes, click Select. The Select Forwarding Class page appears. Select the forwarding class that you want to use for the packet and click OK. The system displays the selected forwarding class in the Forwarding Class field in the Create Term page.</p> <hr/> <p>To create a new forwarding class:</p> <ol style="list-style-type: none"> a. Click Create. The Create Forwarding Class page appears. b. Type a name for the forwarding class—you can use this forwarding class again in the future. c. Select a queue number from the list, and then click OK. The system creates a new forwarding class and displays it in the Forwarding Class field in the Create Term page.

6. Click **OK** to save the term and return to the Create Filter Profile page.

7. Click **Done**.

The new filter is added to the Manage Filter Profile list.

What to Do Next

After you create a Filter profile, you can do one of the following:

- Link the Filter profile as ingress and egress filters to a Port profile. For more information, see [“Creating and Managing Port Profiles” on page 414](#).
- Link the Filter profile as ingress and egress filters to a VLAN profile. For more information, see [“Creating and Managing VLAN Profiles” on page 502](#). You can then assign the VLAN profile to a device or port in case of switching devices and to a device, port, or access point in case of a wireless network.

RELATED DOCUMENTATION

[Understanding Filter Profiles | 540](#)

[Creating and Managing Wireless Filter Profiles | 598](#)

[Network Director Documentation home page](#)

Creating and Managing Wireless Filter Profiles

IN THIS SECTION

- [Managing Wireless Filter Profiles | 598](#)
- [Creating a Wireless Filter Profile | 600](#)
- [Specifying Settings for a Wireless Filter Profile \(WLC\) | 600](#)
- [What To Do Next | 606](#)

Filter profiles are a set of rules that determine whether to accept or discard packets transiting on either a switch or wireless radio interface.

Use the Manage Filter Profiles page to create new wireless Filter Profiles and manage existing wireless Filter Profiles.

This topic describes:

Managing Wireless Filter Profiles

From the Manage Filter Profiles page, you can:

- Create a new wireless Filter profile by clicking **Add**. For directions, see [“Creating a Wireless Filter Profile” on page 600](#).
- Modify an existing wireless Filter Profile by selecting it and clicking **Edit**.
- Assign a wireless Filter Profile to controllers by selecting it and clicking **Assign**. For directions, see [“Assigning a Wireless Filter Profile to Controllers” on page 607](#).
- Reassign a wireless Filter Profile by selecting it and clicking **Edit Assignment**.
- View information about a wireless Filter profile, including the associated interfaces, by either clicking the profile name or by selecting the profile and clicking **Details**.
- Delete a wireless Filter profile by selecting a profile and clicking **Delete**.

TIP: You cannot delete profiles that are in use—that is, profiles assigned to objects or used by other profiles. To see the current assignments for a profile, select the profile and click **Details**.

- Clone a wireless Filter profile by selecting a profile and clicking **Clone**.

Table 125 describes the information provided about wireless Filter profiles on the Manage Filter Profiles page. This page lists all Filter profiles defined for your network, regardless of the scope you selected in the network view.

Table 125: Manage Filter Profile Fields

Field	Description
Profile Name	Name given to the profile when the profile was created.
Family Type	Wireless (WLC)
Description	Description of the profile entered when the profile was created. TIP: To display the entire description, you might need to resize the Description column by clicking the column border in the heading and dragging it.
Filter Family	NA (for wireless).
Assignment State	The assignment state can be: <ul style="list-style-type: none"> • Unassigned—When the profile is not assigned to any object. • Deployed—When the profile is assigned and is deployed from Deploy mode. • Pending Deployment—When the profile is assigned, but not yet deployed in the network.
Creation Time	Date and time when the profile was created.
Last Updated Time	Date and time when the profile was last modified.
User Name	The username of the user who created or modified the profile.

TIP: All columns might not be displayed. To show or hide fields in the table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.


Creating a Wireless Filter Profile

To create a Filter profile, you must provide a filter name and configure at least one term. A term is a collection of one or more match conditions, and actions that the system takes when match conditions are met. A term must have at least one match condition.

To create a wireless Filter Profile:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.
3. From the Tasks pane, expand **Wireless**, expand **Profiles**, and then select **Filter**.
The Manage Filter Profile window opens, displaying a list of currently configured wireless filters under two tabs, **All Profiles** and **Assigned Profiles**.
4. From the Manage Filter Profile window, click **Add** to add a new profile.
The Create Filter Profile for Wireless window opens.
5. Complete the settings described in both the online help and in [“Specifying Settings for a Wireless Filter Profile \(WLC\)” on page 600](#).
6. Click **Done**.

Specifying Settings for a Wireless Filter Profile (WLC)

A Filter profile must have at least one term in it. Each term has only one filtering function. For example, if a term is evaluating the source of packets, then that term cannot also evaluate the protocols used by the packets. Some switch models accommodate multiple terms in one filter. When you have more than one term in a filter, the ordering of the terms is important. The system evaluates multiple filter terms as follows:

- A packet is evaluated against the conditions in the first term of the filter. If the packet matches the conditions in the term, the corresponding action is executed and the evaluation ends. Subsequent terms in the filter are not evaluated.
- If a packet does not match the conditions in the first term, the packet is evaluated against the conditions in the second term. This process continues until either the packet matches conditions in one of the

subsequent terms or there are no more terms in the filter. If a match is found, the corresponding action is executed and the evaluation ends. Subsequent terms in the filter are not evaluated.

- If a packet passes through all the terms in the filter without a match, the packet is discarded.

To configure a Filter profile on a controller:

1. Specify a **Filter Name** and **Description** for the wireless Filter profile.
2. Click **Add** under Terms to add at least one term to the wireless filter. The Create Term window opens.

TIP: The order of the terms within a Filter profile configuration is important. Packets are tested against each term in the order in which the terms are listed in the Filter profile.

3. Enter the settings described in [Table 126](#) to create a wireless term. Required settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 126: Term Fields for Wireless

Field	Description
Term name (all)	Provide a name for this term
Rule type (all)	Select the type of rule (or term) that you want to create. You can create an IP-based rule or a MAC-based rule. The rule type you select affects the rest of the configuration.
IP Type (IP-based rule)	If you are creating an IP-based rule, indicate the IP type, either IPv4 or IPv6 .
Source IP Address (IP-based rule)	Type the source IP address of the term. This parameter specifies the match conditions for packets that originate from the given IP address.
Destination IP Address (IP-based rule)	Type the destination IP address of the term. This parameter specifies the match conditions for packets that terminate at the given IP address.

Table 126: Term Fields for Wireless (continued)

Field	Description
Protocol (IP-based rule)	Select a protocol for the filter term. Select Any to include packets that use any supported protocols to be part of the rule or select None to discard protocol based filtering. Other options are ICMP , TCP , UDP , and Other . The protocol that you select here affects the rest of the settings in this window.

Table 126: Term Fields for Wireless (continued)

Field	Description
	<p>If you selected ICMP as the protocol:</p> <ol style="list-style-type: none"> a. Optionally, change the default protocol number for ICMP, which is 1. b. Indicate an ICMP Code number. ICMP code specifies more specific information than ICMP type. Because the value's meaning depends upon the associated ICMP type, you must specify an ICMP type along with an ICMP code. The keywords are grouped by the ICMP type with which they are associated. c. Indicate an ICMP Type number. ICMP Type specifies the ICMP packet type field. Typically, you specify this match condition in conjunction with the protocol match condition to determine which protocol is being used on the port. d. Indicate a DSCP number. DSCP filters packets by Differentiated Services Code Point (DSCP) value. You can specify a number from 0 to 63, in decimal or binary format. <p>NOTE: You cannot use the DSCP option along with the Precedence and ToS options in the same term.</p> e. Indicate a precedence for the term, either Routine (0), Priority (1), Immediate (2), Flash (3), Flash Override (4), Critical ECP (5), Internet Control (6), Net Control (7) or None (default). f. Indicate a ToS number. The ToS number specifies the type of service (ToS) level to filter packets. Specify one of the following values, or any sum of these values up to 15: <ul style="list-style-type: none"> ● 8—minimum delay ● 4—maximum throughput ● 2—maximum reliability ● 1—minimum monetary cost ● 0—normal <p>For example, a ToS value of 9 filters packets with the ToS levels minimum delay (8) and minimum monetary cost (1).</p> <p>NOTE: You cannot use the DSCP option along with the Precedence and ToS options in the same term.</p> g. Click OK. <p>The term is added to the filter.</p>

Table 126: Term Fields for Wireless (continued)

Field	Description
	<p>If you selected TCP or UDP as the protocol:</p> <ol style="list-style-type: none"> a. Optionally, change the default protocol number for TCP (6) or UDP (17). b. Indicate a source port operator. Options are None (default), Less than, Greater than, Range, Equal, or Not equal. For any option other than None, also indicate a source port name and a source port number. c. Indicate a destination port operator. Options are None (default), Less than, Greater than, Range, Equal, or Not equal. For any option other than None, also indicate a source port name and a source port number. d. Indicate a DSCP number. DSCP filters packets by Differentiated Services Code Point (DSCP) value. You can specify a number from 0 to 63, in decimal or binary format. NOTE: You cannot use the DSCP option with the Precedence and ToS options in the same term. e. Indicate a precedence for the term, either Routine (0), Priority (1), Immediate (2), Flash (3), Flash Override (4), Critical ECP (5), Internet Control (6), Net Control (7) or None (default). f. Indicate a ToS number. The ToS number specifies the type of service (ToS) level to filter packets. Specify one of the following values, or any sum of these values up to 15: <ul style="list-style-type: none"> ● 8—minimum delay ● 4—maximum throughput ● 2—maximum reliability ● 1—minimum monetary cost ● 0—normal <p>For example, a ToS value of 9 filters packets with the ToS levels minimum delay (8) and minimum monetary cost (1)</p> NOTE: You cannot use the DSCP option along with the Precedence and ToS options in the same term. g. Click OK. The term is added to the filter.

Table 126: Term Fields for Wireless (continued)

Field	Description
	<p>If you selected Any as the protocol:</p> <ol style="list-style-type: none"> a. Indicate a DSCP number. DSCP filters packets by Differentiated Services Code Point (DSCP) value. You can specify a number from 0 to 63, in decimal or binary format. <p>NOTE: You cannot use the DSCP option along with the Precedence and ToS options in the same term.</p> b. Indicate a precedence for the term, either Routine (0), Priority (1), Immediate (2), Flash (3), Flash Override (4), Critical ECP (5), Internet Control (6), Net Control (7) or None. c. Indicate a ToS number. The ToS number specifies the type of service (ToS) level to filter packets. Specify one of the following values, or any sum of these values up to 15: <ul style="list-style-type: none"> ● 8—minimum delay ● 4—maximum throughput ● 2—maximum reliability ● 1—minimum monetary cost ● 0—normal <p>For example, a ToS value of 9 filters packets with the ToS levels minimum delay (8) and minimum monetary cost (1)</p> <p>NOTE: You cannot use the DSCP option along with the Precedence and ToS options in the same term.</p> d. Click OK. <p>The term is added to the filter.</p>
Source MAC Address (MAC-based rule)	Type the source MAC address of the term. This parameter specifies the match conditions for packets that originate from the given MAC address.
Destination MAC Address (MAC-based rule)	Type the destination MAC address of the term. This parameter specifies the match conditions for packets that terminate at the given MAC address.
EtherType (MAC-based rule)	Specify EtherType filtering for the term. EtherType indicates the protocol that is encapsulated in the payload of an Ethernet Frame. Select Any to include packets that use any EtherType to be part of the rule or None to discard EtherType-based filtering.

Table 126: Term Fields for Wireless (continued)

Field	Description
Action	<p>Select the action that the system performs on an IP packet if the match conditions that you specified above are met. Possible actions are Discard and Accept. The default action is to discard the packet.</p> <p>NOTE: Forwarding Class is enabled only if you select Accept as the action.</p>
Forwarding Class	<p>Specifies the forwarding class (or output queue) that is to be used for the packet that matches the condition. You can create a new forwarding class or select from a list of available forwarding classes only if you specified the action as Accept.</p> <p>To create a new forwarding class, click Create. The Create Forwarding Class page appears. Specify a name for the forwarding class and the corresponding output queue number and click OK. The system creates a new forwarding class and displays it in the Forwarding Class field in the Create Term page.</p> <p>To select a forwarding class from an existing list of classes, click Select. The Select Forwarding Class page appears. Select the forwarding class that you want to use for the packet and click OK. The system displays the selected forwarding class in the Forwarding Class field in the Create Term page.</p>

Click **OK** to save the wireless filter term and return to the Create Filter Profile page.

What To Do Next

After you create a Filter profile, you can do one of the following:

- Assign filter profile to a device or cluster. For more information, see [“Assigning a Wireless Filter Profile to Controllers” on page 607](#).
- Link the Filter profile as ingress and egress filters to an Authorization profile. For more information, see [“Creating and Managing Wireless Authorization Profiles” on page 395](#).
- Link the Filter profile as ingress and egress filters to a VLAN profile. For more information, see [“Creating and Managing VLAN Profiles” on page 502](#). You can then assign the VLAN profile to a device, port, or access point in case of a wireless network.

RELATED DOCUMENTATION

[Understanding Filter Profiles | 540](#)


[Assigning a Wireless Filter Profile to Controllers | 607](#)

Assigning a Wireless Filter Profile to Controllers

To assign a Filter Profile to a controller:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.
3. In the View pane, select one or more controllers. If you select **Wireless Network**, you select all controllers.
4. From the Tasks pane, expand **Wireless**, expand **Profiles**, and then select **Filter**.
The Manage Filter Profiles page appears, displaying the list of currently configured wireless Filter Profiles.
5. Select a wireless Filter Profile from the list by adding a check mark next to the name, and then click **Assign**.
The Assign Filter Profile wizard opens. The wizard consists of three sections, **Device Selection**, **Profile Assignment**, and **Review**.
6. From **Device Selection**, select one or more devices for Filter Profile assignment—all devices under the selection are also selected.
7. Click **Profile Assignment** to open the next page of the wizard.
The Assign Filter Profile page opens, displaying a list of selected devices.
8. From **Profile Assignment**, select one or more devices from the list by placing a check mark next to the devices. Click either **Assign to Device** or **Assign to Cluster**.

9. Click **Review** and then click **Edit** to make any needed changes.

10. Click **Finish**.

The new assignment appears in the list.

RELATED DOCUMENTATION

[Creating and Managing Wireless Filter Profiles | 598](#)

[Network Director Documentation home page](#)

Configuring Class of Service (CoS)

IN THIS CHAPTER

- Understanding Class of Service (CoS) Profiles | 609
- Creating and Managing Wired CoS Profiles | 613
- Creating and Managing Wireless CoS Profiles | 630
- Assigning a Wireless CoS Profile to Controllers | 635

Understanding Class of Service (CoS) Profiles

IN THIS SECTION

- How Would I Use CoS (also known as QoS)? | 610
- What Wireless Network Traffic Aspects Can I Control Using CoS? | 611
- What CoS Parameters Can I Control? | 611
- What Are the Default CoS Traffic Types? | 612
- Data Center Switching CoS Configuration | 612
- How Do I Implement Class of Service? | 613
- Editing Discovered CoS Profiles | 613

When a network experiences congestion and delay, some packets must be prioritized to avoid random loss of data. Class of service (CoS) (also known as QoS) accomplishes this prioritization by dividing similar types of traffic, such as e-mail, streaming video, voice, large document file transfer, into classes. You then apply different levels of priority, such as those for throughput and packet loss, to each group, and thereby control traffic behavior. For example, when packets must be dropped, you can ensure that packet loss takes place according to your configured rules. CoS also enables you to rewrite the Differentiated Services code point (DSCP), IP precedence, or 802.1p CoS bits of packets exiting a specific interface, thus enabling you to tailor outgoing packets to meet the network requirements of remote peers.

On Data Center Switching devices, CoS can be used to configure Ethernet interfaces to support Fibre Channel over Ethernet (FCoE) traffic.

How Would I Use CoS (also known as QoS)?

On an Ethernet trunk, you can mark frames with a class-of-service (CoS) value. CoS is used to define trunk connections as full-duplex, incoming only, or outgoing only.

Network devices such as routers and switches can be configured to use existing CoS values on incoming packets from other devices (trust mode), or can rewrite the CoS values to something completely different. Layer 2 markings also can extend to the WAN; for example, with a frame relay network. CoS is usually limited to use within an organization's intranet.

With legacy telephone systems, CoS can be used to define the permissions an extension will have on a private branch exchange (PBX) or Centrex. Some users might need extended voicemail message retention or the ability to forward calls to a cell phone, while others have no need to make calls outside the office. Permissions for a group of extensions can be changed by modifying a CoS variable applied to the entire group.

NOTE: CoS configurations can be complicated, so unless it is required, we recommend that you do not alter the default class names or queue number associations.

How Do I Create CoS Groups?

Use 802.1Q tagged VLANs to group users and enable CoS to set priorities supported by downstream devices.

How Is CoS Different From QoS?

CoS operates only on 802.1Q VLAN Ethernet at the data link layer (layer 2), while quality-of-service (QoS) mechanisms operate at the IP network layer (layer 3). 802.1p Layer 2 tagging can be used by QoS to differentiate and shape network traffic.

What Wireless Network Traffic Aspects Can I Control Using CoS?

In addition to separating traffic into classes, you can also optionally configure these settings with CoS:

- Apply a bandwidth limit to the data sessions and to aggregated categories such as trunk interfaces, Layer 3 interfaces, access interfaces, and routed VLAN interfaces.
- Assign the same CoS level to all traffic on the Service profile SSID. This is called static CoS and overrides settings indicated on the 802.1p, overrides DSCP markings in the packets themselves, and disregards any filters that mark CoS. You indicate the value assigned to all user traffic.
- Allow the controller to use the client DSCP for radio ingress traffic and ignore Wi-Fi Multimedia (WMM).
- Specify a traffic class for voice traffic and optionally apply a bandwidth limit to the voice sessions and to aggregated categories. You can also enable static CoS for voice traffic, which overrides settings indicated on the 802.1p, overrides DSCP markings in the packets themselves, and disregards any filters that mark CoS. You indicate the value assigned to all voice traffic.
- Specify which of 11 forwarding queues are used. You can modify the action corresponding to each forwarding queue to suit your requirements. This is referred to as access categories.

What CoS Parameters Can I Control?

You can use CoS profiles to group a set of class of service (CoS) parameters and apply it to one or more interfaces. You can configure the following parameters within a CoS profile:

- Classifiers—Packet classification refers to the examination of an incoming packet. This function associates the packet with a particular CoS servicing level.
- Scheduler maps—Schedulers define the properties of output queues. These properties include the amount of interface bandwidth assigned to the queue, the size of the memory buffer allocated for storing packets, the priority of the queue, and the drop profiles associated with the queue. You associate the schedulers with forwarding classes by means of scheduler maps. You can then associate each scheduler map with an interface, thereby configuring the queues, packet schedulers, and tail drop processes that operate according to this mapping.
- Rewrite values—A rewrite rule modifies the appropriate CoS bits in an outgoing packet. Modification of CoS bits enables the next downstream device to classify the packet into the appropriate service group. Rewriting or marking outbound packets is useful when the device is at the border of a network and must alter the CoS values to meet the policies of the targeted peer.
- Traffic-control profile—Traffic-control profiles enable traffic limitation of a certain class to a specified bandwidth and burst size. Packets exceeding the limits can be discarded, or can be assigned to a different forwarding class, a different loss priority, or both.

What Are the Default CoS Traffic Types?

On EX Series switches, the system provides you with these four predefined traffic types—Data, Voice, Video, and Network Control—with these default traffic configuration and shaping details:

- Data—Forwarding queue 0 (nd_best-effort), Buffer size 50%, Bandwidth reserved 30%
- Voice—Forwarding queue 5 (nd_expedited-forwarding), Buffer size 20%, Bandwidth reserved 0%
- Video—Forwarding queue 4 (nd_video-forwarding), Buffer size 20%, Bandwidth reserved 70%
- Network Control—Forwarding queue 7 (nd_network-control), Buffer size 10%, Bandwidth reserved 0%

For Campus Switching ELS, the system provides you with these four predefined traffic types—Data, Voice, Video, and Network Control—with these default traffic configuration and shaping details:

- Data—Forwarding queue 0 (nd_best-effort), Buffer size 50%, Bandwidth reserved 30%
- Voice—Forwarding queue 1 (nd_expedited-forwarding), Buffer size 20%, Bandwidth reserved 0%
- Video—Forwarding queue 2 (nd_video-forwarding), Buffer size 20%, Bandwidth reserved 70%
- Network Control—Forwarding queue 3 (nd_network-control), Buffer size 10%, Bandwidth reserved 0%

For Campus Switching ELS with *Hierarchical Port Scheduling* (Juniper Networks EX4600 Ethernet switches), Network Director provides you with predefined forwarding classes—nd_cs_best-effort, nd_cs_video-forwarding, nd_cs_expedited-forwarding, and nd_cs_network-control. These forwarding classes are grouped under two priority groups—data_video_pg and voice_control_pg.

On data center switches, the system provides you with forwarding classes—nd_dc_best-effort, nd_dc_network-control, nd_dc_fcoe, nd_dc_no-loss, and nd_dc_mcast. These forwarding classes are grouped under three priority groups—data_control_pg, fcoe_noloss_pg, and multicast_pg.

For both Campus Switching ELS with *Hierarchical Port Scheduling* and Data Center Switching, you can modify and customize each of these priority groups and forwarding classes. For more details, see [“Creating and Managing Wired CoS Profiles”](#) on page 613.

Data Center Switching CoS Configuration

For data center switching devices, these additional CoS features are available:

- Hierarchical Port Scheduling (ETS)—Hierarchical port scheduling (Enhanced Transmission Selection, or ETS) is a two-tier process that provides better port bandwidth utilization and greater flexibility to allocate resources to queues and to groups of queues.
- Priority-based flow control (PFC)—A link-level flow control mechanism.

How Do I Implement Class of Service?

CoS can be implemented from the MSS CLI, from Network Director. RingMaster configures unicast traffic but does not configure multicast traffic. For directions to implement CoS from Network Director, see [“Creating and Managing Wired CoS Profiles” on page 613](#).

Editing Discovered CoS Profiles

Duplicate scheduler configuration is deployed to the device when you edit a CoS profile that are automatically created by Network Director as part of device discovery or out-of-band changes. In CoS configuration, a single classifier can be associated to multiple ports regardless of the other CoS configuration. When Network Director discovers a device with such configuration it will create multiple profiles, based on the difference in other CoS configurations, and mapped to same classifier configuration. If you modify classifier settings in such a CoS profile that is created automatically by Network Director, Network Director cannot modify the configuration because it is mapped to multiple profiles. Whenever you modify such a CoS profile that is created automatically, Network Director will create new classifier settings configuration on the device and map the same to it, without affecting the existing classifier settings. Newly created classifier settings will have a name generated based on the profile name. Even if only one profile is mapped to the classifier settings, Network Director creates new classifier settings and the old settings are orphaned.

NOTE: This behavior is applicable to both hierarchical and non hierarchical profiles, and is applicable for congestion notification profile name, traffic control profile name, scheduler map name, classifier name and rewrite rule settings.

RELATED DOCUMENTATION

[Creating and Managing Wired CoS Profiles | 613](#)
[Network Director Documentation home page](#)

Creating and Managing Wired CoS Profiles

IN THIS SECTION

- [Managing Wired CoS Profiles | 614](#)
- [Using the Default CoS Profiles for Switches | 615](#)

- [Using the Default CoS Profiles for Data Center Switching | 615](#)
- [Creating a Wired CoS Profile | 616](#)
- [Specifying Settings for a Switching and Campus Switching ELS CoS Profile | 617](#)
- [Specifying Settings for a Data Center Switching CoS Profile | 621](#)
- [What to Do Next | 630](#)

CoS profiles enable the grouping of class-of-service (CoS) parameters and apply them to one or more interfaces. Network Director provides you with predefined traffic types for each CoS profile that you create. These traffic types represent the most common types of traffic for the device type. Each of these templates has preconfigured values for all CoS parameters based on the typical application requirements. You can change the preconfigured values of these parameters to suit your requirements.

This topic describes:

Managing Wired CoS Profiles

From the Manage CoS Profiles page, you can:

- Create a new CoS profile by clicking **Add**. For details, see [“Creating a Wired CoS Profile” on page 616](#).
- Modify an existing CoS profile by selecting it and clicking **Edit**.
- View information about a profile by selecting the profile and clicking **Details**.
- Delete a CoS profile by selecting a profile and clicking **Delete**.

TIP: You cannot delete profiles that are in use—that is, assigned to objects or used by other profiles. To see the current assignments for a profile, select the profile and click **Details**.

- Clone an existing CoS profile by selecting it and clicking **Clone**.

[Table 127](#) describes the information provided about wired CoS profiles on the Manage CoS Profiles page. This page lists all CoS profiles defined for your network, regardless of the scope you selected in the network view.

Table 127: Managing Wired CoS Profile Fields

Field	Description
Profile Name	Name given to the profile when the profile was created.

Table 127: Managing Wired CoS Profile Fields (continued)

Field	Description
Family Type	The device family on which the profile was created: EX Series Switches, Campus Switching ELS, or Data Center Switching.
Description	Description of the profile that was entered when the profile was created. If the profile was created by using the CLI and then discovered by Network Director, the description is <i>Profile created as part of device discovery</i> . TIP: To display the entire description, you might need to resize the Description column by clicking the column border in the heading and dragging it.
Creation Time	Date and time when the profile was created.
Update Time	Date and time when the profile was last modified.
User Name	The username of the user who created or modified the profile.

TIP: All columns might not be displayed. To show or hide fields listed in the Manage Authorization Profiles table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

Using the Default CoS Profiles for Switches

When you install Network Director, a default CoS profile (juniper_CoS_template) is added to the Manage CoS Profiles page for EX Series switches and another with the same name is added for Campus Switching ELS. Default CoS profiles have most basic settings preconfigured. For example, the forwarding classes in the default CoS profile have already been assigned with default scheduler values. However, you can use the Edit CoS Profile page to optimize your communication with the network by customizing the bandwidth and buffer size assigned to each of the forwarding classes in the default CoS profile.

Using the Default CoS Profiles for Data Center Switching

When you install Network Director, the following default CoS profiles are installed for Data Center Switching:

- juniper_DC_NonHier_Ethernet_CoS
- juniper_DC_Hier_Ethernet_CoS
- juniper_DC_NonHier_CoS

- juniper_DC_Hier_CoS
- juniper_DC_Hier_FCoE_CoS
- Juniper_DC_Hier_CoS_Fusion


To see the settings configured for a default profile, select it on the Manage CoS Profiles page, then click **Details**.

Creating a Wired CoS Profile

In Network Director, you can create a CoS profile to group a set of Class of Service parameters and apply it to one or more network sessions.

For a CoS profile, you must specify the profile name. You can use defaults for the other values.

To create a wired CoS profile:

1. Click  in the Network Director banner.
2. Under Select View, select one of the following: **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Virtual View** or **Topology View**.

3. From the Tasks pane, expand **Wired**, expand **Profiles**, and then select **CoS**.
4. Click **Add** to add a new profile.
Network Director opens the Device Family Chooser window.
5. From the Device Family Chooser, select the wired device family for which you want to create a profile. The available device families are **Switching (EX)**, **Campus Switching ELS**, and **Data Center Switching**.
6. Click OK.
7. Complete the appropriate settings using the steps mentioned in [“Specifying Settings for a Switching and Campus Switching ELS CoS Profile” on page 617](#), or [“Specifying Settings for a Data Center Switching CoS Profile” on page 621](#).

Specifying Settings for a Switching and Campus Switching ELS CoS Profile

Create a CoS profile for switching by providing a profile name and, optionally, changing any default settings for Traffic Configuration and Shaping.

1. Enter the CoS switching settings described in [Table 128](#). Required settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 128: CoS Profile Settings for EX and Campus Switching ELS

Field	Action
Profile Name	Type the name of the profile. You can use up to 64 characters for profiles created for wired devices. Profile name must not contain special characters or spaces. Note that profiles that are automatically created by Network Director as part of device discovery or out-of-band changes may contain the underscore (_) character.
Description	Type a description of the profile.

2. Network Director includes four predefined traffic types, Data, Voice, Video, and Network Control. You can either modify those traffic types or you can create your own traffic type. Modify and customize any listed traffic type by selecting the traffic type from the list and clicking **Edit**, then changing any of the settings described in [Table 129](#).
3. To create your own traffic type, click **Add** and then configure the settings described in [Table 129](#).

Table 129: Traffic Configuration and Shaping for EX Switching and Campus Switching ELS

Field	Description
Traffic Type	If you are editing a Network Director default traffic type, this field cannot be changed. If you are adding a traffic type, indicate the type of traffic—this can be any value, such as a server name or something to do with your business.

Table 129: Traffic Configuration and Shaping for EX Switching and Campus Switching ELS (continued)

Field	Description
Forwarding Name	<p>If you are editing a Network Director default traffic type, this field cannot be changed. If you are adding a traffic type, you can use one of the predefined forwarding classes for your switch or you can create your own forwarding class. These forwarding classes are always provided: nd_best-effort, nd_network-control, nd_video-forwarding, and nd_expedited-forwarding. To create your own forwarding class, type a name instead of selecting an option.</p> <p>Most switches support the four predefined forwarding classes listed above. The exception is the EX4300 switch, which has eight default forwarding classes, including the standard four classes, plus multicast-network-connect, multicast-assured-forwarding, multicast-expedited-forwarding, and multicast-network-connect.</p>
Forwarding Queue	<p>Existing forwarding classes already have associated queues that cannot be altered. If you defined a new forwarding class by specifying your own Forwarding Name, then select an internal queue number to which forwarding classes are assigned. Most switches support queues 0 - 10. The exception is the EX4300 switch, which supports queues 0 - 11.</p> <p>By default, if a packet is not classified, it is assigned to the class associated with queue 0. You can assign more than one forwarding class to a queue number.</p>
Scheduler Map	
<p>A note in the Scheduler Map section indicates how much buffer size and bandwidth you have available to configure. For example, the message “You have been left with 0 percent buffer size and 0 percent bandwidth.” means that you have no available buffer or bandwidth, and you must reconfigure existing traffic types to free some bandwidth before configuring additional traffic types.</p>	
Low Priority	Enable Low Priority if you want the queue to receive low priority.

Table 129: Traffic Configuration and Shaping for EX Switching and Campus Switching ELS (continued)

Field	Description
Strict High Priority	<p>Enable Strict High Priority if you want the queue to receive preferential treatment over a low-priority queue. Unlimited bandwidth is assigned to a strict-high priority queue.</p> <p>A strict-high priority queue receives preferential treatment over a low-priority queue. Unlimited bandwidth is assigned to a strict-high priority queue. Queues are scheduled according to the queue number, starting with the highest queue, 7, with decreasing priority down through queue 0. Traffic in higher-numbered queues is always scheduled prior to traffic in lower-numbered queues. In other words, in case of two high-priority queues, the queue with the higher queue number is processed first.</p> <p>NOTE: You can modify this field in the Traffic Configuration and Shaping table or from the Traffic Configuration and Shaping window.</p>
Buffer Size (%)	<p>Buffer Size (%) is the size of the memory buffer allocated for storing packets. Use the slider to specify the scheduler Buffer Size percentage.</p> <p>NOTE: You can modify this value by double-clicking this field in the Traffic Configuration and Shaping table or by sliding the bar in the Traffic Configuration and Shaping window.</p>
Bandwidth Reserved (%)	<p>Bandwidth Reserved (%) is the amount of interface bandwidth assigned to the queue. Move the slider to specify the Bandwidth Reserved percentage. Defaults are:</p> <ul style="list-style-type: none"> • Data: 30% • Voice: Strict High • Video: 70% • Network control: 0% <p>If Strict-High is enabled for this traffic type, you cannot reserve bandwidth.</p> <p>NOTE: This field displays the value based on either your input or on the transmit-rate parameter from the switch, if that parameter is configured. While specifying transmit-rate on the EX Series switch, if you choose to specify the value as an exact rate, Network Director converts this value and displays it as a percentage in the Bandwidth Reserved (%) field. You can modify this percentage value from the CoS Profile page.</p>
Shaping Rate	<p>Move the Shaping Rate slider to throttle the rate of packet transmission by setting a maximum bandwidth (rate in bits per second) or a maximum percentage of bandwidth for a queue or a forwarding class.</p>

Table 129: Traffic Configuration and Shaping for EX Switching and Campus Switching ELS (continued)

Field	Description
Traffic Classification	
Behavior aggregate classification classifies packets. The DSCP or DSCP IPv6 precedence bits of the IP header convey the behavior aggregate class information. The information might also be found in the IEEE 802.1ad, or IEEE 802.1p CoS bits.	
Classifier Type	<p>Select a classifier type—DSCP, DSCP-IPv6, INET-precedence, or IEEE-802.1—and associate the corresponding code-point aliases to loss priorities.</p> <p>NOTE: You can specify code-point—loss priority associations for one or more classifier types.</p> <ul style="list-style-type: none"> • DSCP—Differentiated services code point, a field in IPv4 headers, is used to classify traffic. • DSCP-IPv6—Differentiated services code point, a field in IPv6 headers, is used to classify traffic. • INET precedence—Field that indicates class of service rewrite rules are used to classify traffic. • IEEE-802.1—IEEE 802.1ad, or IEEE 802.1p CoS bits are used to classify traffic.
Classifier Code Points	
Code Points	<p>The code points list includes all available and unselected code points for the selected classifier type.</p> <p>Specify one or more code-point aliases or bit sets to associate with a forwarding class by moving the value to one of the two lists, Loss Priority Low or Loss Priority High.</p>
Loss Priority Low	Indicate that packets have low loss priority by selecting code-point aliases from the Code Points table and using the LEFT and RIGHT arrows to move them into the appropriate loss priority table.
Loss Priority Medium-Low	Indicate that packets have medium-low loss priority by selecting code-point aliases from the Code Points table and using the LEFT and RIGHT arrows to move them into the appropriate loss priority table.
Loss Priority Medium-High	Indicate that packets have medium-high loss priority by selecting code-point aliases from the Code Points table and using the LEFT and RIGHT arrows to move them into the appropriate loss priority table.

Table 129: Traffic Configuration and Shaping for EX Switching and Campus Switching ELS (continued)

Field	Description
Loss Priority High	Indicate that packets have high loss priority by selecting code-point aliases from the Code Points table and using the LEFT and RIGHT arrows to move them into the appropriate loss priority table.

- Click **OK** to close the Add Traffic and Classification window and save your configuration.

Your changes are added to this CoS profile.

NOTE: If all bandwidth has already been reserved, your changes are not made. Reduce the bandwidth reserved from another Traffic Type, then repeat the configuration.

- To configure rewrite rules for a forwarding queue, click **Configure Rewrite Rules** at the bottom of the screen. The Configure Rewrite Rules window appears. Specify rewrite rule settings as described below to alter CoS values in outgoing packets on the outbound interfaces of an edge switch:
 - Select the forwarding class for which you want to create or modify rewrite rules. Network Director lists all the forwarding classes that you have used for configuring traffic in the Traffic Configuration and Shaping section.
 - For each classifier's loss priority, select a code-point alias for each loss-priority type—Low, Medium-Low, Medium-High, and High.
- Click **OK** to save the rewrite rules and close the Configure Rewrite Rules window.

The system saves the rewrite rules and returns to the **Create CoS Profile** page.

- Click **Done**.

After you create a CoS profile for switching devices, associate the CoS profile with a Port profile. For directions, see [“Creating and Managing Port Profiles” on page 414](#).

Specifying Settings for a Data Center Switching CoS Profile

You can create a CoS profile by specifying the profile settings and the traffic configuration and shaping details.

To specify the settings for the CoS profile:

1. Enter the settings described in [Table 130](#). Required settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 130: CoS Profile Basic Settings for Data Center Switching

Field	Action
Profile Name	Type the name of the profile. You can use up to 64 characters for profiles created for wired devices. Profile name must not contain special characters or spaces. Note that profiles that are automatically created by Network Director as part of device discovery or out-of-band changes may contain the underscore (_) character.
Description	Type the description of the profile.

2. In the Traffic Classification and Shaping Settings section, select one of these options:
 - **Hierarchical Port Scheduling (ETS)**—Hierarchical port scheduling (Enhanced Transmission Selection, or ETS) is a two-tier process that provides better port bandwidth utilization and greater flexibility to allocate resources to queues and to groups of queues (for QFX and QFabric devices).
 - **Non Hierarchical Port Scheduling**—Non-hierarchical scheduling is a one-tier process that provides port bandwidth utilization and allocates resources to queues (for EX4500 and EX4550 transit switches).
 - **Hierarchical (Fusion)**—Select this scheduling type if you plan to assign the CoS profile to QFX10002 and QFX10008 switches
3. If you selected Hierarchical Port Scheduling (ETS), specify settings in the Priority Group and Traffic Settings section.

The table lists priority groups and the forwarding classes they contain in an expandable list. Priority groups refer to forwarding class sets in the device. You can perform these tasks on priority groups and forwarding classes:

- To add a new priority group, click **Add Priority Group**. The Add Priority Group and Traffic Control Profile Window opens. Enter the settings as described in [Table 131](#).

Table 131: Add Priority Group and Traffic Control Profile Window

Field	Description
Priority Group Name	Enter a name for the priority group.
Traffic Control Profile Settings	
Transmit Rate (%)	Select a transmit rate percentage for the priority group.
Shaping Rate (%)	Select a shaping rate percentage for the priority group.

- To edit a priority group or forwarding class's properties, click the field that you want to edit in the table.
- To edit a forwarding class's properties, click its name. The Edit Traffic Classification and Shaping for priority group window opens. Enter the settings as described in [Table 132](#).

Table 132: Edit and Add Traffic Classification and Shaping for Priority Group Window

Field	Description
Forwarding Class Name	Select or specify a name for the forwarding class.
Forwarding Class Queue	Specify the internal queue numbers to which forwarding classes are assigned.
No Loss	Select to make the forwarding class lossless.
Scheduler Map	
Strict High	Select if you want the queue to receive preferential treatment over a low-priority queue. Unlimited bandwidth is assigned to a strict-high priority queue.
Transmit Rate	Select the percentage of interface bandwidth assigned to the forwarding class. If you have enabled Strict-High , you cannot reserve bandwidth for this traffic type.
Shaping Rate	Select a shaping rate percentage for the forwarding class.
Buffer Size (%)	Select the percentage of the memory buffer allocated for storing packets for the forwarding class.
Traffic Classification	
Classifier Type	Select the classifier type that maps packets to a forwarding class and a loss priority.

Table 132: Edit and Add Traffic Classification and Shaping for Priority Group Window (continued)

Field	Description
Code Points	Specify one or more code-points for associating with a forwarding class.
Loss Priority Low	Indicates that packets have low loss priority. Select code points from the Code Points table and use the LEFT and RIGHT arrows to move them into the appropriate loss priority table.
Loss Priority Medium High	Indicates that packets have medium high loss priority. Select code points from the Code Points table and use the LEFT and RIGHT arrows to move them into the appropriate loss priority table.
Loss Priority High	Indicates that packets have high loss priority. Select code points from the Code Points table and use the LEFT and RIGHT arrows to move them into the appropriate loss priority table.

- To add a forwarding class to a priority group, click the **Add Forwarding Class** link at the end of the priority group's list of forwarding classes. The Add Traffic Classification and Shaping for priority group window opens. Enter the settings as described in [Table 132](#).
 - To remove a priority group or forwarding class, click the **X** at the end of its table row.
4. If you selected Non Hierarchical Port Scheduling, specify settings in the Traffic Configuration and Shaping table.

The table lists forwarding classes. You can perform these tasks on forwarding classes:

- To add traffic configuration and shaping details for different types of traffic, click **Add** in the Traffic Configuration and Shaping box. The Add Traffic Classification and Shaping window opens.
- To modify the details of an existing traffic configuration, select the traffic configuration from the list and click **Edit**. The Edit Traffic Classification and Shaping window opens.

NOTE: You can modify some of the details in the Traffic Configuration and Shaping table without having to open the Edit Traffic Classification and Shaping window—by clicking on the field that you want to modify.

- To delete a traffic configuration entry, select the traffic configuration from the list and click **Remove**. The system deletes the selected traffic configuration entry.

To create your own traffic type, click **Add** and then configure the settings described in [Table 133](#).

Table 133: Traffic Configuration and Shaping for EX Switching and Campus Switching ELS

Field	Description
Traffic Type	If you are editing a Network Director default traffic type, this field cannot be changed. If you are adding a traffic type, indicate the type of traffic—this can be any value, such as a server name or something to do with your business.
Forwarding Name	<p>If you are editing a Network Director default traffic type, this field cannot be changed. If you are adding a traffic type, you can use one of the predefined forwarding classes for your switch or you can create your own forwarding class. These forwarding classes are always provided: nd_best-effort, nd_network-control, nd_video-forwarding, and nd_expedited-forwarding. To create your own forwarding class, type a name instead of selecting an option.</p> <p>Most switches support the four predefined forwarding classes listed above. The exception is the EX4300 switch, which has eight default forwarding classes, including the standard four classes, plus multicast-network-connect, multicast-assured-forwarding, multicast-expedited-forwarding, and multicast-network-connect.</p>
Forwarding Queue	<p>Existing forwarding classes already have associated queues that cannot be altered. If you defined a new forwarding class by specifying your own Forwarding Name, then select an internal queue number to which forwarding classes are assigned. Most switches support queues 0 - 10. The exception is the EX4300 switch, which supports queues 0 - 11.</p> <p>By default, if a packet is not classified, it is assigned to the class associated with queue 0. You can assign more than one forwarding class to a queue number.</p>
Scheduler Map	
<p>A note in the Scheduler Map section indicates how much buffer size and bandwidth you have available to configure. For example, the message “You have been left with 0 percent buffer size and 0 percent bandwidth.” means that you have no available buffer or bandwidth, and you must reconfigure existing traffic types to free some bandwidth before configuring additional traffic types.</p>	
Low Priority	Enable Low Priority if you want the queue to receive low priority.

Table 133: Traffic Configuration and Shaping for EX Switching and Campus Switching ELS (continued)

Field	Description
Strict High Priority	<p>Enable Strict High Priority if you want the queue to receive preferential treatment over a low-priority queue. Unlimited bandwidth is assigned to a strict-high priority queue.</p> <p>A strict-high priority queue receives preferential treatment over a low-priority queue. Unlimited bandwidth is assigned to a strict-high priority queue. Queues are scheduled according to the queue number, starting with the highest queue, 7, with decreasing priority down through queue 0. Traffic in higher-numbered queues is always scheduled prior to traffic in lower-numbered queues. In other words, in case of two high-priority queues, the queue with the higher queue number is processed first.</p> <p>NOTE: You can modify this field in the Traffic Configuration and Shaping table or from the Traffic Configuration and Shaping window.</p>
Buffer Size (%)	<p>Buffer Size (%) is the size of the memory buffer allocated for storing packets. Use the slider to specify the scheduler Buffer Size percentage.</p> <p>NOTE: You can modify this value by double-clicking this field in the Traffic Configuration and Shaping table or by sliding the bar in the Traffic Configuration and Shaping window.</p>
Bandwidth Reserved (%)	<p>Bandwidth Reserved (%) is the amount of interface bandwidth assigned to the queue. Move the slider to specify the Bandwidth Reserved percentage. Defaults are:</p> <ul style="list-style-type: none"> • Data: 30% • Voice: Strict High • Video: 70% • Network control: 0% <p>If Strict-High is enabled for this traffic type, you cannot reserve bandwidth.</p> <p>NOTE: This field displays the value based on either your input or on the transmit-rate parameter from the switch, if that parameter is configured. While specifying transmit-rate on the EX Series switch, if you choose to specify the value as an exact rate, Network Director converts this value and displays it as a percentage in the Bandwidth Reserved (%) field. You can modify this percentage value from the CoS Profile page.</p>
Shaping Rate	<p>Move the Shaping Rate slider to throttle the rate of packet transmission by setting a maximum bandwidth (rate in bits per second) or a maximum percentage of bandwidth for a queue or a forwarding class.</p>

Table 133: Traffic Configuration and Shaping for EX Switching and Campus Switching ELS (continued)

Field	Description
Traffic Classification	
Behavior aggregate classification classifies packets. The DSCP or DSCP IPv6 precedence bits of the IP header convey the behavior aggregate class information. The information might also be found in the IEEE 802.1ad, or IEEE 802.1p CoS bits.	
Classifier Type	<p>Select a classifier type—DSCP, DSCP-IPv6, INET-precedence, or IEEE-802.1—and associate the corresponding code-point aliases to loss priorities.</p> <p>NOTE: You can specify code-point—loss priority associations for one or more classifier types.</p> <ul style="list-style-type: none"> • DSCP—Differentiated services code point, a field in IPv4 headers, is used to classify traffic. • DSCP-IPv6—Differentiated services code point, a field in IPv6 headers, is used to classify traffic. • INET precedence—Field that indicates class of service rewrite rules are used to classify traffic. • IEEE-802.1—IEEE 802.1ad, or IEEE 802.1p CoS bits are used to classify traffic.
Classifier Code Points	
Code Points	<p>The code points list includes all available and unselected code points for the selected classifier type.</p> <p>Specify one or more code-point aliases or bit sets to associate with a forwarding class by moving the value to one of the two lists, Loss Priority Low or Loss Priority High.</p>
Loss Priority Low	Indicate that packets have low loss priority by selecting code-point aliases from the Code Points table and using the LEFT and RIGHT arrows to move them into the appropriate loss priority table.
Loss Priority Medium-Low	Indicate that packets have medium-low loss priority by selecting code-point aliases from the Code Points table and using the LEFT and RIGHT arrows to move them into the appropriate loss priority table.
Loss Priority Medium-High	Indicate that packets have medium-high loss priority by selecting code-point aliases from the Code Points table and using the LEFT and RIGHT arrows to move them into the appropriate loss priority table.

Table 133: Traffic Configuration and Shaping for EX Switching and Campus Switching ELS (continued)

Field	Description
Loss Priority High	Indicate that packets have high loss priority by selecting code-point aliases from the Code Points table and using the LEFT and RIGHT arrows to move them into the appropriate loss priority table.

5. If you selected Hierarchical Port Scheduling (ETS), specify priority-based flow control (PFC) settings in the PFC Settings section. Enter the settings as described in [Table 134](#).

Table 134: PFC Settings for Data Center Switching Hierarchical Port Scheduling (ETS) CoS Profile

Field	Description
Input Cable Length (meter)	Enter the length of the cable attached to the input interface, in meters.
Input	
Add	Click to add an input congestion notification profile (CNP). A new entry appears in the table.
Remove	Click to remove the selected input CNP.
IEEE Code Point	Select the IEEE code point for the input CNP.
Maximum Receive Size (bytes)	Enter the maximum receive unit (MRU) on an interface for traffic that matches the PFC priority, in bytes.
Output	
Add	Click to add an output CNP. A new entry appears in the table.
Remove	Click to remove the selected output CNP.
IEEE Code Point	Select the IEEE code point for the output CNP.
Queue List	Select output queues on which to enable flow control (PFC pause).

6. If you selected Non-Hierarchical Port Scheduling, specify priority-based flow control (PFC) settings in the PFC Settings section. Enter the settings as described in [Table 135](#).

Table 135: PFC Settings for Data Center Switching Non-Hierarchical Port Scheduling CoS Profile

Field	Description
Input	
Add	Click to add an input congestion notification profile (CNP). A new entry appears in the table.
Remove	Click to remove the selected input CNP.

7. If you selected Hierarchical Port Scheduling (ETS), specify rewrite rule settings in the Rewrite Rule Settings section as described in [Table 136](#).

Table 136: Rewrite Rule Settings for Data Center Switching CoS Profile

Field	Description
Forwarding Name	The name of the forwarding class.
Queue	The number corresponding to the forwarding queue. You cannot modify this field.
Rewrite Type	Select a rewrite-rules mapping for the traffic that passes through the various queues on the interface.
Egress Code Point - Loss Priority Low	Specify a code-point for association with a forwarding class for loss priority low.
Egress Code Point - Loss Priority Medium High	Specify a code-point for association with a forwarding class for loss priority medium high.
Egress Code Point - Loss Priority High	Specify a code-point for association with a forwarding class for loss priority high.

8. If you selected Non-Hierarchical Port Scheduling, click **Configure Rewrite Rules** at the bottom of the screen to configure rewrite rules for a forwarding queue. The Configure Rewrite Rules window appears. Specify rewrite rule settings as described below to alter CoS values in outgoing packets on the outbound interfaces of an edge switch:
- Select the forwarding class for which you want to create or modify rewrite rules. Network Director lists all the forwarding classes that you have used for configuring traffic in the Traffic Configuration and Shaping section.

- b. For each classifier's loss priority, select a code-point alias for each loss-priority type—Low, Medium-Low, Medium-High, and High.

9. If you selected Hierarchical (Fusion) Scheduling

10. Click **Done** to save the changes to the profile.

What to Do Next

After you have created a CoS profile for switching devices, you can associate the CoS profile to a Port profile.

RELATED DOCUMENTATION

[Understanding Class of Service \(CoS\) Profiles | 609](#)

[Creating and Managing Wireless CoS Profiles | 630](#)

Creating and Managing Wireless CoS Profiles

IN THIS SECTION

- [Managing Wireless CoS Profiles | 631](#)
- [Creating a Wireless CoS Profile | 632](#)
- [Specifying Settings for a Wireless CoS Profile | 633](#)
- [What To Do Next | 635](#)

CoS profiles enable the grouping of class of service (CoS) parameters and apply them to one or more wireless interfaces. Network Director provides you with predefined traffic types for each wireless CoS profile that you create. These traffic types represent the most common types of traffic for the device type. Each of these templates has preconfigured values for all CoS parameters based on the typical application requirements. You can change the preconfigured values of these parameters to suit your requirements.

This topic describes:

Managing Wireless CoS Profiles

From the Manage CoS Profiles page, you can:

- Create a new wireless CoS profile by clicking **Add**. For directions, see [“Creating a Wireless CoS Profile” on page 632](#).
- Modify an existing wireless CoS profile by selecting it and clicking **Edit**.
- Assign an existing wireless CoS profile to controllers by selecting it and clicking **Assign**. For directions, see [“Assigning a Wireless CoS Profile to Controllers” on page 635](#).
- Reassign an existing wireless CoS profile by selecting it and clicking **Edit Assignment**.
- View information about a profile by selecting the profile and clicking **Details**.
- Delete a wireless CoS profile by selecting a profile and clicking **Delete**.

TIP: You cannot delete profiles that are in use—that is, assigned to objects or used by other profiles. To see the current assignments for a profile, select the profile and click **Details**.

- Clone an existing wireless CoS profile by selecting it and clicking **Clone**.

[Table 137](#) describes the information provided about wireless CoS profiles on the Manage CoS Profiles page. This page lists all wireless CoS profiles defined for your network, regardless of the scope you selected in the network view.

Table 137: Manage Wireless CoS Profile Fields

Field	Description
Profile Name	Name given to the profile when the profile was created.
Family Type	The device family on which the profile was created, either EX Series Switches, Wireless LAN Controllers, Campus Switching ELS, or Data Center Switching.
Description	Description of the profile that was entered when the profile was created. If the profile was created by using the CLI and then discovered by Network Director, the description is <i>Profile created as part of device discovery</i> . TIP: To display the entire description, you might need to resize the Description column by clicking the column border in the heading and dragging it.
Creation Time	Date and time when the profile was created.
Last Updated Time	Date and time when the profile was last modified.

Table 137: Manage Wireless CoS Profile Fields (continued)

Field	Description
User Name	The username of the user who created or modified the profile.

TIP: All columns might not be displayed. To show or hide fields, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.


Creating a Wireless CoS Profile

In Network Director, you can create a wireless CoS profile to group a set of Class of Service parameters and apply it to one or more wireless network sessions.

For a CoS profile, you must specify the profile name. You can optionally reconfigure the default settings.

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.
3. In the Tasks pane, expand **Wireless**, expand **Profiles**, and then select **CoS**.
The Manage CoS Profiles window opens.
4. In the Manage CoS Profiles, click **Add** to add a new wireless CoS profile.
5. Specify the wireless CoS Profile settings as described in both the online help and in [“Specifying Settings for a Wireless CoS Profile” on page 633](#).
6. Click **Done** to save the wireless CoS profile.
The system saves the CoS profile and displays it on the Manage CoS Profiles page.

Specifying Settings for a Wireless CoS Profile

To configure a wireless CoS profile:

1. Enter the settings described in [Table 138](#). Required settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 138: CoS Profile Basic Settings for Wireless

Field	Action
Profile Name	Type a unique name that identifies the profile. You can use up to 32 characters for profiles created for wireless devices. Profile name must not contain special characters or spaces. Note that profiles that are automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character.
Description	Type a description of the profile.
Session CoS	
Enable Bandwidth Limit	Select Enable Bandwidth Limit if you want to specify a bandwidth limit for a given session, and then type the Maximum Bandwidth (Kb/s) (full duplex rate) for aggregates of access categories (ACs) for a wireless client. Downstream packets are shaped and upstream packets are policed.
Enable Static CoS	Select Enable Static CoS to assign the same CoS level to all traffic on the Service profile SSID, regardless of 802.1p or DSCP markings in the packets themselves, and regardless of any filters that mark CoS. This option provides a simple way to configure an SSID for priority traffic such as VoIP traffic. Select the CoS Value that you want the controller to assign to all user traffic.
Trust Client DSCP	Select Trust Client DSCP to enable the controller to use the client DSCP (Differentiated Services code point) for radio ingress traffic and ignore Wi-Fi Multimedia (WMM).
Voice CoS	

Table 138: CoS Profile Basic Settings for Wireless (continued)

Field	Action
Traffic Class	<p>Specify the traffic class for the data type voice:</p> <ul style="list-style-type: none"> ● None—Select to indicate that no traffic class is required for voice traffic. ● VoIP-Data—Select to indicate that you want to enable a traffic class for voice traffic. <p>If you selected VoIP-Data, you can additionally select Enable Bandwidth Limit if you want to specify a bandwidth limit for voice traffic, and then type the maximum bandwidth (full duplex rate) for aggregates of access categories (ACs) for a wireless client. Downstream packets are shaped and upstream packets are policed.</p> <p>If you selected VoIP-Data, you can additionally select Enable Static CoS to assign the same CoS level to all traffic on the Service profile SSID, regardless of 802.1p or DSCP markings in the packets themselves, and regardless of any firewall filters that mark CoS. This option provides a simple way to configure an SSID for priority traffic such as VoIP traffic. Select the CoS value that you want the controller to assign to all user traffic.</p>
Access Categories	<p>Access Categories includes up to four access categories for QoS. You can modify the action (Permit or Demote) corresponding to each forwarding queue to suit your requirements.</p> <p>If you remove any of the four default types (Background, Best Effort, Voice, and Video) by selecting it and clicking Remove, you can add additional categories until there are four. To add a category:</p> <ol style="list-style-type: none"> a. Click Add in the Access Categories box. The Add Access Category window appears. b. Select an Access Category, Background, Best-Effort, Voice, or Video. c. Select an action to be applied to packets in the forwarding queue, either Permit or Demote. d. Click OK. <p>The Add Access Category window closes and the access category is added to the list of Access Categories.</p>

2. Click **Done**.

What To Do Next

After you create a CoS profile for wireless devices, associate the CoS profile with an Authorization profile during Authorization profile creation. This Authorization profile is then associated with a WLAN profile (during WLAN profile creation) to apply the CoS settings to all the users who connect to the WLAN SSID. For directions, see [“Creating and Managing Wireless Authorization Profiles” on page 395](#) and [“Creating and Managing a WLAN Service Profile” on page 1090](#).

RELATED DOCUMENTATION

[Understanding Class of Service \(CoS\) Profiles | 609](#)

[Assigning a Wireless CoS Profile to Controllers | 635](#)

[Creating and Managing Wired CoS Profiles | 613](#)

[Network Director Documentation home page](#)

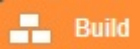
Assigning a Wireless CoS Profile to Controllers

CoS profiles enable the grouping of class of service (CoS) parameters and apply them to one or more wireless interfaces. Network Director provides you with predefined traffic types for each wireless CoS profile that you create. These traffic types represent the most common types of traffic for the device type. You must have an existing CoS Profile to complete this task—for directions, see [“Creating and Managing Wireless CoS Profiles” on page 630](#).

To assign a CoS Profile to a controller:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.
3. In the View pane, select one or more controllers. If you select **Wireless Network**, you select all controllers.
4. In the Tasks pane, expand **Wireless**, expand **Profiles**, and then select **CoS**.

The Manage CoS Profiles page appears, displaying the list of currently configured CoS Profiles.

5. Select a CoS Profile from the list, and then click **Assign**.

The Assign CoS Profile wizard opens. The wizard consists of three sections, **Device Selection**, **Profile Assignment**, and **Review**.

6. From **Device Selection**, select one or more devices for CoS Profile assignment—all devices under the selection are also selected.

7. Click **Profile Assignment** to open the next page of the wizard.

The Assign CoS Profile page opens, displaying a list of selected devices.

8. From **Profile Assignment**, select one or more devices from the list by placing a check mark next to the devices. Click either **Assign to Device** or **Assign to Cluster**.

9. Click **Review** and then click **Edit** to make any needed changes.

10. Click **Finish**.

The new assignment appears in the list.

RELATED DOCUMENTATION

[Creating and Managing Wireless CoS Profiles | 630](#)

[Understanding Class of Service \(CoS\) Profiles | 609](#)

[Creating and Managing Wired CoS Profiles | 613](#)

[Network Director Documentation home page](#)

Configuring Media Access Control Security (MACsec)

IN THIS CHAPTER

- [Media Access Control Security Overview | 637](#)
- [Configuring and Managing MACsec Profiles | 638](#)
- [Assigning the MACsec Profiles | 644](#)

Media Access Control Security Overview

Media Access Control Security (MACsec) is an industry-standard security technology that provides secure communication on Ethernet links. MACsec enables you to secure Ethernet links between two MACsec-capable devices. You can enable MACsec on point-to-point Ethernet links using static secure association key (SAK) security mode or static connectivity association key (CAK) security mode.

When you enable MACsec using the static CAK security mode, a connectivity association key and a randomly generated secure association key are exchanged between the devices on each point-to-point Ethernet link. After the matching pre-shared keys are successfully exchanged, MACsec enables MKA protocol on the devices. The MKA protocol maintains MACsec on the link, and decides which switch on the point-to-point link becomes the key server. The key server then creates an SAK that is shared with the switch at the other end of the point-to-point link only, and that SAK is used to secure all data traffic traversing the link. The key server will continue to periodically create and share a randomly-created SAK over the point-to-point link for as long as MACsec is enabled. A pre-shared key includes a connectivity association name (CKN) and its own connectivity association key (CAK). The CKN and CAK are configured by the user in the connectivity association and must match on both ends of the link to initially enable MACsec.

NOTE: A pre-shared key includes a connectivity association name (CKN) and its own connectivity association key (CAK). You can configure the CKN and CAK in the connectivity association and these values must match on both ends.

When you enable MACsec using static SAK security mode, you must configure the secure channels between the point-to-point Ethernet link. The secure channels are responsible for transmitting and receiving data

on the MACsec-enabled link, and also responsible for transmitting SAKs across the link to enable and maintain MACsec. A typical connectivity association when MACsec is enabled using SAK security mode contains two secure channels—one secure channel for inbound traffic and another secure channel for outbound traffic. You must configure the SAK settings manually, there is no key server or other tool that creates SAKs. Security is maintained on the point-to-point Ethernet link by periodically rotating between the two security keys. Each security key name and value must have a corresponding matching value on the interface at the other end of the point-to-point Ethernet link to maintain MACsec on the link.

MACsec is widely used in campus deployments to secure network traffic between endpoints and access switches. You can enable MACsec on extended ports in a Junos Fusion Enterprise topology to provide secure communication between the satellite device and connected hosts. Network Director supports MACsec configuration for a Junos Fusion Enterprise setup. You can create a profile for the MACsec configuration and assign the profiles to the extended ports of the satellite devices in a Junos Fusion Enterprise setup.

For more information about MACsec, see [Understanding Media Access Control Security \(MACsec\)](#).

RELATED DOCUMENTATION

| [Understanding Junos Fusion Enterprise](#) | 710

Configuring and Managing MACsec Profiles

IN THIS SECTION

- [Creating a MACsec Profile](#) | 639
- [Specifying Settings for a MACSsec Profile](#) | 640
- [What to Do Next](#) | 643

From the MACsc Profile page of the Network Director UI you can create and manage MACsec profiles that specify MACsec settings for the extended ports in the aggregation device in a Junos Fusion Enterprise device. From the Manage MACsec Profile page, you can:

- Create a new MACsec profile by clicking **Add**.
- Modify an existing MACsec profile by selecting the profile and clicking **Edit**.
- Associate a profile to the extended ports by selecting the profile and clicking **Assign**.

- Change current assignments for a profile by selecting the profile and clicking **Edit Assignment**.
- Delete a MACsec profile by selecting the profile and clicking **Delete**.
- Clone an existing MACsec profile by selecting the profile and clicking **Clone**.
- View information about a profile by selecting the profile and clicking **Details**.

[Table 139](#) describes the information provided about wired MACsec profiles on the Manage MACsec Profiles page. This page lists all the MACsec profiles defined for the Junos Fusion Enterprise device, regardless of the scope you selected in the network view.

Table 139: Managing MACsec Profile Fields

Field	Description
Profile Name	Name of the profile.
Connection Association Name	Name of the MACsec connectivity association.
Description	Description of the profile.
MACsec Mode	Static secure association key (static-SAK) security mode or static connectivity association key (static-CAK) using which you enabled MACsec on the device.
Assignment State	Profile assignment state. One of the following: <ul style="list-style-type: none"> • Deployed—The profile has been assigned and the configuration has been deployed on the devices. • Pending Deployment—The profile has been assigned or its previous assignments have been changed, but the new or modified configuration has not yet been deployed on the devices. • Unassigned—The profile has not yet been assigned.
User Name	The username of the user who created or modified the profile.


This topic describes:

Creating a MACsec Profile

To create a MACsec profile:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View**, or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.
3. In the Tasks pane, expand **Wired**, expand **Profiles**, and then select **MACsec**.
The Manage MACsec Profile page appears, displaying the list of currently configured MACsec profiles.
4. Click **Add** to add a new profile.
The Create MACsec Profile page appears.
5. Enter the MACsec settings described in [“Specifying Settings for a MACSsec Profile” on page 640](#).
6. Click **Done**.

Specifying Settings for a MACSsec Profile

[Table 140](#) describes the MACsec Profile settings. Required settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 140: MACsec Profile Settings

Field	Action
Profile Name	Type the name of the profile.
Description	Type a description of the profile.
Connection Association Name	Type the name for the MACsec connectivity association.
MACsec Mode	Select the mode using which you can enable MACsec on the device. The available modes are static secure association key (static-SAK) security mode or static connectivity association key (static-CAK) security mode.
CAK Settings	If you want to enable MACsec by using the CAK mode, configure the CAK settings specified in Table 141 .

Table 140: MACsec Profile Settings (continued)

Field	Action
SAK Settings	If you want to enable MACsec by using the SAK mode, configure the SAK settings specified in Table 142 for the inbound and outbound secure channels.

Table 141: CAK Settings

Field	Description
Connectivity Association Key Name	Type a name for the connectivity association key that you want to use for enabling MACsec.
Connectivity Association Key	Specify the key to exchange with the other end of the link on the secure channel. You must use a hexadecimal string of 32 digits.
Confirm Connectivity Association Key	Specify the connectivity association key again. If there is a mismatch (between the connectivity association keys), an error message is shown.
Enable Include Secure Channel Identifier	Enable Include Secure Channel Identifier tagging on a device that is enabling MACsec on an Ethernet link connecting to an Junos Fusion Enterprise device.
Key Server Priority	Specify the MACsec Key Agreement (MKA) server election priority number. You can specify a value between 0 and 255. The lower the number, the higher the priority.
Transmit Interval (milli sec)	Specify the transmit interval for MACsec Key Agreement (MKA) protocol data units (PDUs). The MKA transmit interval setting sets the frequency for how often the MKA PDU is sent to the directly connected device to maintain MACsec on a point-to-point Ethernet link. A lower interval increases bandwidth overhead on the link; a higher interval optimizes the MKA protocol data unit exchange process. The default transmit interval is 2000 milliseconds
Disable Encryption	Select this option if you want to disable the MACsec encryption for a connectivity association that has MACsec already enabled on it.

Table 141: CAK Settings (continued)

Field	Description
Offset	<p>Specify the offset 0, 30, or 50 for all the packets traversing the link. The default offset is 0. All traffic in the connectivity association is encrypted when encryption is enabled and an offset is not set.</p> <p>When the offset is set to 30, the IPv4 header and the TCP/UDP header are unencrypted while encrypting the rest of the traffic.</p> <p>When the offset is set to 50, the IPv6 header and the TCP/UDP header are unencrypted while encrypting the rest of the traffic.</p> <p>You would typically forward traffic with the first 30 or 50 octets unencrypted if a feature needed to see the data in the octets to perform a function, but you otherwise prefer to encrypt the remaining data in the frames traversing the link. Load balancing features, in particular, typically need to see the IP and TCP/UDP headers in the first 30 or 50 octets to properly load balance traffic.</p>
Replay Window Size	<p>Specify the size of the replay protection window.</p> <p>NOTE: When this variable is set to 0, all packets that arrive out-of-order are dropped.</p>
Exclude Protocols	<p>Specify the name of the protocol that should not be MACsec-secured. Options include:</p> <ul style="list-style-type: none"> • cdp—Cisco Discovery Protocol. • lacp—Link Aggregation Control Protocol. • lldp—Link Level Discovery Protocol.

Table 142: SAK Settings

Field	Description
Secure Channel name	Type a name for the secure channel.
MAC address	Specify a MAC address on which you want to enable MACsec using static secure association key (SAK) security mode. The mac-address variables must match on the sending and receiving ends of a link to enable MACsec using static SAK security mode.
Port	<p>Specify the port ID number in a secure channel when enabling MACsec using static secure association key (SAK) security mode. The port IDs must match on a sending and receiving secure channel on each side of a link to enable MACsec.</p> <p>After the port numbers match, MACsec is enabled for all traffic on the connection.</p>

Table 142: SAK Settings (*continued*)

Field	Description
Enable Encryption	<p>Select this option if you want to Enable MACsec encryption within an outbound secure channel.</p> <p>NOTE: You can enable MACsec without enabling encryption. If a connectivity association with an outbound secure channel that has not enabled MACsec encryption is associated with an interface, traffic is forwarded across the Ethernet link in clear text. You are, therefore, able to view this unencrypted traffic when you are monitoring the link.</p>
Offset	<p>Specify the number of octets in an Ethernet frame that you want to send in unencrypted plain-text when encryption is enabled for MACsec.</p> <p>Setting the offset to 30 allows a feature to see the IPv4 header and the TCP/UDP header while encrypting the remaining traffic. Setting the offset to 50 allows a feature to see the IPv6 header and the TCP/UDP header while encrypting the remaining traffic.</p>
Secure Association	<p>Specify the secure association keys corresponding to the secure association number. The key string is a 32-digit hexadecimal number.</p> <p>Re-enter the secure association key for every secure association number. If there is a mismatch between the connectivity association key and their respective confirmation keys, an error message is shown.</p>

What to Do Next

After you create the MACsec profile, you must assign the profile to the Junos Fusion Enterprise satellite device by using the Manage MacSec Profile page and then deploy the Device profile by using the **Deploy** mode.

To assign a MACsec Settings profile to a device, see [“Assigning the MACsec Profiles” on page 644](#). For information about deploying the configurations, see [“Deploying Configuration to Devices” on page 1180](#).

NOTE: You can assign the MACsec profile to the extended ports on Junos Fusion Enterprise Aggregation Device.

In the CAK mode, if you change the connection association key name of a deployed MACsec profile, you must re-configure the connectivity association key and the confirmation key for that profile. Similarly, in the SAK mode, if you change the inbound or outbound channel names of the deployed MACSec profiles, you must re-configure the key and the confirmation key for that profile.

RELATED DOCUMENTATION

[Media Access Control Security Overview | 637](#)

[Understanding Junos Fusion Enterprise | 710](#)


Assigning the MACsec Profiles

IN THIS SECTION

- [Assigning a MACsec Profile to a Device | 644](#)
- [Editing the MACsec Profile Assignments | 645](#)

Assigning a MACsec Profile to a Device

You can assign a MACsec profile to extended ports of satellite devices only and not to the native ports on the aggregation devices. Therefore, the tree view shows only the satellite devices (standalone and that are part of the cluster) to which you can assign a profile.

1. Click  in the Network Director banner.
2. Under Select View, select one of the following views: **Logical View**, **Location View**, **Device View** or **Custom Group**.

TIP: Do not select **Datacenter View** or **Topology View**.

3. In the Tasks pane, select **Wired > Profiles > MACsec**.

The Manage MACsec Profile page is displayed.

4. Select the MACsec profile that you want to assign and then click **Assign**.

The Assign MACsec profile page appears displaying a list of Junos Fusion Enterprise devices managed by Network Director.

Editing the MACsec Profile Assignments

Use the Edit Assignments page to change MACsec profile assignments. To edit an existing assignment:

1. Select a profile from the **Manage MACsec Profile Settings** page and click **Edit Assignment**.

The Edit Assignments page for the selected device appears.

2. Expand the **Devices** cabinet and make the desired change from the **Operation** column of the table.

3. Click **Apply** once you are done with the changes.

The Manage MACsec Profile page is displayed.

RELATED DOCUMENTATION

[Media Access Control Security Overview | 637](#)

[Configuring and Managing MACsec Profiles | 638](#)

[Understanding Junos Fusion Enterprise | 710](#)

Configuring Link Aggregation Groups (LAGs)

IN THIS CHAPTER

- Understanding Link Aggregation | 646
- Managing and Creating a Link Aggregation Group | 647
- Understanding Multichassis Link Aggregation | 653
- Creating and Managing Multichassis Link Aggregation Groups (MC-LAGs) | 654
- Creating and Managing ESI Link Aggregation Groups (ESI-LAGs) | 670

Understanding Link Aggregation

You can combine multiple physical Ethernet ports to form a logical point-to-point link, known as a link aggregation group (LAG) or bundle. A LAG provides more bandwidth than a single Ethernet link can provide. Additionally, link aggregation provides network redundancy by load-balancing traffic across all available links. If one of the links fails, the system automatically load-balances traffic across all remaining links. In a Virtual Chassis, LAGs can be used to load-balance network traffic between member switches.

The maximum number of interfaces that can be grouped into a LAG and the maximum number of LAGs supported on a switch varies according to the switch model and the version of and the version of Juniper Networks Junos operating system (Junos OS) that is running on that switch. [Table 143](#) lists the maximum number of interfaces per LAG and the maximum number of LAGs that are supported on EX Series switches running Junos OS Release 12.3. If your switch is running a different version of Junos OS, refer to the device specific documentation, [EX Series Ethernet Switches](#), before implementing LAG in your network.

Table 143: Maximum Interfaces per LAG and Maximum LAGs per Switch

Switch	Maximum Interfaces per LAG	Maximum LAGs
EX2200	8	32
EX3200	8	32
EX3300 (Standalone and Virtual Chassis)	8	111

Table 143: Maximum Interfaces per LAG and Maximum LAGs per Switch (*continued*)

Switch	Maximum Interfaces per LAG	Maximum LAGs
EX4200 (Standalone and Virtual Chassis)	8	111
EX4500 and EX4550 (Standalone and Virtual Chassis)	8	111
EX6200	8	111
EX8200	12	255
EX8200 Virtual Chassis	12	239

RELATED DOCUMENTATION

[Managing and Creating a Link Aggregation Group | 647](#)

[Network Director Documentation home page](#)

Managing and Creating a Link Aggregation Group

IN THIS SECTION

- [Link Aggregation Group Options | 649](#)
- [Creating a Link Aggregation Group | 650](#)
- [Managing ICCP Settings | 651](#)
- [What To Do Next | 652](#)

IEEE 802.3ad link aggregation enables you to group Ethernet interfaces to form a single link layer interface, also known as a link aggregation group (LAG) or bundle.

Aggregating multiple links between physical interfaces creates a single logical point-to-point trunk link or a LAG. Link Aggregation Control Protocol (LACP), a component of IEEE 802.3ad, provides additional functionality for LAGs.

LACP ensures that both ends of the Ethernet link are functional and are members of the aggregation group before the link is added to the LAG. If you use LACP, make sure that LACP is enabled at both the local and remote ends of the link. When LACP is configured, it detects misconfigurations on the local end or the remote end of the link. Thus, LACP can help to prevent communication failure. When LACP is not enabled, a local LAG might attempt to transmit packets to a remote single interface, which causes the communication to fail. However, when LACP is enabled, a local LAG cannot transmit packets unless a LAG with LACP is also configured on the remote end of the link.

The maximum number of interfaces that can be grouped into a LAG and the maximum number of LAGs supported on a switch varies according to the switch model and the version of Juniper Networks Junos operating system (Junos OS) that is running on that switch. Be aware of the maximum number of interfaces per LAG and the maximum number of LAGs that are supported on your switches by referring to your device specific documentation before implementing LAG in your network.

NOTE: You only see the Manage Lag option under Device Management when a qualified switch is selected in the View Pane.

When creating LAGs, follow these guidelines:

- You must configure the LAG on both sides of the link.
- You must set the interfaces on either side of the link to the same speed.
- You can configure and apply firewall filters on a LAG.

NOTE: You only see the Manage Lag option under Device Management when a qualified switch is selected in the View Pane.

NOTE: MC-LAG, or Multi-Chassis Link Aggregation Group, is a type of LAG with constituent ports that terminate on separate chassis, thereby providing node-level redundancy. Unlike link aggregation in general, MC-LAG is not covered under IEEE 802.1AX-2008. Its implementation varies by vendor. For directions to create an MC-LAG, see [“Creating and Managing Multichassis Link Aggregation Groups \(MC-LAGs\)” on page 654](#).

This topic includes:

Link Aggregation Group Options

From the Manage LAG page, you can:

- Create a new Link Aggregation by clicking **Create**. The Create Link Aggregation window opens—for directions, see [“Creating a Link Aggregation Group” on page 650](#).
- Modify an existing Link Aggregation by selecting it and clicking **Edit**. The Modify Link Aggregation window opens. You can modify all the fields in the Modify Link Aggregation window, except the Interface Name field.
- Delete a Link Aggregation Group by selecting it and clicking **Delete**.
- Manage ICCP settings for the selected device by clicking **Manage ICCP Settings**. See [“Managing ICCP Settings” on page 651](#) for more information.

[Table 144](#) describes the information provided about the link aggregation configurations on the LACP (Link Aggregation Control Protocol) Configuration page. This page lists all link aggregation groups defined on the selected device.

Table 144: LACP (Link Aggregation Control Protocol) Configuration Fields

Field	Description
Logical Interface Name	Name given to the aggregated interface when the LAG was created.
Member Interfaces	Names of individual member interfaces.
LACP Mode	Mode in which LACP packets are exchanged between the interfaces. The possible modes are: <ul style="list-style-type: none"> • Active—Indicates that the interface initiates transmission of LACP packets • Passive—Indicates that the interface responds only to LACP packets.
Description	The description for the LAG. TIP: If you cannot view the entire description, you can resize the Description column by clicking the column border in the heading and dragging it.
Deployment State	The deployment state of the link aggregation. Deployment state can be: <ul style="list-style-type: none"> • Pending Deployment—Indicates that the LAG is not yet deployed on the device. • Deployed—Indicates that the LAG is deployed on the device. • Pending Removal—Indicates that the LAG is deleted.
Creation Time	Date and time when this profile was created.
Update Time	Date and time when this profile was last modified.

Table 144: LACP (Link Aggregation Control Protocol) Configuration Fields (*continued*)

Field	Description
User Name	The username of the user who created or modified the profile.

TIP: All columns might not be displayed. To show or hide fields in the LACP (Link Aggregation Control Protocol) Configuration table, click the DOWN arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

Creating a Link Aggregation Group


You can create one or more LAGs in Network Director. The number of interfaces that can be grouped into a LAG and the total number of LAGs supported on a switch varies according to switch model.

TIP: You can also create one or more MC-LAGs for Virtual Chassis—see [“Creating and Managing Multichassis Link Aggregation Groups \(MC-LAGs\)”](#) on page 654.

To create a link aggregation group:

1. In the View pane, select a switch for link aggregation.

NOTE: The Manage LAG task is only available when a qualified switch is selected in the View pane.

2. Click  in the Network Director banner.

The Manage LAG page opens.
3. Select **Wired > Manage LAG** in the Tasks pane.

The Manage LAG page opens.
4. Click **Create**.

The Create Link Aggregation window opens.
5. Use the up and down arrows to select an AE Name for the aggregation interface. The interface name begins with *ae* followed by an interface number.

6. Select the mode in which LACP packets are to be exchanged between interfaces, either **Active** or **Passive**.
 - **Active**—Indicates that the interface initiates transmission of LACP packets
 - **Passive**—Indicates that the interface responds only to LACP packets.
7. Enter a description for the link aggregation.
8. Configure up to eight available interfaces on the LAG. Select one or more interfaces from the Available list and then click the RIGHT arrow to move them to the Selected list.

NOTE: The Available interfaces list displays only those interfaces that are not part of any link aggregation.

9. If the device is capable of using MC-LAGs, an MC-LAGs section also appears in the Create Link Aggregation window. For information about MC-LAG configuration, see [“MC-LAG Settings” on page 654](#).

10. Click **OK** to save the link aggregation configuration.

A message confirms that the link aggregation is created successfully and ready to be deployed to a device. If the configuration contains an error, the message instead indicates the error.

11. Click **OK** to close the information message.

The LAG appears in the Manage LAG list.

Managing ICCP Settings

When a QFX Series device is the selected scope, you can use the ICCP LAG Settings window to manage ICCP on the selected device. [Table 145](#) describes the fields in this window.

Table 145: ICCP Settings

Field	Description
Disable (Delete) ICCP Settings	Disable ICCP on the device.
AE Name	Select the aggregated Ethernet interface to use for the ICCP connection.
Local IP	Configure the local IP address to be used by all switches hosting the MC-LAG.
Peer IP	Configure the IP address of the ICCP peer.

Table 145: ICCP Settings (continued)

Field	Description
VLAN	Enter the name of the VLAN to use for the ICCP connection.
VLAN ID	Enter the ID of the VLAN to use for the ICCP connection.
Liveness detection min receive interval	Configure the minimum interval at which the switch must receive a reply from the other switch with which it has established a Bidirectional Forwarding Detection (BFD) session.
Liveness detection min transmit interval	Configure the minimum transmit interval during which a switch must receive a reply from a switch with which it has established a BFD session.
Liveness detection backup peer IP	Configure the IP address of the liveness detection backup.
Session establish hold time	Configure the time during which an ICCP connection must succeed between the switches hosting the MC-LAG. Configured session establishment hold time results in faster ICCP connection establishment. The recommended value is 50 seconds.

What To Do Next

The configuration changes that you make in the Build mode are not deployed to devices automatically. After you create a link aggregation group, you must manually deploy the changes to the switches in Deploy mode. For details, see [“Deploying Configuration to Devices” on page 1180](#).

TIP: Even though link aggregation configuration is not contained within a profile, you can view the link aggregation groups assigned to a switch by using the View Assigned Profiles task in Build mode.

RELATED DOCUMENTATION

[Understanding Link Aggregation | 646](#)

[Viewing Profiles Assigned to a Device | 1144](#)

[Creating and Managing Multichassis Link Aggregation Groups \(MC-LAGs\) | 654](#)

[Network Director Documentation home page](#)

Understanding Multichassis Link Aggregation

Layer 2 networks are increasing in scale mainly because of technologies such as virtualization. Protocol and control mechanisms that limit the disastrous effects of a topology loop in the network are necessary. Spanning Tree Protocol (STP) is the primary solution to this problem because it provides a loop-free Layer 2 environment. STP has gone through a number of enhancements and extensions, and although it scales to very large network environments, it still provides only one active path from one device to another, regardless of the number of actual connections existing in the network. Although STP is a robust and scalable solution to redundancy in a Layer 2 network, the single logical link creates two problems: At least half of the available system bandwidth is off-limits to data traffic, and network topology changes occur. The Rapid Spanning Tree Protocol (RSTP) reduces the overhead of the rediscovery process and allows a Layer 2 network to reconverge faster, but the delay is still high.

Link aggregation (IEEE 802.3ad) solves some of these problems by enabling users to use more than one link connection between switches. All physical connections are considered one logical connection. The problem with standard link aggregation is that the connections are point to point.

Multichassis link aggregation groups (MC-LAGs) enable a client device to form a logical LAG interface between two MC-LAG peers (QFX3500 and QFX3600 devices). An MC-LAG provides redundancy and load balancing between the two MC-LAG peers, multihoming support, and a loop-free Layer 2 network without running STP.

On one end of an MC-LAG, there is an MC-LAG client device, such as a server, that has one or more physical links in a link aggregation group (LAG). This client device does not need to have an MC-LAG configured. On the other side of the MC-LAG, there are two MC-LAG peers. Each of the MC-LAG peers has one or more physical links connected to a single client device.

The MC-LAG peers use Interchassis Control Protocol (ICCP) to exchange control information and coordinate with each other to ensure that data traffic is forwarded properly.

Link Aggregation Control Protocol (LACP) is a subcomponent of the IEEE 802.3ad standard. LACP is used to discover multiple links from a client device connected to an MC-LAG peer. LACP must be configured on all member links for an MC-LAG to work correctly.

RELATED DOCUMENTATION

[Creating and Managing Multichassis Link Aggregation Groups \(MC-LAGs\) | 654](#)

[Network Director Documentation home page](#)

Creating and Managing Multichassis Link Aggregation Groups (MC-LAGs)

IN THIS SECTION

- [Accessing the MC-LAG Page | 655](#)
- [\[xref target has no title\]](#)
- [Creating an MC-LAG | 655](#)
- [MC-LAG Automation Parameters | 661](#)
- [Editing an MC-LAG | 663](#)
- [Deleting an MC-LAG | 669](#)
- [Managing an MC-LAG Created Through CLI Mode | 669](#)

Multichassis link aggregation groups (MC-LAGs) enable a device to form a logical link aggregation group (LAG) interface between two switches. An MC-LAG provides redundancy and load balancing between the two switches, multihoming support, and a loop-free Layer 2 network without running Spanning Tree Protocol (STP).

MC-LAG peer switches use the Inter-Chassis Control Protocol (ICCP) to exchange control information and interchassis link (ICL) to exchange data.

At one end of an MC-LAG are the MC-LAG client devices, such as servers or switches, that have one or more physical links in a LAG. Client devices do not need to detect the MC-LAG. At the other end of the MC-LAG are two peer devices. Each of these switches has one or more physical links connected to a single client device. The switches coordinate with each other to ensure that data traffic is forwarded properly.

You can create MC-LAGs using QFX Series and EX9200 devices. ELS and non-ELS devices are supported. However, both the peer devices must be the same type and must be either ELS or non-ELS devices. Network Director can manage MC-LAG devices that are created and configured through the CLI mode also. If MC-LAG devices are configured through the CLI mode, ensure that LLDP is enabled on MC-LAG, ICCP LAG, ICL LAG, and client LAG links.

Supported devices in an MC-LAG:

- Peer devices: QFX3500, QFX3600, QFX5100, QFX10002, and EX9200 switches
- Client devices: All standalone and Virtual Chassis devices managed by Network Director except Virtual Chassis Fabric (VCF) members, QFabric and IP Fabric devices, and MX Series devices

For detailed steps on creating MC-LAGs using Network Director, follow the procedure given below or the steps shown in this video-based tutorial:



Video: [Configuring MC-LAG using Network Director](#)

This topic includes:

Accessing the MC-LAG Page

To access the MC-LAG page:

1. Click the Build mode icon  in the Network Director banner.
2. Select **Wired** > **Tasks** > **Manage MC-LAG** in the Tasks pane.

The Manage MC-LAG page opens, which displays the existing MC-LAG peers and enables you to create, edit, or delete an MC-LAG. In the Manage MC-LAG page, the peer devices for each MC-LAG that is created using Network Director or created using the CLI mode and discovered for management by Network Director, are listed. The Manage MC-LAG page displays the device name, device model, deployment status, and local IP address of the MC-LAG peer devices. If any peer device is not managed by Network Director, the MC-LAG Peer displays as *Unknown*. Click the peer devices of any MC-LAG to view details of the MC-LAG, such as, descriptions of the peer devices, peer-to-peer link details, and client-to-peer link details.

Creating an MC-LAG

IN THIS SECTION

- [Selecting Peer Devices and Configuring Peer-to-Peer Link Settings | 656](#)
- [Selecting Client Devices and Configuring Client-to-Peer Link Settings | 658](#)
- [Saving MC-LAG Settings | 660](#)
- [Deploying MC-LAG Configuration | 661](#)

To create an MC-LAG:

1. Click **Create MC-LAG** in the Manage MC-LAG page.

The Create MC-LAG page opens. It displays two tabs—Peer Devices and Client Devices. By default, the Peer Devices tab is selected and displays in orange color.

On the left of the Create MC-LAG page, the Peer Devices tab lists QFX Series and EX9200 devices that are managed by Network Director. These are the available devices from which you can select the peer devices for the MC-LAG you create. On the right, a schematic diagram of the two peer devices PEER1 PEER2 and a representation of the client devices as boxes are displayed.

Creating an MC-LAG involves four tasks:

Selecting Peer Devices and Configuring Peer-to-Peer Link Settings

To select the peer devices and configure peer-to-peer link settings:

1. From the list of devices in the Peer Devices tab in the Create MC-LAG page, select a device, and drag and drop it into one of the boxes labeled PEER1 or PEER2 2.

After you drag and drop the first peer device, the list refilters and displays only devices that qualify to be the second peer.

For example, if you select a QFX10002 switch as one of the peer devices, then only QFX10002 switches are listed for you to select as the second peer device.

2. Select the second device from the refiltered list of peer devices and drag and drop it into the second peer box.

The Peer to Peer Link Settings window opens. The Client Devices tab is automatically enabled in the background in the Create MC-LAG page.

3. In the Peer to Peer Link Settings window, select **Combine Data and Control Links** if you want to combine the data and control links, that is, if you want a single link to act as both the control link and data link between the two ports that you selected. Network Director configures this link as an ICCP link.

If you want to have ICCP (control) and ICL (data) links separately between the peer devices, do not select this option.

NOTE: By default, Combine Data and Control Links is not selected. If you select this, you must specify the **VLAN Name** and **VLAN ID** in the respective fields in the Peer to Peer Link Settings window.

The physically connected ports of the peer devices are displayed in the Data and Control Link Ports* table in the Peer to Peer Link Settings window, if you have refreshed the topologies of the peer devices

in the Topology View in Network Director. If the LLDP or topology information of the peer devices are not available for Network Director, port details are not displayed.

TIP: To refresh the topology, select Topology View in Views and then select Discovery-Topology > Refresh Topology in the Tasks pane. For the topology to refresh, LLDP must be enabled on the interfaces that are connected to the peer and client devices.

4. Click **Add Port**.

A new row is added to the table under Data and Control Link Ports*, where you must enter the port details for the peer devices.

5. From the drop-down menu for the PEER₂ device you selected, select a port to assign to the MC-LAG.

6. From the drop-down menu for the PEER₁ device you selected, select a port to assign to the MC-LAG.

7. Specify the type of link between the ports on the two peer devices by selecting **Data**, **Control**, or **Data & Control** from the Port Type list.

NOTE: If you have selected Combine Data and Control Links in Step 3, Data & Control is the default port type. If you have not selected Combine Data and Control Links in Step 3, the available options are Data and Control.

For the Data & Control port type, a single link between the peer devices acts as both the control and data link. If you select Data as the port type, then you must add a new pair of ports in the next row by clicking Add Ports, and then selecting Control as the port type for the control between the peer devices. If you select Control as the port type, then you must add a new pair of ports in the next row by clicking Add Ports, and then selecting Data as the port type for the control between the peer devices.

NOTE: You must specify at least one link between the peer devices.

8. Click **Update**.

A new row is added with the port details.

9. Enter the **IPv4 Address** and mask.

This IPv4 address is configured for the control link inet address and used as the local IP address for the ICCP. Network Director configures the peer IP addresses from this local IP address internally.

10. Click **OK**.

The Peer to Peer Link Settings window closes, and the Create MC-LAG page is displayed.

The Client Devices tab is selected by default. In the schematic diagram, the links that you configured between the peer devices changes to display in green, indicating that the links are successfully configured. The color does not indicate the operational status of the link.

Selecting Client Devices and Configuring Client-to-Peer Link Settings

To select a client device and configure client-to-peer link settings:

1. In the Client Devices tab on the Create MC-LAG page, select the device type for the client you want to be part of the MC-LAG by clicking an option from the drop-down menu in the **Type** field. Options available are: Switches, Bare Metal Servers, and Hypervisors.

The list of client devices displays the devices (switches, hypervisors, or bare metal servers) depending on the type that you selected. By default, only switches are listed.

2. Select a device from the list of client devices, and drag and drop it into one of the boxes labeled *Drag & Drop Clients here to add*.

NOTE: If you select a Virtual Chassis switch, the client box shows a graphical representation of Virtual Chassis; if you select a bare metal server or hypervisor server, the client box shows the graphical representation of that respective type of server.

The Client to Peer Link Settings window opens.

3. Select **MC-AE Mode**. The modes available are Active-Active and Active-Standby.

NOTE: Only EX9200 and QFX10002 devices support both Active-Active and Active-Standby modes. The other devices support only the Active-Active mode.

- **Active-Active mode:** If the client-to-peer setting mode is set to Active-Active mode, all peer port links will be active in the MC-LAG. In this mode, MAC addresses discovered in one MC-LAG peer device is propagated to the other peer device. Traffic is load balanced, and convergence is faster.
- **Active-Standby mode:** If the client-to-peer setting mode is set to Active-Standby mode, only one of the MC-LAG peer devices is active at any given time. The other peer device is in backup, that is standby, mode.

The ports that are physically connected between the client and peer devices are displayed in PEER? 1 and PEER? 2 if you have refreshed the topologies of the peer devices in the Topology View in Network Director. If the LLDP or topology information of the peer devices are not available for Network Director, the port details are not displayed.

TIP: To refresh the topology, select Topology View in Views and then select Discovery-Topology > Refresh Topology in the Tasks pane. For the topology to refresh, LLDP must be enabled on the ports that are connected to the peer and client devices.

4. Click **Add Port** to select the client and peer ports.

A new row is added to the table, where you must enter the port details for the peer and client devices.

5. Select the client port from the drop-down menu corresponding to the **Client Port**.

NOTE: If you selected Switches as the type of client device, then Client Port is a mandatory field. If you selected Bare Metal Servers or Hypervisors, then the drop-down menu does not display any client port, as Network Director does not enable you to configure VLANs or ports in the servers.

6. From the drop-down menu select the peer port you want to connect to the client port.

7. Click **Update**. A row is added that displays the client port and peer port.

NOTE: If you have selected Peer? 1 Port and linked it to a client port first, then select Peer? 2 Port and link it to a client port. Both Peer? 1 Port and Peer? 2 Port cannot be selected in one row. The client device must be connected to both peer devices.

8. In the Client to Peer VLANs* table in the Client to Peer Link Settings window, Network Director displays all the VLANs of the client. If the client has the same VLAN ID as that of a peer or the peers, Network Director automatically populates the Routed Interface Address and VRRP Attributes for those peers in the respective fields. If there are no VLANs displayed in the table, add a VLAN by clicking Select VLAN or Add VLAN. This VLAN is configured in the PEER? 1 and PEER? 2 devices to ensure connectivity and data flow between the peers. You can configure multiple clients.

To edit a VLAN, click the fields of the VLAN that you want to edit.

NOTE:

- If you want to select a VLAN other than the VLANs displayed in the Client to Peer Link Settings window, click **Select VLAN** and select the VLANs from the list that displays in the Choose VLAN Profile pop-up window.
- You can remove a VLAN that you have created, but not deployed, in the client device by selecting the VLAN and clicking **Remove VLAN**.
- Do not remove VLANs that are deployed in the devices.

Network Director Release 2.5 supports Layer 3 routing. To enable Layer 3 routing, configure the Routed Interface Address and VRRP Attributes by clicking the respective fields.

Select the IP type by clicking the arrow in the **IP Type** field. The available options are IPv4 and IPv6.

9. Enter the IP addresses and mask for the peer devices in the corresponding fields. The IP addresses must be the IP addresses of the integrated routing and bridging interface.
10. Enter the VRRP group ID in **Group ID** and enter the virtual IP address in **Virtual IP** to assign the virtual IP that is shared between each switch in the VRRP group.
11. Click **Update**.

To add a VLAN, click **Add VLAN**. A new row is created in the Client to Peer VLANs* table. Enter the VLAN ID and VLAN name in their corresponding fields, and perform Steps 8 through 11.

To remove a VLAN, select the VLAN, and click **Remove VLAN**.

12. Click **OK** to submit the settings that you entered in the Client to Peer Link Settings window and to close the window.

The Client to Peer Link Settings window closes.

Network Director configures different IP addresses on IRB interfaces on the MC-LAG peers and runs the VRRP on the IRB interfaces. The virtual IP address is the gateway IP address for the MC-LAG clients. To provide Layer 3 routing functions to downstream clients, the MC-LAG network peers must be configured to provide the same gateway address to the downstream clients.

Saving MC-LAG Settings

To save the MC-LAG settings that you configured:

1. Click **Save** in the Create MC-LAG page.

Network Director saves the MC-LAG settings and displays the message **MC-LAG save is successful and is ready to be deployed to the devices**.

2. Click **OK**.

The Manage MC-LAG page lists the MC-LAG that you created. By default, the Deployment State for the MC-LAG displays as Pending Deployment.

Deploying MC-LAG Configuration

To deploy a new or edited MC-LAG configuration:

1. In the **Deploy** mode, click **Configuration Deployment > Deploy Configuration Changes** in the Tasks pane.

The Devices with Pending Changes page opens, displaying devices that have pending configuration changes.

2. In the list in the Devices with Pending Changes page, select the devices that you configured as the peer and client devices of the MC-LAG.

NOTE: To view the deployment information for a device, select the device and click **View**. The Configuration window opens, which shows the CLI and XML view of the configuration that will be deployed in the device.

3. Click **Deploy Now** to deploy the configuration.

The Device Configuration window opens. The Deployment Status shows the status as **INPROGRESS** and changes to **SUCCESS** once the deployment is successfully completed.

MC-LAG Automation Parameters

Network Director configures a number of parameters internally and automates the creation or modification of MC-LAGs.

[Table 146](#) describes the parameters that are internally configured by Network Director.

Table 146: MC-LAG Automation Parameters

Parameter	Description
LAG	LAG is created for ICCP, ICL, and MC-AE in peer devices, and a LAG is created for client devices.
mc-ae-id	Specifies which MC-LAG the aggregated Ethernet interface belongs to.

Table 146: MC-LAG Automation Parameters (continued)

Parameter	Description
redundancy-group (supported only in QFX10002 and EX9200 devices)	Used by ICCP to associate multiple chassis that perform similar redundancy functions. It is used to establish a communication channel so that applications running on the peer devices can exchange messages.
init-delay-time:240ms:	Specifies the delay in number of seconds to bring the MC-LAG interface back to the Up state when an MC-LAG peer is rebooted.
chassis-id 0 for Peer 1, and 1 for Peer 2	Used by LACP for calculating the port number of the MC-LAG physical member links. Each MC-LAG peer must have a unique chassis ID.
status-control Active for Peer 1, and Standby for Peer 2	Specifies whether this node becomes active or goes into standby mode when an ICL failure occurs; must be active on one node and standby on the other node.
LACP active	Configured in ICL LAG, ICCP LAG, MC-LAG and client switch LAG. LACP is used to discover multiple links from a client device connected to an MC-LAG peer. LACP must be configured on all member links for an MC-LAG to work correctly.
LACP system-id and admin-key	Configures the same LACP system ID and admin-key for the MC-LAG on each MC-LAG peer. This displays Peer1 and Peer2 as a single switch to the edge switch when negotiating LACP.
LACP periodic fast	Configured on ICCP LAG, ICL LAG and MC-LAG. LACP fast periodic is achieved by configuring fast intervals (in seconds) for periodic transmission of LACP.
Hold time Up 100000 down 0 for interfaces used for MC LAG. Up 0 down 2000 for interfaces used for ICL LAG.	Specifies the hold-time value to use to damp interface transitions. When an interface goes down, it is not broadcast to the rest of the system till it remains down for the hold-time period. Similarly, an interface is not broadcast as being Up till it remains up for the hold-time period.
multi-chassis-protection	Specifies the peer's ICCP IP address and the ICL link used for protection if the MC-AE interface goes down.
session-establishment-hold-time 300	Establishes ICCP connection quickly.
backup-liveness-detection: management IP of peer device	Is invoked when the ICCP link goes down. With backup liveness detection enabled, the MC-LAG peers establish an out-of-band channel through the management network in addition to the ICCP channel.

Table 146: MC-LAG Automation Parameters (*continued*)

Parameter	Description
liveness-detection minimum-receive-interval 500, multiplier 3, transmit-interval 500	Determines whether a peer is up or down by exchanging keepalive messages over the management link between the two ICCP peers.
RSTP	Is enabled on peer devices MC-LAG and switch client LAG in point to point mode . If client is a server, then it enables bpdu-block-on-edge and edge on MC-LAG peer devices. Bridge-priority is set to 0 on both the peer devices.
ARP, MAC, arp-l2-validate, l2-interface ICL LAG on IRB	Provides IRB-to-IRB connectivity across the ICL. Using the VRRP over IRB method to enable Layer 3 functionality, it configures static ARP entries through the ICL for the IRB interface of the remote MC-LAG peer, which enables routing protocols to run over the IRB interfaces.

Editing an MC-LAG

IN THIS SECTION

- [Managing Peer Devices and Peer-to-Peer Link Settings | 664](#)
- [Managing Client Devices and Client-to-Peer Link Settings | 665](#)

In the Manage MC-LAG page, you can add, edit or delete peer ports, edit existing peer-to-peer link settings, add client, remove client, and edit client-to-peer link settings. You cannot add or delete peer devices if both the peers are part of MC-LAG.

1. Click **Edit** corresponding to the MC-LAG peers that you want to modify, in the Manage MC-LAG page.

The Edit MC-LAG page opens. It displays two tabs—Peer Devices and Client Devices. If both the peer devices of the MC-LAG are already configured as part of the MC-LAG, the Client Devices tab is selected, and it displays in orange color. On the left of the Edit MC-LAG page, a list of client devices are displayed.

If one of the peer devices is *Unknown*, the Peer Devices tab is selected, and it displays in orange color. On the left of the Edit MC-LAG page, a list of peer devices, that are of the same type and ELS capability as of the discovered peer, are displayed.

On the right of the Edit MC-LAG page, a schematic diagram of the existing two peer devices PEER¹ and PEER² and a representation of the client devices as boxes are displayed.

Managing Peer Devices and Peer-to-Peer Link Settings

To add, edit, or delete a peer port, or edit peer-to-peer link settings:

1. Click **Control Link** or **Data Link** that is displayed between PEER¹ and PEER² in the schematic diagram.

The Peer to Peer Link Settings window opens.

NOTE: The Combine Data and Control Links option is unavailable.

The peer ports that you already configured are displayed in the table Data and Control Links Ports*.

2. To add a port, click **Add Port**.

A new row is added to the table, where you must enter the port details for the peer devices.

3. From the drop-down menu for the PEER¹ device, select a port to assign to the MC-LAG.
4. From the drop-down menu for the PEER² device, select a port to assign to the MC-LAG.
5. Specify the type of link between the ports on the two peer devices by selecting **Data**, **Control**, or **Data & Control** from the Port Type list.

NOTE: If you have selected Combine Data and Control Links, Data & Control is the default port type. If you have not selected Combine Data and Control Links, the available options are Data and Control.

For the Data & Control port type, a single link between the peer devices act as both the control and data link. If you select Data as the port type, then you must add a new pair of ports in the next row by clicking Add Ports, and then selecting Control as the port type for the control between the peer devices. If you select Control as the port type, then you must add a new pair of ports in the next row by clicking Add Ports, and then selecting Data as the port type for the control between the peer devices.

NOTE: You must specify at least one link between the peer devices.

NOTE: To delete a peer port, select the port that you want to remove from the MC-LAG, and click **Remove Port**.

To edit a peer port, click the port and modify the port details.

6. Click **Update**.

7. If you want to edit the control link IPv4 address, edit the **IPv4 Address** and mask fields in the Control Link table.

This IPv4 address is configured for the control link inet address and is used as the local IP address for the ICCP. Network Director internally configures the peer IP addresses from this local IP address.

If the data and control links are combined, VLAN ID and VLAN name are displayed and can be edited here.

8. Click **OK**.

The Peer to Peer Link Settings window closes, and the Edit MC-LAG page is displayed.

Managing Client Devices and Client-to-Peer Link Settings

To add or remove client devices, and edit client-to-peer link settings:

1. Click Client Devices tab in the Edit MC-LAG page.

The Client Devices tab on the Create MC-LAG page lists switches that are managed by Network Director. On the right, a schematic diagram of the two peer devices PEER1 and PEER2 and a representation of the client devices as boxes are displayed.

2. Select the device type for the client you want to be part of the MC-LAG by clicking an option from the drop-down menu in the **Type** field. Options available are: Switches, Bare Metal Servers, and Hypervisors.

The list displays the devices (switches, hypervisors, or bare metal servers) depending on the type that you selected. By default, only switches are listed.

3. Select a device from the list of client devices, and drag and drop it into one of the boxes labeled as *Drag & Drop Clients here to add*.

NOTE: If you select a Virtual Chassis switch, the client box shows a graphical representation of Virtual Chassis; if you select a bare metal server or hypervisor server, the client box shows the graphical representation of that respective type of server.

To delete a client device from an MC-LAG configuration, click the **x** mark on the client device in the carousel. The client device is removed from the carousel and the Deployment State changes to Pending Deployment in the Manage MC-LAG page.

The Client to Peer Link Settings window opens.

4. Select the **MC-AE Mode**. The available options are: Active-Active and Active-Standby.

NOTE:

- The MC-AE mode is enabled if you are adding a client device.
 - Only EX9200 and QFX10002 devices support both Active-Active and Active-Standby modes. The other devices support only the Active-Active mode.
- **Active-Active mode:** If the client-to-peer setting mode is set to Active-Active mode, all peer port links will be active in the MC-LAG. In this mode, MAC addresses discovered in one MC-LAG peer device is propagated to the other peer device. Traffic is load balanced, and convergence is faster.
 - **Active-Standby mode:** If the client-to-peer setting mode is set to Active-Standby mode, only one of the MC-LAG peer devices is active at any given time. The other peer device is in backup, that is standby, mode.

The ports that are physically connected between the client and peer devices are displayed in PEER¹ and PEER² if you have refreshed the topologies of the peer devices in the Topology View in Network Director. If the LLDP or topology information of the peer devices are not available for Network Director, the port details are not displayed.

TIP: To refresh the topology, select Topology View in Views and then select Discovery-Topology > Refresh Topology in the Tasks pane. For the topology to refresh, LLDP must be enabled on the ports that are connected to the peer and client devices.

5. Click **Add Port** to select the client and peer ports.

A new row is added to the table, where you must enter the port details for the peer and client devices.

6. Select the client port from the drop-down menu corresponding to the **Client Port**.

NOTE: If you selected Switches as the type of client device, then Client Port is a mandatory field. If you selected Bare Metal Servers or Hypervisors, then the drop-down menu does not display any client port, as Network Director does not enable you to configure VLANs or ports in the servers.

7. Select **Peer 1 Port** or **Peer 2 Port** from the drop-down list in the corresponding fields.

NOTE: The client port is displayed only if you have selected Switches as the device type for the client device.

8. Click **Update**.

9. To add a new peer port and link it to a client port:

Click **Add Port**.

A blank row is added in the Client to Peer Ports table.

10. Select the client port by clicking the drop-down menu corresponding to the **Client Port**.

11. Select **Peer 1 Port** or **Peer 2 Port**, from the drop-down menu in the corresponding fields.

NOTE: If you have selected Peer 1 Port and linked it to a client port earlier, then select Peer 2 Port and link it to the a client port. Both Peer 1 Port and Peer 2 Port cannot be selected in one row.

12. Click **Update**.

13. In the Client to Peer VLANs* table, in the Client to Peer Link Settings window, Network Director displays all the VLANs of the client. If the client has the same VLAN ID as that of a peer or the peers, Network Director automatically populates the Routed Interface Address and VRRP Attributes for those peers in the respective fields. If there are no VLANs displayed in the table, add a VLAN by clicking Select VLAN or Add VLAN. This VLAN is configured in the Peer 1 and Peer 2 devices to ensure connectivity and data flow between the peers. You can configure multiple clients.

NOTE:

- If you want to select a VLAN other than the VLANs displayed in the Client to Peer Link Settings window, click **Select VLAN** and select the VLANs from the list that displays in the Choose VLAN Profile pop-up window.
- You can remove a VLAN that you have created, but not deployed, in the client device by selecting the VLAN and clicking **Remove VLAN**.
- Do not remove VLANs that are deployed in the devices.

Network Director Release 2.5 supports Layer 3 routing. To enable Layer 3 routing, configure the Routed Interface Address and VRRP Attributes in the respective fields.

Select the IP type by clicking the arrow in the **IP Type** field. The available options are IPv4 and IPv6.

Enter the IP addresses and mask for the peer devices in the corresponding fields.

NOTE: While editing an existing client device link settings, you cannot edit the VLAN Name, VLAN ID, Routed interface Address, and VRRP Attributes if they are already configured. If they are not configured, you can add Routed interface Address and VRRP Attributes.

14. Enter the VRRP group ID in **Group ID** and enter the virtual IP address in **Virtual IP** to assign the virtual IP address that is shared between each switch in the VRRP group.

15. Click **Update**.

To add a VLAN, click **Add VLAN**. A new row is created. Enter the VLAN ID and VLAN name in their corresponding fields, and perform Steps 13 through 15.

To remove a VLAN, select the VLAN and click **Remove VLAN**.

16. Click **OK** to submit the settings that you entered in the Client to Peer Link Settings window.

The Client to Peer Link Settings window closes.

17. Click **Save** in the Manage MC-LAG page.

Network Director saves the MC-LAG settings and displays the message **MC-LAG save is successful and is ready to be deployed to the devices**.


18. Click **OK**.

The Manage MC-LAG page lists the newly created MC-LAG. By default, the Deployment State for the newly created MC-LAG displays as Pending Deployment.

To deploy the edited MC-LAG, see [“Deploying MC-LAG Configuration” on page 661](#)

Deleting an MC-LAG

To delete MC-LAG:

1. Click the Build mode icon  in the Network Director banner.
2. Select **Wired > Tasks > Manage MC-LAG** in the Tasks pane.
The Manage MC-LAG page opens, displaying the existing MC-LAG peers and enables you to delete MC-LAGs. The MC-LAGs displayed can be MC-LAGs that are created using Network Director or through the CLI mode.
3. Click **Delete** for the corresponding MC-LAG Peers that you want to delete, in the Manage MC-LAG page.

NOTE: If you delete an MC-LAG, Network Director removes the MC-LAG configuration settings from the peer devices and also deletes the LAG configuration from the client devices. The Deployment State changes to Pending Removal if the MC-LAG is already deployed. If it is not deployed, that is, if it is Pending Deployment, then the MC-LAG is removed from the Manage MC-LAG page.

Managing an MC-LAG Created Through CLI Mode

IN THIS SECTION

- [MC-LAG Peer Pairing | 669](#)
- [Mapping Client Devices to Peer Devices | 670](#)
- [Ports Mapping Between Peer-to-Peer and Client-to-Peer Devices | 670](#)

MC-LAG Peer Pairing

Once the MC-LAG devices are discovered by Network Director and Network Director successfully retrieves the MC-LAG configuration from the peer devices, Network Director pairs the MC-LAG peers based on the ICCP local IP address and peer IP address. For example, if Peer 1 is configured with ICCP local IP address 192.0.2.1 and Peer IP address 192.0.2.2, and Peer 2 is configured with ICCP local IP address 192.0.2.2

and Peer IP address 192.0.2.1, then based on the local IP address of Peer 1, Network Director searches for devices that have the same peer IP address as the local IP address. Because Peer 2 has the same peer IP address as the Peer 1 IP address, these two devices form MC-LAG peers. If in case, the local IP address is not found, then Network Director displays one of the peer devices in the MC-LAG pair as *Unknown* in the MC-LAG Manage page.

Mapping Client Devices to Peer Devices

If LLDP is enabled in the connected ports of the peer devices and client devices, after refreshing the topology, the Edit MC-LAG page displays the Network Director managed client switches connected to peer devices. If the client device is not managed by Network Director, or if the client device is not a switch (it is a bare metal server or a hypervisor), or if the topology information is not available for the devices, then the Edit MC-LAG page displays the client device as **Client_MC-AE ID (Unknown)**, where MC-AE ID specifies which MC-LAG the aggregated Ethernet port belongs to.

Ports Mapping Between Peer-to-Peer and Client-to-Peer Devices

On refreshing the topology, peer-to-peer link settings display the port mapping between the peer devices, and client-to-peer link settings display the port mapping between the client and peer devices.

RELATED DOCUMENTATION

[Understanding Link Aggregation | 646](#)

[Managing and Creating a Link Aggregation Group | 647](#)

[Viewing Profiles Assigned to a Device | 1144](#)

[Network Director Documentation home page](#)

Creating and Managing ESI Link Aggregation Groups (ESI-LAGs)

IN THIS SECTION

- [Accessing the ESI-LAG Page | 671](#)
- [Creating an ESI-LAG | 671](#)
- [Editing an ESI-LAG | 676](#)
- [Deleting an ESI-LAG | 679](#)
- [ESI-LAG Automation Parameters | 680](#)

Ethernet Switch Identifier (ESI) refers to the set of Ethernet links that connect one or more access devices (called client devices) to a pair of core devices (called as peers) in a campus environment. ESI link aggregation groups (ESI-LAGs) enable one or more client devices to form a logical link aggregation group (LAG) interface with the peers. The peer should already be connected with each other before forming an ESI-LAG between them.

You can create ESI-LAGs by using EX9200 devices as the core devices.

NOTE: Network Director supports an ESI-LAG configuration only if the ESI-LAG is created by using Network Director.

Supported devices in an ESI-LAG:

- Peer devices in a core network: EX9200
- Client devices in an access network: EX4300, EX4200, EX3400, EX2300

For creating an ESI-LAG, follow the procedure described in this topic:

This topic includes:

Accessing the ESI-LAG Page

To access the ESI-LAG page:

1. Click the Build mode icon  in the Network Director banner.
2. Select **Wired > Tasks > Manage ESI-LAG** in the Tasks pane.

The Manage ESI-LAG page opens, which displays the existing ESI-LAG peers and enables you to create, edit, or delete an ESI-LAG. The Manage ESI-LAG page also displays the device name, device model, deployment status, and local IP address of the ESI-LAG peer devices. Click the peer devices of any ESI-LAG to view details such as, descriptions of the peer devices, peer-to-peer link details, and client-to-peer link details, of the ESI-LAG.

Creating an ESI-LAG

IN THIS SECTION

- [Selecting Peer Devices and Configuring Peer-to-Peer Link Settings | 672](#)
- [Selecting Client Devices and Configuring Client-to-Peer Link Settings | 674](#)

- Saving ESI-LAG Settings | 675
- Deploying ESI-LAG Configuration | 676

To create an ESI-LAG:

1. Click **Create ESI-LAG** on the Manage ESI-LAG page.

The Create ESI-LAG page opens. It displays two tabs—Peer Devices and Client Devices. By default, the Peer Devices tab is selected and displays in orange color.

On the left of the Create ESI-LAG page, the Peer Devices tab lists EX9200 devices that are managed by Network Director. These are the available devices from which you can select the peer devices for the ESI-LAG you create. On the right, a schematic diagram of the two peer devices PEER1 and PEER2 and boxes representing the client devices is displayed.

Creating an ESI-LAG involves the following tasks:

Selecting Peer Devices and Configuring Peer-to-Peer Link Settings

To select the peer devices and configure peer-to-peer link settings:

1. From the list of devices in the Peer Devices tab on the Create ESI-LAG page, select a device, and drag and drop it into the box labeled PEER1 or PEER2.

After you drag and drop the first peer device, the list of devices refilters and displays only devices that qualify to be the second peer.

For example, if you select an EX9200 device as one of the peer devices, then other EX9200 devices that are discovered by Network Director are listed for you to select as the second peer device.

2. Select the second device from the refiltered list of peer devices and drag and drop it into the second peer box.

The Peer to Peer Link Settings window opens. The Client Devices tab is automatically enabled in the background on the Create ESI-LAG page.

3. In the Peer to Peer Link Settings window, click **Add Port** to add ports of the peer devices to be used in the LAG.

A new row is added to the Peer to Peer Ports* table, where you must enter the port details for the peer devices.

4. From the drop-down menu for the Peer 1 device, select a port to assign to the ESI-LAG.

5. From the drop-down menu for the Peer 2 device, select a port to assign to the ESI-LAG.
6. In the Loop Back section, configure the loopback IPv4 address for Peer 1 and Peer 2 in the **Peer 1 IPv4 Address** and **Peer 2 IPv4 Address** text fields, respectively.

The loopback address ensures that the peer devices are reachable to management applications and other entities that want to communicate with the devices.

7. In the Logical Interface section,
 - Provide the logical IPv4 address for Peer 1 and Peer 2 devices in the **Peer 1 IPv4 Address** and **Peer 2 IPv4 Address** text fields, respectively.

The logical interfaces are created on the interfaces on which the ESI-LAG is to be configured.

- (Optional) In the **BGP Group Name** text field, edit the name for the BGP group to which the peer devices will be assigned.

When the ESI-LAG is deployed on your network, a BGP group with the name assigned is created. Other BGP-related parameters are autogenerated and assigned to the peer devices.

- (Optional) In the **Virtual Switch Instance** text field, edit the name of the virtual switch instance assigned to the peer devices.

When the ESI-LAG is deployed on your network, a virtual switch instance with the name assigned is created in the routing instance of the peer devices. You can configure only one routing instance by using Network Director.

- (Optional) In the **VRF Instance Name** text field, edit the name of the VRF instance assigned to the peer devices.

When the ESI-LAG is deployed on your network, a VRF instance by the name assigned is created and assigned to the peer devices.

- In the **Autonomous System Number** text field, enter the autonomous system number to which the peer devices are assigned as part of ESI-LAG.

8. In the Peer To Peer VLAN table:

- Click **Add VLAN** to add VLANs. This VLAN is configured between the peers to ensure connectivity and data flow between the peers.

When you click Add VLAN, enter the following values for the VLAN:

- **VLAN ID**
- **VLAN Name**
- **VNI ID**

- **IP Type, IP Address** of the peers and **Mask** for the routed interface address of the VLAN
- **Anycast Gateway** for the VLAN

9. Click **Update**.

A row is added that displays the VLAN associated with the peer devices.

10. (Optional) Click a VLAN and click **Remove VLAN** to remove any VLAN that you do not want to be part of the ESI-LAG configuration.

11. Click **OK**.

The Peer to Peer Link Settings window closes, and the Create ESI-LAG page appears.

The Client Devices tab is selected by default. In the schematic diagram, the links that you configured between the peer devices changes to green, indicating that the links are successfully configured. The color does not indicate the operational status of the link.

Selecting Client Devices and Configuring Client-to-Peer Link Settings

To select a client device and configure client-to-peer link settings:

1. In the Client Devices tab on the Create ESI-LAG page, select the client device.
2. Select a device from the list of client devices, and drag and drop it into one of the boxes labeled *Drag & Drop Clients here to add*.

The Client to Peer Link Settings window opens.

3. Click **Add Port** to select the client and peer ports.

A new row is added to the table, where you must enter the port details for the peer and client devices.

4. Select the client port from the **Client Port** drop-down list.

5. From the drop-down menu, select the peer port you want to connect to the client port.

NOTE: You can configure only one interface to connect the client to a peer.

6. Click **Update**. A row is added that displays the client port and peer port.

NOTE: If you have selected Peer 1 Port and linked it to a client port first, then select Peer 2 Port and link it to a client port. Both Peer 1 Port and Peer 2 Port cannot be selected in one row. The client device must be connected to both peer devices.

7. In the Client to Peer VLANs* table in the Client to Peer Link Settings window, click **Select VLAN** to assign the client to one or more VLANs.

The Choose VLANs pop-up window appears listing the VLANs to which the Peers are assigned.

8. Select one or more VLANs from the Choose VLANs pop-up window to which you want to assign the client.

9. Click **OK**.

The selected VLANs are added to the Client to Peer VLANs* table.

10. (Optional) Remove a VLAN that you have created, in the client device by selecting the VLAN and clicking **Remove VLAN**.

11. Click **Update**.

The client is assigned to the selected VLANs.

12. Click **OK** to submit the settings that you entered in the Client to Peer Link Settings window and close the window.

The Client to Peer Link Settings window closes.

Saving ESI-LAG Settings

To save the ESI-LAG settings that you configured:

1. Click **Save** on the Create ESI-LAG page.

Network Director saves the ESI-LAG settings and displays the message **ESI-LAG save is successful and is ready to be deployed to the devices**.

2. Click **OK**.

The Manage ESI-LAG page lists the ESI-LAG that you created. By default, the Deployment State for the ESI-LAG displays as Pending Deployment.

Deploying ESI-LAG Configuration

To deploy a new or edited ESI-LAG configuration:

1. In the **Deploy** mode, click **Configuration Deployment > Deploy Configuration Changes** in the Tasks pane.

The Devices with Pending Changes page opens, displaying devices that have pending configuration changes.

2. In the list on the Devices with Pending Changes page, select the devices that you configured as the peer and client devices of the ESI-LAG.

NOTE: To view the deployment information for a device, select the device and click **View**. The Configuration window opens, which shows the CLI and XML view of the configuration that will be deployed on the device.

3. Click **Deploy Now** to deploy the configuration.

The Device Configuration window opens. The Deployment Status shows the status as **INPROGRESS** and changes to **SUCCESS** once the deployment is successfully completed.

Editing an ESI-LAG

IN THIS SECTION

- [Managing Peer Devices and Peer-to-Peer Link Settings | 677](#)
- [Managing Client Devices and Client-to-Peer Link Settings | 678](#)

On the Manage ESI-LAG page, you can add, edit, or delete peer ports, edit existing peer-to-peer link settings, add client, remove client, and edit client-to-peer link settings. However, you cannot add or delete peer devices of the ESI-LAG.

NOTE: You cannot edit an ESI-LAG after it is configured and the devices are deployed on the network (that is, the state of the devices is DEPLOYED.)

1. On the Manage ESI-LAG page, click **Edit** corresponding to the ESI-LAG peers that you want to modify.

The Edit ESI-LAG page opens. It displays two tabs—Peer Devices and Client Devices. If both the peer devices of the ESI-LAG are already configured as part of the ESI-LAG configuration, the Client Devices tab is selected, and it displays in orange color. On the left of the Edit ESI-LAG page, a list of client devices are displayed.

If one of the peer devices is *Unknown*, the Peer Devices tab is selected, and it displays in orange color. On the left of the Edit ESI-LAG page, a list of peer devices, that are of the same type and ELS capability as of the discovered peer, are displayed.

On the right of the Edit ESI-LAG page, a schematic diagram of the existing two peer devices PEER¹ 1, PEER² 2 and a representation of the client devices as boxes are displayed.

Managing Peer Devices and Peer-to-Peer Link Settings

To add, edit, or delete a peer port, or edit peer-to-peer link settings:

1. Click **EVPN-VXLAN** link that is displayed between PEER¹ 1 and PEER² 2 in the schematic diagram.

The Peer to Peer Link Settings window opens.

The peer ports that you already configured are displayed in the Peer to Peer Ports table.

2. Do one of the following:

- To add a port, click **Add Port**.

A new row is added to the table, where you must enter the port details for the peer devices.

From the drop-down menu for the PEER¹ 1 device or PEER 2 device, select a port to assign to the ESI-LAG.

- To edit a peer port, click the port and edit the port.
- To delete a peer port, select the port that you want to remove from the ESI-LAG, and click **Remove Port**.

NOTE: You must specify at least one link between the peer devices.

3. Click **Update**.
4. (Optional) Edit the Peer 1 or Peer 2 loop back address, Peer 1 or Peer 2 logical interfaces. BGP Group Name, Virtual Instance Switch Name, VRF Instance Name, or Autonomous System Number.
5. (Optional) Edit the Peer To Peer VLANs as follows:
 - To add a VLAN, click **Add VLAN**.
A new row is added to the Peer To Peer VLANs table. Enter the values for VLAN ID, VLAN Name, VNI ID, IP type, IP addresses and mask for the integrated routing and bridging interface of the VLAN and the anycast gateway in the corresponding fields. Click **Update** to save the new VLAN in the table.
 - To edit a VLAN, click on the VLAN and edit one or more attributes of the VLAN—VLAN ID, VLAN Name, VNI ID, IP type, IP addresses and mask for the integrated routing and bridging interface of the VLAN and the anycast gateway.
 - To remove a VLAN, select the VLAN and click **Remove VLAN**.
The VLAN is removed from the table.
6. Click **OK**.
The Peer to Peer Link Settings window closes, and the Edit ESI-LAG page is displayed.

Managing Client Devices and Client-to-Peer Link Settings

To add or remove client devices, and edit client-to-peer link settings:

1. Click Client Devices tab on the Edit ESI-LAG page.
The Client Devices tab on the Create ESI-LAG page lists switches that are managed by Network Director. On the right, a schematic diagram of the two peer devices PEER1 PEER2 and a representation of the client devices as boxes is displayed.
2. Select a device from the list of client devices, and drag and drop it into one of the boxes labeled as *Drag & Drop Clients here to add*.

NOTE: To delete a client device from an ESI-LAG configuration, click the x mark on the client device in the carousel. The client device is removed from the carousel.

You cannot delete a client device if the client device is already deployed.

The Client to Peer Link Settings window opens.

3. In the Client to Peer Ports table, you can edit or delete the configured client and peer ports.

To remove the configured client and peer ports, click on a row and click **Remove Port**. The port is removed from the Client to Peer Ports table.

To add a port, click **Add Port** and select the client and peer ports. Click **Update** to save the port.

To edit a port, click on the port and edit the port. Click **Update** to save the edited port values.

NOTE: You can add only one interface connection between a client and a peer.

4. In the Client to Peer VLANs* table, in the Client to Peer Link Settings window, Network Director displays the VLAN configured in the ESI-LAG

In this table, you can edit the configured VLAN, add a new VLAN or remove a configured VLAN.

To add a new VLAN, click **Select VLAN** and select the VLANs from the list that displays in the Choose VLAN Profile pop-up window.

To edit a VLAN, click on the VLAN and edit the VLAN ID or VLAN Name.

You can remove a VLAN that you have created, but not deployed, in the client device by selecting the VLAN and clicking **Remove VLAN**.

5. Click **Update**.
6. Click **OK** to submit the settings that you entered in the Client to Peer Link Settings window.

The Client to Peer Link Settings window closes.

7. Click **Save** in the Manage ESI-LAG page.

Network Director saves the ESI-LAG settings and displays the message **ESI-LAG save is successful and is ready to be deployed to the devices**.

8. Click **OK**.

The Manage ESI-LAG page lists the newly created ESI-LAG. By default, the Deployment State for the edited ESI-LAG displays as Pending Deployment.

To deploy the edited ESI-LAG, see [“Deploying ESI-LAG Configuration” on page 676](#).

Deleting an ESI-LAG

To delete ESI-LAG:

1. Click the Build mode icon  **Build** in the Network Director banner.

2. Select **Wired > Tasks > Manage ESI-LAG** in the Tasks pane.

The Manage ESI-LAG page opens, displaying the configured ESI-LAG peers and enables you to delete ESI-LAGs.

3. Click **Delete** for the corresponding ESI-LAG Peers that you want to delete, in the Manage ESI-LAG page.

NOTE: If you delete an ESI-LAG, Network Director removes the ESI-LAG configuration settings from the peer devices and also deletes the LAG configuration from the client devices. The Deployment State changes to Pending Removal if the ESI-LAG is already deployed. If it is not deployed, that is, if it is Pending Deployment, then the ESI-LAG is removed from the Manage ESI-LAG page.

ESI-LAG Automation Parameters

Network Director configures a number of parameters internally and automates the creation or modification of ESI-LAGs.

[Table 147](#) describes the parameters that are internally configured by Network Director.

Table 147: ESI-LAG Automation Parameters

Parameter	Description
LAG	Used to create a LAG between peer devices, and between the client and peer devices.
LACP active	Used to configure LACP in peer devices and the client device. LACP is used to discover multiple links from a client device connected to peers. LACP must be configured on all member links to work properly.
LACP periodic fast	Used to configure LACP periodic fast in Peer switches and client switch. LACP fast periodic is achieved by configuring fast intervals (in seconds) for periodic transmission of LACP.
Loopback Address	Used to configure loopback address. The loopback address ensures that the device provides an IP address to management applications as the device must always be available to hosts attempting to route packets to the device. Setting a loopback address ensures that the device can receive packets addressed to the loopback address as long as the device is reachable through any entry (ingress) interface.

Table 147: ESI-LAG Automation Parameters (*continued*)

Parameter	Description
ESI ID	Used to configure an Ethernet Segment Identifier (ESI) on a per-interface basis. All interfaces configured with the same ESI, on any devices within the same EVPN domain, appear as part of the same L2 segment or LAG.
ESI mode	Used to configure ESI all-active mode to enable Active-Active Multihoming in peers.
Policy options	Used to specify routing policy evpn-pplb for EVPN.
Routing option	Used to specify forwarding table with per-packet load balancing (PPLB) export policy for EVPN, autonomous system number and router-id.
Virtual Switch Configuration	Used to create a virtual switch routing instance and auto-generate related parameters.
vtep-source-interface	Used to specify the source interface for a Virtual Extensible LAN (VXLAN) tunnel and configure a logical interface unit 0 (lo0.0) on the loopback interface.
instance-type - (virtual-switch)	Used to provide support for Layer 2 bridging. This routing instance type is used to isolate a LAN segment with its Spanning Tree Protocol (STP) instance and to separate its VLAN identifier space.
route distinguisher id	Used to specify a route distinguisher for the routing instance.
vrf-target	Used to specify a virtual routing and forwarding (VRF) target community
Protocol	Used to enable the Ethernet VPN (EVPN) protocol.
VLAN-VxLAN Mapping	Used to map the VLAN and VxLAN configuration to the virtual switch
VPN routing and forwarding (VRF) instance	Used to create a VRF routing instance and auto-generate the VRF routing parameters.
instance-type (vrf)	Used to provide support for Layer 3 VPNs, where interface routes for each instance goes only into the corresponding forwarding table
vrf-target	Used to specify a virtual routing and forwarding (VRF) target community
route distinguisher id	Used to assign a route distinguisher to the routing instance automatically.
BGP Protocol	Used to enable BGP protocol on peer devices and auto-generates the BGP parameters.

Table 147: ESI-LAG Automation Parameters (continued)

Parameter	Description
BGP sessions	Used to auto-generate internal group type, set the loopback address as the local address and the peer loop back address as the neighbor address.
VPN family	Used to auto-generate inet, inet-vpn, and evpn signaling for BGP.
mutipath	Used to allow load sharing among multiple eBGP and multiple iBGP paths.
OSPF Protocol	Used to enable OSPF on peer devices and configur egaArea IP as 0.0.0.0 on loopback and LAG Interface.

RELATED DOCUMENTATION

[Understanding Link Aggregation | 646](#)

[Managing and Creating a Link Aggregation Group | 647](#)

[Viewing Profiles Assigned to a Device | 1144](#)

[Network Director Documentation home page](#)

Configuring Fibre Channel Gateways

IN THIS CHAPTER

- [Configuring Fibre Channel Gateways | 683](#)
- [Creating and Managing FC Gateway Service Profiles | 684](#)
- [Assigning an FC Gateway Service Profile to Ports | 688](#)

Configuring Fibre Channel Gateways

IN THIS SECTION

- [Using the FC Gateway Service Profile to Configure FC Gateways | 683](#)
- [Using a Combination of Profiles to Configure FC Gateways | 684](#)

This topic describes the methods Network Director provides to configure Fibre Channel (FC) gateways on data center switching devices.

This topic describes:

Using the FC Gateway Service Profile to Configure FC Gateways

An FC Gateway Service profile provides a quick way to configure Fibre Channel (FC) gateways on Data Center Switching devices. It is intended for FC gateway devices that are connected directly to the FCoE network, without transit switches. In addition to the settings you can specify, the profile creates an FC gateway configuration with some default settings that cannot be modified. For example, you cannot specify any CoS settings for the FC gateway; only the default settings are available.

For information about using the FC Gateway Service profile, see [“Creating and Managing FC Gateway Service Profiles” on page 684](#).

Using a Combination of Profiles to Configure FC Gateways

If you want to configure an FC gateway that does not meet the requirements of the FC Gateway Service profile, you can use a combination of other profiles to configure the components of the FC gateway and assign them to the appropriate devices, ports, and profiles.

The profiles you use to configure an FC gateway include:

- CoS profile—Configures CoS on the FC gateway Ethernet interfaces that is required to support FCoE. You assign the CoS profile to the port profiles you assign to the Ethernet ports in the FC gateway.
- VLAN profiles—Configure the FCoE VLANs used to carry FCoE traffic on the FC gateway and the native VLAN used to transport FIP traffic. You assign the FCoE VLAN profile to the Fabric profile and the Port profiles assigned to the FC and Ethernet ports in the FC gateway. You assign the native VLAN profile to the port profile for the ports in the FC gateway.
- Port profiles—Configure FC and Ethernet ports as FC gateway ports, and configure Ethernet port VLAN membership and FIP settings. You assign Port profiles to the FC and Ethernet ports in the FC gateway. The Port profile includes the CoS and VLAN profiles assigned to it.
- Fabric profile—Configures the FC gateway fabric. The Fabric profile includes the FCoE VLAN profile assigned to it.

RELATED DOCUMENTATION

[Creating and Managing FC Gateway Service Profiles | 684](#)

[Creating and Managing Wired CoS Profiles | 613](#)

[Creating and Managing Fabric Profiles | 695](#)

[Creating and Managing Port Profiles | 414](#)

[Creating and Managing VLAN Profiles | 502](#)

[Network Director Documentation home page](#)

Creating and Managing FC Gateway Service Profiles

IN THIS SECTION

● [Managing FC Gateway Service Profiles | 685](#)

● [Creating FC Gateway Service Profiles | 686](#)

- Specifying Settings for an FC Gateway Service Profile | 687
- What to Do Next | 688

An FC Gateway Service profile provides a quick way to configure Fibre Channel (FC) gateways on Data Center Switching devices. It is intended for FC gateway devices that are connected directly to the FCoE network, without transit switches. In addition to the settings you can specify, the profile creates an FC gateway configuration with some default settings that cannot be modified. For example, you cannot specify any CoS settings for the FC gateway; only the default settings are available.

If you want to configure an FC gateway that does not meet these requirements, you can use other profiles to configure the components of the FC gateway and assign them to the appropriate devices, ports, and profiles. See “[Configuring Fibre Channel Gateways](#)” on page 683 for more information.

To manage or create FC Gateway Service profiles: In Build mode, select **Wired > System > FC Gateway Service** in the Tasks pane. The Manage Fabric Profile page appears.

This topic describes:

Managing FC Gateway Service Profiles

From the FC Gateway Service Profile page, you can:

- Create a new profile by clicking **Add**.
- Modify an existing profile by selecting it and clicking **Edit**.
- Associate a profile to specific interfaces by selecting it and clicking **Assign**.

During the assignment process, you will have the option to configure interface-specific settings.

- Change a profile’s current interface assignments by selecting it and clicking **Edit Assignments**.
- View information about a profile by selecting the group and clicking **Details** or by clicking the profile name.
- Clone a profile by selecting a profile and clicking **Clone**.
- Delete profiles by selecting the profiles and clicking **Delete**.

TIP: You cannot delete profiles that are in use—that is, assigned to objects or used by other profiles. To see the current assignments for a profile, select the profile and click **Details**.

Table 148 describes the information provided about FC Gateway Service profiles on the Manage FC Gateway Service Profile page. This page lists all FC Gateway Service profiles defined for your network, regardless of your current selected scope in the network view.

Table 148: FC Gateway Service Profile Information

Field	Description
Profile Name	Name given to the profile when it was created.
Fabric ID	Fabric ID number assigned to the profile.
FCoE VLAN Name	Name of the FCoE VLAN assigned to the profile.
FCoE VLAN Id	ID of the FCoE VLAN assigned to the profile.
Native VLAN Name	Name of the native VLAN assigned to the profile.
Native VLAN Id	ID of the native VLAN assigned to the profile.
Description	Description given to the profile when it was created.
Assignment State	Profile assignment state. One of the following: <ul style="list-style-type: none"> • Deployed—The profile has been assigned and the configuration has been deployed on the devices. • Pending Deployment—The profile has been assigned or its previous assignments have been changed, but the new or modified configuration has not yet been deployed on the devices. • Unassigned—The profile has not yet been assigned.

TIP: All columns might not be displayed. To show or hide fields in the table, click the down arrow on the field header, select Columns, and select or clear the check box adjacent to the field that you want to show or hide.

Creating FC Gateway Service Profiles

To create an FC Gateway Service profile:

1. From the Network Director Banner, select **Build** mode.
2. In the Tasks pane, select **Wired > System > FC Gateway Service**.

The FC Gateway Service Profiles page appears.

3. Click **Add**.

The Create Gateway Profile page appears.

4. Enter settings for the FC Gateway Service profile as described in [“Specifying Settings for an FC Gateway Service Profile” on page 687](#).
5. Click **Done**.

Specifying Settings for an FC Gateway Service Profile

Use the Create Gateway Profile page to define specify settings for an FC Gateway Service profile, which you can then assign to interfaces on Data Center Switching devices.

[Table 149](#) describes the settings.

Table 149: FC Gateway Service Profile Settings

Field	Description
Profile Name	Enter a name for the profile.
Description	Enter a description for the profile.
Fabric ID	Enter an FC gateway fabric ID number from 1 through 4094.
FCoE VLAN Name and ID	Enter an FCoE VLAN Name and ID.
Native VLAN Name and ID	Enter a native VLAN Name and ID.
FC Map	The FC-MAP value is a unique MAC address prefix an FC switch uses to identify FCoE traffic for a given FC fabric (traffic on a particular FCoE VLAN). Range: 0x0EFC00 through 0x0EFCFF. Default: 0xEFC00
Priority	Set the global priority value associated with the switch FCF-MAC. CNAs use the priority value to determine the switch with which they will perform FIP FLOGI. The lower the value, the higher the priority. The switch advertises this value to the server ENodes on the FCoE network. Range: 0 through 255. Default: 128

Table 149: FC Gateway Service Profile Settings (continued)

Field	Description
FCoE Trusted	Select to enable the interface to trust FCoE traffic. If an interface is connected to another switch such as an FCoE forwarder (FCF) or a transit switch, you can configure the interface as trusted so that the interface forwards FCoE traffic from the switch to the FCoE devices without installing FIP snooping filters.
Port Max Packet Size	Select the maximum packet size. Range: 256-9216. Default: 2500

What to Do Next

Once the FC Gateway Service profile is created, you must assign the profile to the required devices by using the Assign Device Profile page and then deploy the device configuration changes by using Deploy mode. To assign an FC Gateway Service profile, see [“Assigning an FC Gateway Service Profile to Ports” on page 688](#). For information about deploying your configurations, see [“Deploying Configuration to Devices” on page 1180](#).

RELATED DOCUMENTATION

[Assigning an FC Gateway Service Profile to Ports | 688](#)

[Network Director Documentation home page](#)

Assigning an FC Gateway Service Profile to Ports

IN THIS SECTION

- [Assigning an FC Gateway Service Profile | 689](#)
- [Editing the Assignments of an FC Gateway Service Profile | 692](#)

You assign an FC Gateway Service profile to the ports on Data Center Switching devices that you want to include in a Fibre Channel (FC) gateway. You must assign the profile to at least one FC port and at least one Ethernet or aggregated Ethernet interface on the device. In a QFabric fabric, an FC gateway cannot span multiple nodes.

This topic describes:

Assigning an FC Gateway Service Profile

To assign an FC Gateway Service profile:

1. From the Network Director Banner, select **Build** mode.
2. In the Tasks pane, select **Wired > Profiles > FC Gateway Service**.

The FC Gateway Service Profiles page appears. The page displays all the FC Gateway Service profiles that you configured as well as the system-created profiles detected during device discovery.

3. Select a profile from the list of profiles and click **Assign**.

The Assign FC Gateway Service Common Settings page opens to the Object Selection page.

NOTE: If Network Director fails to read the configuration of one or more devices after the device discovery, such devices are not displayed in the Assign Profile page. You will not be able to assign profiles to such devices. The Manage Jobs page in System mode displays details of the device discovery jobs. Use the information displayed on this page to take appropriate corrective steps to enable Network Director to reread the configuration of the failed device. For more information, see [“Discovering Devices in a Physical Network” on page 204](#).

4. Enable either **Select Devices** or **Select Port Groups**.
5. If you enabled **Select Port Groups**, select one or more port groups from the Select Port Group list.
6. If you enabled **Select Devices**, expand the list of objects and select the objects that contain the devices and interfaces you want to assign by clicking the check box next to them.

If you select a container node, all devices under that node are selected. The list of objects is filtered to include only devices that match the profile’s family type. Be sure that a check mark appears in the corresponding check box—highlighting the object name is not sufficient.

7. Click either **Next** or **Profile Assignment**.

The Profile Assignment page displays the list of existing assignments in the Assignments table.

8. To assign the profile to ports on a device:

a. Select one or more container nodes or devices from the Assignments list:

- To assign the profile to nonconsecutive interfaces or to aggregated Ethernet interfaces, select a single device.
- To assign the profile to interfaces in the same consecutive interface range (for example, ge-0/0/0 through ge-0/0/15) on one or more devices, select one or more devices. To make multiple selections, press Shift or Ctrl while making the selections.
- To assign a profile to ports within a QFabric system select the member node group or groups that contain the ports.
- To assign a profile to ports within a Virtual Chassis Fabric (VCF), you can select any container nodes or member devices within the VCF, including the VCF container node.
- To assign a profile to aggregated Ethernet ports within a Virtual Chassis or VCF, select the Virtual Chassis or VCF container node. To assign a profile to physical device ports within a Virtual Chassis or VCF, select one or more member devices.
- Channelized ports are only applicable for Data Center Switching ELS devices and only XE interfaces can be used as channelized ports.

b. Click the **Assign to Port** button from the list.

The Assign Profile to Ports window opens.

c. Select either **Ports** (default) or **Port Range**. If you selected multiple devices in the previous step, you cannot choose the Port option.

d. If you selected the Port option, select the ports from the list of ports.

By default, aggregated Ethernet interfaces are listed after the ge- and xe- interfaces in the list of ports. Members of aggregated Ethernet interfaces are not included in the port list.

e. If you selected the Port Range option, enter the port range:

- i. In the Normal Ports section, enter a first and last port name in the text boxes, then click the **Add** button. The port range appears in the Selected Port Range section.
- ii. Repeat the add process to add any additional port ranges.
- iii. To delete a port range, select its check box, then click the **Delete** button.

At least one port within the port range must be available on each selected device for the port range to succeed. Assignments are created for the ports within the port range that are available. You can

assign the profile to the same interface on multiple devices by entering the interface name in both fields of the port range.

9. Click **Assign** to complete the port assignments and close the window.
10. On the Profile Assignment page, you can view the assignment details for the selected device and also remove any assignments:
 - To view the assignment details, select the device and click **View Assignments**.
The Profile Details page for selected device appears. Expand the **Device** name to view the details of the assignment. The assignment status displays the status whether the device is deployed or is pending device update, and so on.
 - To delete a device from the assignment list, select the device from the Assignments table and click **Remove**.
11. Click **Next** or click **Review** from the top wizard workflow to review the assignments. Use the **Edit** button if you want to edit the profile assignment.
12. Click **Finish** once you are done reviewing the profile assignment.

After you click Finish, the Create Profile Assignments Job Details dialog box appears, which reports on the status of the profile assignment job. If you have assigned the profile to a large number of objects, the profile assignment job can take some time to complete. Instead of waiting for the Job Details dialog box to report job completion status, you can close it and check the details of the profile assignment job at a later time using the Manage Job task in System mode.

NOTE: If any assignment fails, the profile assignment job fails and none of the assignments are created. Check the details for the profile assignment job for information about why the assignment failed.

After you assign an FC Gateway Service profile to a device, you can deploy the profile from the **Deploy** mode. For details, see [“Deploying Configuration to Devices” on page 1180](#).

To view the details of a profile, select the profile from the Manage Fabric Profiles page and click the **Details** button.

Editing the Assignments of an FC Gateway Service Profile

To edit an FC Gateway Service profile's port assignments:

1. Select a profile from the Manage FC Gateway Service Profile page and click **Edit Assign**.
The Edit Assignments page opens.
2. Expand the nodes in the list to locate the devices and ports you want to change.
3. To delete a device or port from a profile, click **Delete** in the **Operation** column of the device or port.
An X appears in the Record Status column.
4. To delete multiple devices or ports from a profile, select the devices or ports, then click the **Delete** button.
An X appears in the Record Status column of the deleted devices or ports.
5. When you are done making assignment changes, click **Apply**.
The Manage FC Gateway Service Profile page opens.

RELATED DOCUMENTATION

[Creating and Managing FC Gateway Service Profiles | 684](#)

[Deploying Configuration to Devices | 1180](#)

[Network Director Documentation home page](#)

Creating Configurations for Fabrics

IN THIS CHAPTER

- Understanding Fabric Profiles | 693
- Creating and Managing Fabric Profiles | 695
- Assigning a Fabric Profile to Devices and Ports | 701

Understanding Fabric Profiles

A Fabric profile contains configuration settings for a gateway Fibre Channel (FC) fabric. You assign Fabric profiles to QFX Series devices that act as a FCoE-FC gateway, to configure gateway FC fabrics.

A gateway FC fabric is a QFX Series configuration construct. It is not the same thing as an FC fabric in the storage area network (SAN); the gateway FC fabric is local to the switch. It creates associations that connect FCoE devices with converged network adapters (CNAs) on the Ethernet network to an FC switch on the Fibre Channel network. A gateway FC fabric consists of:

- A unique fabric name.
- A unique fabric ID.
- At least one dedicated VLAN for FCoE traffic. VLANs that carry FCoE traffic must not carry any other type of traffic.

NOTE: On a QFX3500 or QFabric system Node device, the same VLAN cannot be used in both transit switch mode and FCoE-FC gateway mode.

- At least one FCoE VLAN interface (Layer 3 VLAN interface) that includes one or more 10-Gigabit Ethernet interfaces connected to FCoE devices. The FCoE VLANs transport traffic between the FCoE servers and the FCoE-FC gateway. Each FCoE VLAN must carry only FCoE traffic. You cannot mix FCoE traffic and standard Ethernet traffic on the same VLAN.

The 10-Gigabit Ethernet interfaces that connect to FCoE devices must include a native VLAN to transport FIP traffic because FIP VLAN discovery and notification frames are exchanged as untagged packets.

Each FCoE VLAN interface can present multiple VF_Port interfaces to the FCoE network.

NOTE: Storm control must be disabled on all Ethernet interfaces that belong to the FCoE VLAN to prevent FCoE traffic from being dropped.

- One or more native FC interfaces. The native FC interfaces transport traffic between the gateway and the FC switch.

TIP: If the network does not use a dual-rail architecture for redundancy, configure more than one native FC interface for each FC fabric to create redundant connections between the FCoE devices and the FC switch. If one physical link goes down, any sessions it carried can log in again and connect to the FC switch on a different interface. Even in dual-rail architecture networks, creating redundant connections between the QFabric system and the FC switch is the best practice.

You can configure FIP parameters for the fabric or accept the default FIP parameters. VN_Port to VF_Port (VN2VF_Port) FIP snooping is automatically enabled on all server-facing ports because all ports are untrusted by default. You can disable VN2VF_Port FIP snooping on a port-by-port basis by marking a port as an FCoE trusted interface. You can disable VN2VF_Port FIP snooping on all Ethernet ports in an FC fabric by configuring the fabric as FCoE trusted.

Because the switch has 12 native FC ports and each FC fabric requires a minimum of one native FC port, the switch supports a maximum of 12 FC fabrics. However, as a best practice for redundancy, we recommend that you assign at least two native FC interfaces to each FC fabric.

On a QFabric system, all of the FC and FCoE traffic that belongs to a particular gateway FC fabric must ingress and egress the same gateway Node device. Gateway FC fabrics do not span across Node devices. All of the native FC interfaces and the Ethernet interfaces that belong to the FCoE VLAN must reside on the same gateway Node device to be included in an FC fabric on that Node device.

Traffic from FC and FCoE devices that are not in the same FC fabric remain separate and cannot communicate with each other through the gateway.

RELATED DOCUMENTATION

[Creating and Managing Fabric Profiles | 695](#)

[Network Director Documentation home page](#)

Creating and Managing Fabric Profiles

IN THIS SECTION

- [Managing Fabric Profiles | 695](#)
- [Creating Fabric Profiles | 697](#)
- [Specifying Settings for a Fabric profile | 697](#)
- [What to Do Next | 701](#)

A Fabric profile contains configuration settings for a gateway Fibre Channel (FC) fabric. A gateway FC fabric creates associations that connect Fibre Channel over Ethernet (FCoE) devices with converged network adapters (CNAs) on the Ethernet network to an FC switch on the Fibre Channel network. You can assign Fabric profiles to QFX Series devices that act as a FCoE-FC gateway, to configure gateway FC fabrics.

To manage or create FC Gateway Service profiles: In Build mode, select **Wired** > **System** > **Fabric** in the Tasks pane. The FC Gateway Service Profile page appears.

This topic describes:

Managing Fabric Profiles

From the Manage Fabric Profiles page, you can:

- Create a new profile by clicking **Add**.
- Modify an existing profile by selecting it and clicking **Edit**.
- Associate a profile to specific interfaces by selecting it and clicking **Assign**.

During the assignment process, you will have the option to configure interface-specific settings.

- Change a profile's current interface assignments by selecting it and clicking **Edit Assignments**.
- View information about a profile by selecting the group and clicking **Details** or by clicking the profile name.
- Clone a profile by selecting a profile and clicking **Clone**.
- Delete profiles by selecting the profile and clicking **Delete**.

TIP: You cannot delete profiles that are in use—that is, assigned to objects or used by other profiles. To see the current assignments for a profile, select the profile and click **Details**.

[Table 150](#) describes the information provided about Fabric profiles on the Manage Fabric Profiles page. This page lists all Fabric profiles defined for your network, regardless of your current selected scope in the network view.

Table 150: Fabric Profile Information

Field	Description
Fabric Name	Name given to the profile when it was created.
Fabric ID	Number from 1 through 4094 assigned to the Fabric profile when it was created.
Device Family	Device family to which the profile applies.
Description	Any description added when the profile was created.
Assignment State	One of the following: <ul style="list-style-type: none"> • Deployed—The profile has been assigned to interfaces and the configuration has been deployed on the devices. • Pending Deployment—The profile has been assigned to interfaces or its previous assignments have been changed, but the new or modified configuration has not yet been deployed on the devices. • Unassigned—The profile has not yet been assigned to interfaces.
Creation Time	Date and time when this profile was created.
Last Updated Time	Date and time when this profile was last modified.
Username	The username of the user who created or modified the profile.

TIP: All columns might not be displayed. To show or hide fields in the table, click the down arrow on the field header, select Columns, and select or clear the check box adjacent to the field that you want to show or hide.

Creating Fabric Profiles

To create a Fabric profile:

1. From the Network Director Banner, select **Build** mode.
2. In the Tasks pane, select **Wired > System > Fabric**.
The Manage Fabric profiles page appears.
3. Click **Add**.
The Create Fabric Profile page appears.
4. Enter settings for the Fabric profile as described in [“Specifying Settings for a Fabric profile” on page 697](#).
5. Click **Done**.

Specifying Settings for a Fabric profile

Use the Create Fabric profile page to define a common set of attributes in a Fabric profile, which you can then apply to Fibre Channel (FC) interfaces on Data Center Switching devices.

TIP: You can reference VLAN profiles in a Fabric profile. Create these profiles before you create Fabric profiles.

After you create a Fabric profile, you can assign it to devices, individual interfaces, or port groups. During this process, you can also configure interface-specific attributes. You can assign only one Fabric profile to an interface or port group.

[Table 151](#) describes the Fabric profile settings. Mandatory settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 151: Fabric Profile Settings

Field	Action
Fabric Name	Type a name for the profile. The value must be a string between 2 and 64 characters long, begin with a letter, and consist of letters, numbers, periods, dashes, and underscores only.
Fabric ID (1-4094)	Type an ID number from 1 through 4094.

Table 151: Fabric Profile Settings (continued)

Field	Action
Description	Type a description of the Fabric profile, which will appear on the Manage Fabric Profiles page.
Max Login Sessions (128-2500)	Set the maximum number of FCoE initialization protocol (FIP) session logins permitted for the FCoE-FC gateway fabric. The range is 128 through 2500. The default is 2500.
VLAN Profiles	
Add	Click Add to select FCoE VLAN profiles to assign to the Fabric profile. Select the VLAN profiles from the list, then click OK .
Remove	To remove VLAN profiles, select one or more profiles, then click Remove .
Proxy Settings	

Table 151: Fabric Profile Settings (continued)

Field	Action
Load Balance Algorithm	<p>Set the load-balancing algorithm that the QFX Series uses to distribute FCoE sessions (FLOGI and FDISC sessions from the FCoE devices in the Ethernet network) among the NP_Port links to the FC switch.</p> <p>NOTE: Changing the load-balancing algorithm when FCoE sessions are running forces the FCoE sessions to log out, and then log in again.</p> <p>Select one of the following:</p> <ul style="list-style-type: none"> ● simple—Load balancing is based on the weighted utilization (load) of the NP_Ports connected to an FC fabric. Each new FLOGI or FDISC is assigned to the least-loaded link. When a link load rebalance occurs, the system minimizes disruption by using an algorithm to log out only the sessions that need to be moved to other links to balance the link load. To further minimize disruption, the algorithm logs out the sessions with the fewest dependencies (for example, FDISC sessions are logged out before FLOGI sessions). When the sessions log in again, they are placed on NP_Port interfaces in a manner that balances the link loads. This is the default load-balancing algorithm. ● enode-based—Load balancing is based on the ENode FLOGI. When an ENode logs in to the fabric, all subsequent FDISC sessions (VN_Port sessions) associated with that ENode are placed on the same link as the ENode FLOGI session, regardless of the link load. New ENode FLOGIs are placed on the least-loaded link. When a link load rebalance occurs, the system logs off all sessions. The sessions log in again and are placed on NP_Port interfaces in a balanced manner. ● flogi-based—FLOGI-based load balancing is similar to ENode-based load balancing, but the behavior when the loads are rebalanced is different. Load balancing is based on the ENode FLOGI. When an ENode logs in to the fabric, all subsequent FDISC sessions associated with that ENode are placed on the same link as the ENode FLOGI session, regardless of the link load. New ENode FLOGIs are placed on the least-loaded link. When a link load rebalance occurs, the system minimizes disruption by using an algorithm to log out only the sessions that need to be moved to other links to balance the link load. When the logged out sessions log back in, they are placed on NP_Port interfaces in a manner that balances the link loads.
Disable Worldwide Name Verification (default is disabled)	<p>Select to Disable the fabric worldwide name (WWN) verification check in the fabric login accept message (FLOGI-ACC) for implicit FLOGIs.</p> <p>If you enable this option, when a QFX Series NP_Port performs a FLOGI to the FC fabric, the QFX Series does not verify the fabric WWN in the FLOGI-ACC against the current fabric WWN.</p> <p>NOTE: Disabling or enabling the fabric WWN verification check logs out all FCoE sessions.</p>

Table 151: Fabric Profile Settings (continued)

Field	Action
FIP Settings	
FCoE Trusted (default is disabled)	<p>Select to enable the specified 10-Gigabit Ethernet interface to trust Fibre Channel over Ethernet (FCoE) traffic. If an interface is connected to another switch such as an FCoE forwarder (FCF) or a transit switch, you can configure the interface as trusted so that the interface forwards FCoE traffic from the switch to the FCoE devices without installing FIP snooping filters.</p> <p>Configure the specified local Fibre Channel fabric to trust FCoE traffic on all ports in the fabric. Changing the fabric ports from untrusted to trusted removes any existing FIP snooping filters from the ports. Changing the fabric ports from trusted to untrusted by removing the fcoe-trusted configuration from the fabric forces all of the FCoE sessions on those ports to log out so that when the ENodes and VN_Ports log in again, the switch can build the appropriate FIP snooping filters.</p>
FCoE Mapped Address	<p>Set the FCoE mapped address prefix (FC-MAP) value for the FCoE VLAN to match the FC switch (or FCoE forwarder) FC-MAP value for the FC fabric. The FC-MAP value is a unique MAC address prefix an FC switch uses to identify FCoE traffic for a given FC fabric (traffic on a particular FCoE VLAN).</p> <p>Enter an FC-MAP value, which is a hexadecimal value preceded by "0x". The range is 0x0EFC00 through 0x0EFCFF. The default value is 0xEFC00.</p> <p>You can configure the FC-MAP value or use the default value. The FC switch provides the FC-MAP value to FCoE nodes (ENodes) in the FIP discovery advertisement message. If the EX Series switch or the QFX Series FCoE VLAN FC-MAP value does not match the FC switch FC-MAP value, neither device discovers the FC switch on that VLAN, and the ENodes on that VLAN cannot access the FC switch. The FC switch accepts only FCoE traffic that uses the correct FC-MAP value as part of the VN_Port MAC address.</p> <p>When the QFX Series acts as an FCoE-FC gateway, the FC-MAP value for the gateway and the FCoE devices must match the FC switch FC-MAP value in order to communicate with the FC switch.</p> <p>NOTE: Changing the FC-MAP value causes all logins to drop and forces the ENodes to log in again.</p>
Keep Alive Interval (default is 8000)	Set the global interval between periodic FIP keepalive advertisements, in milliseconds. The range is 250 through 90000 milliseconds. The default is 8000 milliseconds.

Table 151: Fabric Profile Settings (continued)

Field	Action
Priority (default is 128)	<p>Set the global priority value associated with the switch FCF-MAC. CNAs use the priority value to determine the switch with which they will perform FIP FLOGI. The lower the value, the higher the priority. The switch advertises this value to the server ENodes on the FCoE network.</p> <p>The range is 0 through 255. The default value is 128.</p>
Max Sessions Per Node	<p>Set the maximum number of FCoE login sessions (FLOGI plus FDISC) from a single ENode allowed on the gateway FC fabric (the fabric configured on the QFabric system). The maximum number of logins per ENode is 2000 sessions.</p> <p>NOTE: A session is a FLOGI or FDISC login to the FC SAN fabric. Session does not refer to end-to-end storage sessions. There is no limit to the number of end-to-end storage sessions.</p>

What to Do Next

Once the Fabric profile is created, you must assign the profile using the Assign Device Profile page to the required devices and then deploy the device configuration changes using the Deploy mode. To assign a Fabric profile, see [“Assigning a Fabric Profile to Devices and Ports” on page 701](#). For information about deploying your configurations, see [“Deploying Configuration to Devices” on page 1180](#).

RELATED DOCUMENTATION

[Understanding Fabric Profiles | 693](#)

[Assigning a Fabric Profile to Devices and Ports | 701](#)

[Network Director Documentation home page](#)

Assigning a Fabric Profile to Devices and Ports

IN THIS SECTION

- [Assigning a Fabric Profile | 702](#)
- [Editing the Assignments of a Fabric Profile | 705](#)

A Fabric profile contains configuration settings for a gateway Fibre Channel (FC) fabric. You can assign an existing user-created or system-created Fabric profile to devices and FC ports to configure FC fabrics, using the steps described in this topic. A Fabric profile can be assigned to a QFX3500 device in standalone mode or to a QFabric Node.

Gateway FC fabrics on a QFabric system are called local FC fabrics. A local FC fabric does not span Node devices and does not span the fabric Interconnect device. Local FC fabrics are entirely contained on a single Node device. You can assign multiple interfaces on one Node device to a gateway FC fabric.

This topic describes:

Assigning a Fabric Profile

To assign a Fabric profile:

1. From the Network Director Banner, select **Build** mode.
2. In the Tasks pane, select **Wired > Profiles > Fabric**.

The Manage Fabric profiles page appears. The page displays all the Fabric profiles that you configured as well as the system-created profiles detected during device discovery.

3. Select a profile from the list of profiles and click **Assign**.

The Assign Fabric Profile wizard opens to the Device Selection page.

NOTE: If Network Director fails to read the configuration of one or more devices after the device discovery, such devices are not displayed in the Assign Profile page. You will not be able to assign profiles to such devices. The Manage Jobs page in System mode displays details of the device discovery jobs. Use the information displayed on this page to take appropriate corrective steps to enable Network Director to reread the configuration of the failed device. For more information, see [“Discovering Devices in a Physical Network” on page 204](#).

4. Enable either **Select Devices** or **Select Port Groups**.
5. If you enabled **Select Port Groups**, select one or more port groups from the Select Port Group list.
6. If you enabled **Select Devices** during Object Selection, expand the list of objects and select the objects that contain the devices and interfaces you want to assign by clicking the check box next to them.

If you select a container node, all devices under that node are selected. The list of objects is filtered to include only devices that match the profile’s family type. Be sure that a check mark appears in the corresponding check box—highlighting the object name is not sufficient.

7. Click either **Next** or **Profile Assignment**.

The Profile Assignment page displays the list of existing assignments in the Assignments table.

8. To assign the profile to an entire device, select the device from the Assignments table, then click the **Assign** button, then select **Assign to Device**.

NOTE: If you assign the profile to a device but not to ports on the device, the fabric will be created on the device, but no FC ports will be assigned to the fabric. You must assign FC ports to the profile to add them to the fabric.

9. To assign the profile to ports within the selected objects:

- a. Select one or more container nodes or devices from the Assignments list:

- To assign the profile to nonconsecutive interfaces or to aggregated Ethernet interfaces, select a single device.
- To assign the profile to interfaces in the same consecutive interface range (for example, ge-0/0/0 through ge-0/0/15) on one or more devices, select one or more devices. To make multiple selections, press Shift or Ctrl while making the selections.
- To assign a profile to ports within a QFabric system select the member node group or groups that contain the ports.
- To assign a profile to ports within a Virtual Chassis Fabric (VCF), you can select any container nodes or member devices within the VCF, including the VCF container node.
- To assign a profile to aggregated Ethernet ports within a Virtual Chassis or VCF, select the Virtual Chassis or VCF container node. To assign a profile to physical device ports within a Virtual Chassis or VCF, select one or more member devices.
- Channelized ports are only applicable for Data Center Switching ELS devices and only XE interfaces can be used as channelized ports.

- b. Click the **Assign to Port** button.

The Assign Profile to Ports window opens.

- c. Select either **Ports** (default) or **Port Range**. If you selected multiple devices in the previous step, you cannot choose the Port option.
- d. If you selected the Port option, select the ports from the list of ports.

- e. If you selected the Port Range option, enter the port range:
 - i. In the Normal Ports section, enter a first and last port name in the text boxes, then click the **Add** button. The port range appears in the Selected Port Range section.
 - ii. Repeat the add process to add any additional port ranges.
 - iii. To delete a port range, select its check box, then click the **Delete** button.

At least one port within the port range must be available on each selected device for the port range to succeed. Assignments are created for the ports within the port range that are available. You can assign the profile to the same interface on multiple devices by entering the interface name in both fields of the port range.

- f. Click **Assign** to complete the port assignments and close the window.

10. To configure object-specific attributes for an interface:

- a. Click **Define** from the interface's **Attributes** column in the Assignments table.

The Configure attributes page opens.

- b. Configure the attributes.

For information about these attributes, see [Table 152](#).

- c. Click **Save**.

11. On the Profile Assignment page, you can view the assignment details for the selected device and also remove any assignments:

- To view the assignment details, select the device and click **View Assignments**.

The Profile Details page for selected device appears. Expand the **Device** name to view the details of the assignment. The assignment status displays the status whether the device is deployed or is pending device update, and so on.

- To delete a device from the assignment list, select the device from the Assignments table and click **Remove**.

12. Click **Next** or click **Review** from the top wizard workflow to review the assignments. Use the **Edit** button if you want to edit the profile assignment.

13. Click **Finish** once you are done reviewing the profile assignment.

After you click Finish, the Create Profile Assignments Job Details dialog box appears, which reports on the status of the profile assignment job. If you have assigned the profile to a large number of objects,

the profile assignment job can take some time to complete. Instead of waiting for the Job Details dialog box to report job completion status, you can close it and check the details of the profile assignment job at a later time using the Manage Job task in System mode.

NOTE: If any assignment fails, the profile assignment job fails and none of the assignments are created. Check the details for the profile assignment job for information about why the assignment failed.

Table 152: Fabric Profile Object-Specific Interface Attributes

Field	Description
Keep Alive Interval	Set the global interval between periodic FIP keepalive advertisements, in milliseconds. The range is 250 through 90,000 milliseconds. The default is 8000 milliseconds.
Priority	Set the global priority value associated with the switch FCF-MAC. CNAs use the priority value to determine the switch with which they will perform FIP FLOGI. The lower the value, the higher the priority. The switch advertises this value to the server ENodes on the FCoE network. The range is 0 through 255. The default value is 128.

After you assign a Fabric profile to a device, you can deploy the profile from the **Deploy** mode. For details, see [“Deploying Configuration to Devices” on page 1180](#).

To view the details of a profile, select the profile from the Manage Fabric Profiles page and click the **Details** button.

Editing the Assignments of a Fabric Profile

To edit a Fabric profile’s device and port assignments:

1. Select a profile from the Manage Fabric Profile page and click **Edit Assign**.
The Edit Assignments page opens.
2. Expand the nodes in the list to locate the devices and ports you want to change.
3. To delete a device or port from a profile, click **Delete** in the **Operation** column of the device or port.
An X appears in the Record Status column.
4. To delete multiple devices or ports from a profile, select the devices or ports, then click the **Delete** button.

An X appears in the Record Status column of the deleted devices or ports.

5. To change a device or port's attributes, click **Define** in the **Attributes** column of the device or port, then modify the attributes using the Configure attributes window.

A pencil icon appears in the Record Status column of the edited device or port.

6. When you are done making assignment changes, click **Apply**.

The Manage Fabric Profile page opens.

RELATED DOCUMENTATION

[Understanding Fabric Profiles | 693](#)

[Creating and Managing Fabric Profiles | 695](#)

[Deploying Configuration to Devices | 1180](#)

[Network Director Documentation home page](#)

Creating and Managing Datacenter Fabrics

IN THIS CHAPTER

- Understanding Junos Fusion | 707
- Understanding Junos Fusion Enterprise | 710
- Understanding Junos Fusion Data Center | 713
- Software Requirements for Junos Fusion | 715
- Creating and Managing Fusion Configuration Templates | 716
- Managing Fusion Fabrics | 732
- Creating and Managing Satellite Software Upgrade Groups | 739
- Understanding Layer 3 Fabrics | 742
- User Privileges Required for the DHCP and File Server While Using Zero Touch Provisioning | 743
- Managing Layer 3 Fabrics | 744
- Creating Layer 3 Fabrics | 746
- Editing Layer 3 Fabrics | 761
- Viewing Layer 3 Fabric Connectivity | 764
- Performing Layer 3 Fabric Connectivity Checks | 765
- Setting Up Virtual Chassis Fabrics | 766
- Managing Virtual Chassis Fabrics | 774
- Understanding QFabric System Setup in Network Director | 779
- Setting Up QFabric Systems | 780

Understanding Junos Fusion

Junos Fusion technology, based on the IEEE 802.1BR standard, is a rich, open framework that makes networks highly versatile, extensible, and responsive in multivendor environments. With Junos Fusion technology, network administrators can reduce network complexity and operational expenses by collapsing underlying network elements into a single, logical point of management using QFX Series and EX Series switches running the Junos operating system.

Junos Fusion consists of two major components—*aggregation devices* and *satellite devices*.

Aggregation devices serve as the core of a Junos Fusion fabric and are responsible for almost all management tasks, including interface configuration for every satellite device interface in the topology. An aggregation device runs Junos OS for the entire Junos Fusion. The network-facing interfaces on the satellite devices—extended ports—are configured from the aggregation device and support features that are supported on Junos OS running on the aggregation device. A Junos Fusion fabric can have one aggregation device (single-home Junos Fusion) or two aggregation devices (multihome or dual-home Junos Fusion).

Satellite devices form the access layer of a Junos Fusion fabric. These devices, which are connected through uplink ports to the aggregation devices, need not be individually managed as the control plane resides on the aggregation device.

Network Director supports two Junos Fusion technologies — Junos Fusion Enterprise and Junos Fusion Data Center.

[Table 153](#) lists the device models that are supported as aggregation and satellite devices for each Junos Fusion technology..

Table 153: Devices Supported in Junos Fusion

Junos Fusion Technology	Aggregation Device	Satellite Device
Junos Fusion Enterprise	EX9200	EX2300
		EX3400
		EX4300
Junos Fusion Data Center	QFX10002-46Q	EX4300
	QFX10002-72Q	QFX5100

Network Director provides a single pane of glass (SPOG) solution for managing Junos Fusion Data Center and Junos Fusion Enterprise fabrics, enabling network agility and reducing costs.

Setting up a Junos Fusion Enterprise or Junos Fusion Data Center using Network Director involves the following tasks:

1. Create a configuration template for the instance (Junos Fusion Enterprise or Junos Fusion Data Center). In a configuration template, you specify details such as the fusion topology, ports to be used, software image to be used for satellite devices, and so on.
2. Apply the template to one or more Junos Fusion systems. While applying a template, you select the aggregation devices, specify the software image to be used for the aggregation devices, the DHCP and file server details that Network Director uses to bring up the aggregation devices, and the ICL or ICCP port details for multi-home Junos Fusion.

Network Director creates the Fusion solution based on the template that you applied, and displays the Fusion instance in the Manage Fusion Fabric page.

RELATED DOCUMENTATION

- [Understanding Junos Fusion Enterprise | 710](#)

- [Understanding Junos Fusion Data Center | 713](#)

- [Creating and Managing Fusion Configuration Templates | 716](#)

Understanding Junos Fusion Enterprise

For enterprise networks, Junos Fusion Enterprise provides automated network configuration and simplifies scalability for medium to large enterprise networks with the Juniper Networks EX9200 line of Ethernet switches, EX4300, EX2300, and EX3400 switches. Junos Fusion Enterprise technology can be deployed across a building, or multiple buildings, to connect large numbers of devices in a fabric that can be managed as a single device.

Figure 23 displays a typical Junos Fusion Enterprise topology.

In Junos Fusion Enterprise deployments, satellite devices do not need to be individually connected to aggregation devices. Up to 10 satellite devices can be interconnected through standard 10-Gigabit Ethernet or 40-Gigabit Ethernet interfaces to form a satellite cluster (as shown in Figure 23), which in turn can be connected to the aggregation devices over a pair of fiber uplink ports.

Figure 23: Junos Fusion Enterprise Topology

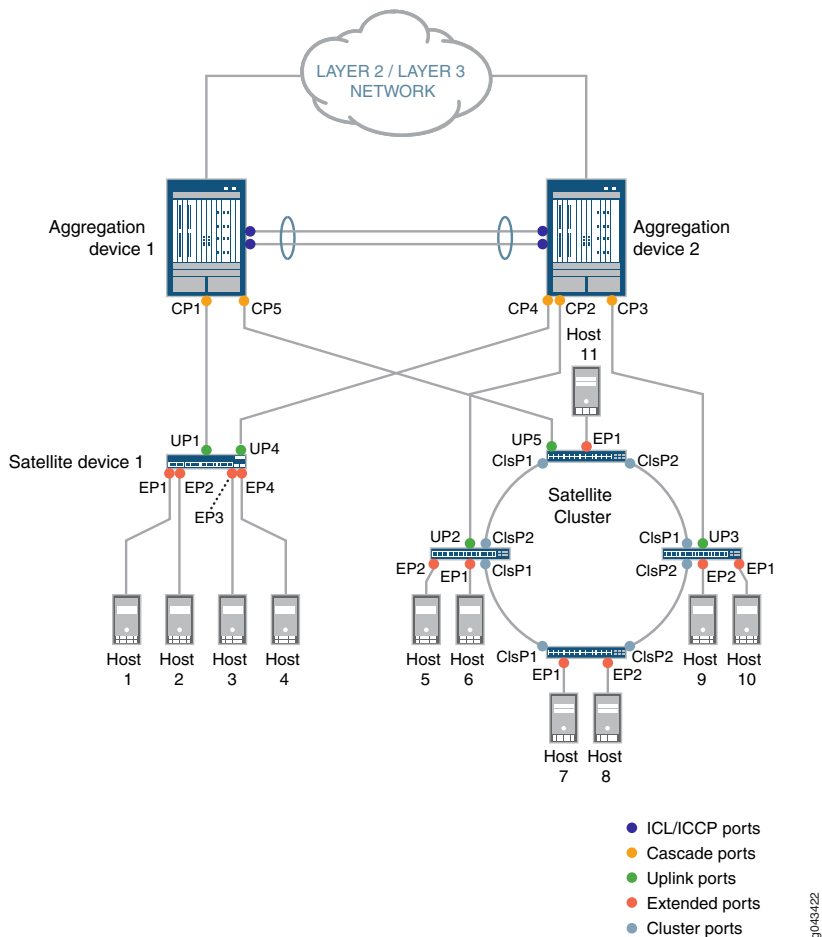


Figure 23 displays a multihome Junos Fusion Enterprise topology with two aggregation devices—Aggregation

device 1 and Aggregation device 2.

Connections—Aggregation device 1 is connected to *Satellite device 1* and a satellite device in the satellite cluster. Satellite device 1 is connected to hosts 1 through 4. Aggregation device 2 is connected to Satellite device 1 and a satellite cluster comprising four satellite devices connected in a ring topology. In a satellite cluster, it is not necessary that all satellite devices are directly connected to an aggregation device. In this topology, we have two satellite devices in the cluster that are directly connected to Aggregation device 2 and one satellite device that is directly connected to Aggregation device 1. The satellite devices, in turn, are connected to each other using cluster ports—cluster port 1 (ClSP1) and cluster port 2 (ClSP2) of each device connects to similar ports on the neighboring satellite devices. One of the satellite device is connected to Aggregation Device 1 through uplink port UP5. The satellite devices in the cluster are connected to hosts 5 through 11.

Port Usage—Both the aggregation devices are connected using ICL and ICCP ports on both ends. ICL ports are used to forward data traffic, whereas ICCP ports are used to exchange control information. You can also choose to use a single link for both ICL and ICCP traffic.

Aggregation devices use cascade ports (CP1, CP2, CP3, CP4, and CP5) to connect to the satellite devices. Satellite device 1 uses uplink port UP1 to connect to the Aggregation Device 1 and UP4 to connect to Aggregation device 2. The satellite cluster uses uplink ports UP2 and UP3 to connect to Aggregation device 2 and UP5 to connect to Aggregation device 1. The satellite devices in the cluster are connected to each other using cluster ports ClSP1 and ClSP2. The satellite devices use extended ports (EP1 through EP11) to connect to the hosts.

To set up a Junos Fusion Enterprise similar to the topology shown in [Figure 23](#) using Network Director, you must perform the following tasks:

1. Create a configuration template. While creating the template, you specify that this is a Junos Fusion for a multihome enterprise fabric; plan the chassis for the aggregation device; identify the ICL, ICCP, and cascade ports on the aggregation device and the cluster ports on the satellite devices.
2. Apply the template to a Junos Fusion fabric. The device to which you apply the template should:
 - Support Junos Fusion. For the list of supported devices see, [Juniper Networks Devices Supported by Junos Space Network Management Platform](#).
 - Have pre-configured with Multichassis Link Aggregation Groups (MC-LAGs) configuration (as MAC-LAG provides multihome support) and should be managed by Network Director.
 - Have the required DMI schema uploaded. For the list of supported DMI schema, see [Uploading DMI Schemas](#)

You can apply the template to a device that is already managed by Network Director or to a new unmanaged device, to make the device the aggregation device for the fusion fabric. When you apply the configuration template to one or more unmanaged devices, you specify the software image that must be installed on the aggregation devices and the Zero Touch Provisioning details for the aggregation devices. Whereas, when you apply the template to a managed device (device that is already managed by Network Director), you select the device that you want to convert to an aggregation device in your

fusion fabric and manually refresh the device topology. To refresh the topology, navigate to Topology View and click **Refresh Topology** under the Tasks menu. Network Director converts the device to a fusion fabric aggregation device and deploys all the necessary configurations on the device.

After the template is applied successfully to a Junos Fusion fabric, there are quite a few tasks that Network Director performs internally that makes building your Junos Fusion fabric simple and error-free. Network Director converts the device that you specified in the Apply Template workflow to an aggregation device and applies the port settings on the various ports that you specified in the configuration template. When an EX4300 device is connected to one of the configured cascade ports, a *link up* event is triggered. The link up event initiates a syslog message to Network Director and Network Director initiates a *topology refresh* job. Network Director then installs the appropriate satellite software on the satellite device and performs the necessary configurations. If a second satellite device is connected to the first satellite device to form a satellite cluster, or another satellite device is connected to the aggregation device, another link up event is triggered and the same steps are repeated. This process continues for all additional satellite devices that are connected to the aggregation device.

Network Director lists the fusion fabric in the Manage Fusion Fabrics page. You can see the details and status of the fabric in the Manage Fusion Fabrics page. You can also edit the fusion fabric, download the cabling plan, and view the fabric connectivity using the Manage Fusion Fabrics page.

RELATED DOCUMENTATION

[Understanding Junos Fusion | 707](#)

[Media Access Control Security Overview | 637](#)

Understanding Junos Fusion Data Center

Junos Fusion Data Center provides automated network configuration and enhanced scalability for medium to large data centers with the Juniper Networks QFX10002-36Q switches, QFX10002-72Q switches, QFX5100 switches, and EX4300 switches.

Figure 24 displays the typical topology of a Junos Fusion in a data center.

Figure 24: Junos Fusion Data Center Topology




Figure 24 displays a multi-home Junos Fusion Data Center topology with two aggregation devices—*Aggregation Device 1* and *Aggregation Device 2*.

Connections—*Aggregation Device 1* is connected to *Satellite Device 1* and *Aggregation Device 2* is connected to *Satellite Device 2*. *Satellite Device 1* and *Satellite Device 2* are in turn connected to server through the extended ports. *Satellite Devices 1* and *2* need to be connected to both aggregation devices for redundancy. For the list of supported aggregation and satellite devices in a Junos Fusion Data Center, see “[Understanding Junos Fusion](#)” on page 707.

Port Usage—Both the aggregation devices are connected using ICL ports with LAGs on both ends. ICL ports are used to forward data traffic. You can choose to use a single link for ICL traffic.

Aggregation devices use cascade ports to connect to the satellite devices.

Satellite devices use uplink ports to connect to the aggregation devices and extended ports to connect to the servers.

To set up a Junos Fusion Data Center similar to the topology that is shown in [Figure 24](#) using Network Director, perform the following tasks:

- a. Create a configuration template. While creating the template, you need to:
 - Specify that this is a Junos Fusion for a multi-home data center fabric.
 - Identify the ICL and cascade ports on the aggregation device.
 - Specify the default satellite image for the satellite devices.
 - Upload the satellite image in the Network Directory image repository as a pre-requisite to create a template.
- b. Apply the template to a Junos Fusion fabric. The device to which you apply the template should:
 - Support Junos Fusion. For the list of supported devices see, [Juniper Networks Devices Supported by Junos Space Network Management Platform](#).
 - Have pre-configured with Multichassis Link Aggregation Groups (MC-LAGs) configuration (as MAC-LAG provides multihome support) and should be managed by Network Director.
 - Have the required DMI schema uploaded. For the list of supported DMI schema, see [Uploading DMI Schemas](#)

You can apply the template to a device that is already managed by Network Director or to a new unmanaged device, to make the device the aggregation device for the fusion fabric. When you apply the configuration template to one or more unmanaged devices, you specify the software image that is to be installed on the aggregation devices and the zero touch provisioning details for the aggregation devices. Whereas, when you apply the template to a managed (device that is already managed by Network Director) device, you select the device that you want to convert to an aggregation device in your fusion fabric and manually refresh the device topology. To refresh the topology, navigate to Topology View and click **Refresh Topology** under Tasks menu. Network Director converts the device to a fusion fabric aggregation device and deploys all the necessary configurations on the device. The devices that you want to convert to as aggregation devices should run the Junos software version that is supported by Junos Fusion. For the list of supported software version, see . In the above topology, you must discover both the *Aggregation Device 1* and *Aggregation Device 2* and manage them with the MC-LAG configuration between these aggregation devices.

After the template is applied successfully, Network Director lists the fusion fabric in the Manage Fusion Fabric page. You can see the details and status of the fabric in the Manage Fusion Fabrics page. You can also edit the fusion fabric, download the cabling plan, and view the fabric connectivity from the Manage Fusion Fabrics page.

RELATED DOCUMENTATION

| [Understanding Junos Fusion](#) | 707

Software Requirements for Junos Fusion

An aggregation device in a Junos Fusion always runs Junos OS software and is responsible for almost all management tasks, including configuring all network-facing ports—the *extended ports*—on all satellite devices in the Junos Fusion. The extended ports in a Junos Fusion, therefore, support all features that are supported by the Junos OS running on the aggregation device.

An aggregation device in a Junos Fusion runs the same Junos OS regardless of whether it is or is not part of a Junos Fusion. Hence, Junos OS is acquired, installed, and managed on an aggregation device in a Junos Fusion in the same manner that it is acquired, installed, and managed on a standalone device that is not part of a Junos Fusion.

The satellite devices in a Junos Fusion run satellite software that has the built-in intelligence to extend the feature set on Junos OS to the satellite device. The satellite software is a Linux-based operating system that enables the satellite devices to communicate with the aggregation device for control plane data while also passing network traffic. Satellite software is also known as satellite network operating system software. Make sure that the devices that are to be converted as satellite devices run the requisite Junos OS version that Junos Fusion supports. For more details on supported Junos OS version, see *Network Director Release Notes, Release 3.2*.

All satellite devices in a Junos Fusion must run the satellite software. The satellite software, notably, applies the feature set on Junos OS to the aggregation device to the satellite device. The satellite software enables the satellite device to participate in the Junos Fusion, but does not provide any other software features for the satellite device.

You can run the same version of satellite software on satellite devices that are different hardware platforms. For instance, if your Junos Fusion included EX4300 and QFX5100 switches as satellite devices, the EX4300 and QFX5100 switches acting as satellite devices can install the satellite software from the same satellite software package.

RELATED DOCUMENTATION

[Understanding Junos Fusion](#) | 707

Creating and Managing Fusion Configuration Templates

IN THIS SECTION

- [Create a Configuration Template for Junos Fusion Enterprise | 717](#)
- [Create a Configuration Template for Junos Fusion Data Center | 722](#)
- [Clone a Configuration Template | 726](#)
- [Apply Configuration Template to Devices | 726](#)
- [View Details about a Configuration Template | 731](#)
- [Delete a Configuration Template | 732](#)

Large campus and data center networks might have many similar network topology instances. These instances can be Virtual Chassis Fabrics (VCFs), Layer 3 Fabrics, or Junos Fusion Enterprise, or Junos Fusion Data Center. Creating each of these instances afresh can be tedious, time-consuming, and error-prone.

Network Director enables you to create a configuration template for a Junos Fusion Fabric and reuse it in all similar instances. A configuration template combines the common settings that apply to an instance, such as the chassis details, ports to be used as cascade and cluster ports, software image for the satellite devices, and so on. However, you might still specify some additional configuration attributes while applying the template to a fabric or Junos Fusion setup.

Before you create a configuration template:

- Understand the software requirements for the aggregation devices and the satellite devices. For more information, see [“Software Requirements for Junos Fusion” on page 715](#).
- Ensure that the software images for the aggregation devices and the satellite devices are uploaded to Network Director using the **Image Management > Manage Image Repository** in the Deploy mode.


You can perform the following tasks from the Manage Fusion Configuration Templates page.

Create a Configuration Template for Junos Fusion Enterprise

To create a Junos Fusion Enterprise configuration template:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View**, or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.
3. From the Tasks menu, select **Network Builder > Manage Fusion Config. Templates**.
The Manage Fusion Configuration Templates page opens.
4. Click **Create**.
The Create Fusion Configuration Template wizard opens.
5. In the Template Type page of the Create Fusion Configuration Template wizard, select the type of deployment for which you want to create the template. Select **Campus** to create a template for Junos Fusion Enterprise.
6. Select one of the following depending on the type of Junos Fusion topology:
 - **Fusion Enterprise Single Aggregation Device**—Indicates that the one or more satellite devices are connected to a single aggregation device.
 - **Fusion Enterprise Multiple Aggregation Devices**—Indicates that each satellite device is connected to two aggregation devices forming an MC-LAG cluster at the aggregation layer.
7. For a Junos Fusion Enterprise, select one of the following depending on the type of Junos Fusion topology:
 - **Fusion Enterprise Single Aggregation Device**—Indicates that the one or more satellite devices are connected to a single aggregation device.
 - **Fusion Enterprise Multiple Aggregation Devices**—Indicates that each satellite device is connected to two aggregation devices forming an MC-LAG cluster at the aggregation layer.

For a Junos Fusion Data Center, select **Fusion Data Center Multiple Aggregation Devices**.
8. Click + in the Available Satellite Images box to ensure that the appropriate satellite software image is available.

Network Director lists the satellite software image only if you have uploaded the software image in Network Director Image Repository.

9. Click **Next**.

The Settings page opens.

10. Enter a name for the configuration template.

If you chose to configure with multiple aggregation devices, the Settings page displays two tabs—*Aggregation Device 1* and *Aggregation Device 2*.

11. In the Aggregation Device 1 tab, select the device model that you want to use as the aggregation device.

Network Director supports the following as aggregation devices:

- EX9204, EX9208, and EX9214 in a Junos Fusion Enterprise setup.

NOTE: In a Fusion Enterprise Multiple Aggregation Device topology, when you select a device as the first aggregation device, Network Director considers the second aggregation device also to be of the same device model.

12. When you select a device model in a Junos Fusion Enterprise setup, you need to create a new chassis for the aggregation device model that you selected.

To create a new chassis:

a. Click **Build New Chassis**.

The Build Chassis window opens. The Build Chassis window has two panes—the *Available line cards* pane and the *Chassis: FPC slots* pane.

The Available line cards pane lists the all the EX9200 line card models and the Chassis: FPC slots pane lists the available FPC slots on the device.

NOTE: Network Director does not allow you to import chassis details. You must create a chassis for each configuration template irrespective of whether the template is for a single-home or a multihome Junos Fusion topology.

b. Drag and drop the line cards that you want to add to the chassis from the Available line cards pane to the appropriate FPC slots in the Chassis: FPC slots pane.

From all the EX9200 line cards that are listed in the Available Line Card list, select at least one line card that supports the cascade port. If you do not choose any line card that is supported by cascade port, Network Director returns an error message that prompts you to select a line card that is supported by cascade port.

- c. Click **Set** after you have added all the required line cards to the FPC slots.

The Build Chassis window closes.

After you setup the chassis, the line cards that support the cascade ports are listed in the Select Cascade Ports window. All the other line cards are filtered out. For the list of EX9200 line cards that support cascade ports see, [Line Cards on EX9200 Switch Cascade Port Support](#).

- d. Mouse over **Preview** to preview the chassis with the line cards that you added.

13. A cascade port is a port on an aggregation device that sends and receives control and network traffic from an attached satellite device. Click **Select Cascade Ports** and perform the following steps to select cascade ports for the Junos Fusion Enterprise:

- a. To select ports individually from the ports list, click **Identify Port > By List** or to select ports by specifying a range, click **Identify Port > By Range**. The Select Ports window opens.

NOTE: In the Select Ports window, Network Director displays only the Gigabit Ethernet interfaces, 10-Gigabit Ethernet interfaces, and the 40-Gigabit Ethernet interfaces as these are the recommended interfaces for a cascade port.

- b. If you want to select ports individually, select the ports that you want to use as cascade ports on each FPC. Use the left navigation pane to view and select the ports on each FPC slot.
- c. If you want to select ports using a port range, specify the starting and ending port ID.

- d. Click **Add**.

Network Director adds the selected ports or the ports that are available on the FPC from the range that you specified, to the list of cascade ports in the Create Template page.

- e. Select a cascade port and optionally, you can enter the corresponding FPC slot ID of the satellite device.

Specify the FPC slot ID in the range 65-254.

NOTE: When you add a satellite-enabled device in a Junos Fusion setup, Network Director checks for the FPC slot ID that you specify and uses this ID while Network Director configures the satellite devices. If you have not provided an FPC slot ID, Network Director automatically generates the FPC slot ID for satellite device from the allowed range and provisions the same on the aggregation device.

14. For a Fusion Enterprise Multiple Aggregation Device topology, you must select the port type as an **ICL Port** to configure an Inter-Chassis Link (ICL port). However, it is optional to select the **ICCP PORT** as the port type to configure an Inter-Chassis Control Protocol (ICCP) port. Select the **ICCP PORT** as the port type, to manually configure the ICCP port. If you do not select the **ICCP PORT** as the port type, Junos platform automatically assigns one of the ports as an ICCP port and deploys all the necessary configurations to the port. At the time of applying the template to one or more aggregation devices, the ICCP configuration is pushed to the aggregation devices.

NOTE: Automatic ICCP provisioning is enabled by default and if you manually configure an ICCP parameter that is normally set by default, your configuration automatically overrides the default parameter. If you decide to configure ICCP, you must configure matching configurations on both the aggregation devices.

An ICL port is used to forward data traffic across the aggregation devices. This link provides redundancy when a link failure occurs in one of the active links.

An ICCP port is used to exchange the control information between two aggregation devices to ensure that data traffic is forwarded properly.

Click **Select ICL & ICCP Ports** and perform the following steps to select the ports:

- a. To select ports individually from the ports list, click **Identify Port > By List** or to select ports by specifying a range, click **Identify Port > By Range**. The Select Ports window opens.
- b. If you want to select ports individually, select the ports that you want to use as ICL ports on each FPC. Use the left navigation pane to view and select the ports on each FPC slot.
- c. If you want to select ports using a port range, specify the starting and ending port ID.
- d. Click **Add**. Network Director adds the selected ports or the ports that are available on the FPC from the range that you specified, to the list of ICL ports in the Create Template page.
- e. For a Junos Fusion Enterprise, specify the port type as **ICL port** or **ICCP port** depending on the port usage.

In Select ICL & ICCP window, you can select more than one ICL and ICCP port as a LAG interface. The LAG interface provides redundancy and load balancing between the two aggregation devices.

When you select an ICL port as the LAG interface, there is no additional configuration required for the LAG on the selected ICL port. Network Director generates a default configuration for the aggregation devices. This configuration is pushed to the aggregation device at the time of link up event when the connection between the aggregation devices on the ICL interfaces are made.

When you select an ICCP port as the LAG interface, you must specify the local IP address of the ICCP port. For information on specifying the ICCP local address, see [“Apply Configuration Template to Devices” on page 726](#). If you do not select the ICCP port, Network Director automatically configures the ICCP port to establish ICCP session between the connected aggregated devices.

NOTE: If ICL or ICCP physical connections are not made between aggregation devices when you apply the template to aggregation devices, Network Director displays the each aggregation device as separate single home Junos Fusion fabrics. When the ICL and ICCP connections between aggregation devices are established, a link up is triggered. Network Director deletes the device, rediscovers the device, and converts the Junos Fusion to a multihome fabric.

15. If the Junos Fusion Enterprise uses a satellite cluster, you must select the ports on satellite devices that will act as cluster ports. In a satellite cluster, up to 10 satellite devices can be interconnected using the standard 10-Gigabit Ethernet or 40-Gigabit Ethernet interfaces to form a cluster, which in turn can be connected to the aggregation devices over a pair of fiber uplinks. In a cluster, all the satellite devices do not need to be directly connected to the aggregation device. One or two satellite devices in a satellite cluster connects to the aggregation device through cluster ports.

Click **Select Satellite Cluster Ports** to select the cluster ports that the satellite devices use to connect to other satellite devices in a satellite cluster. Perform the following steps to add cluster ports:

NOTE: If you have multiple satellite devices that connect to one or more aggregation devices, the ports that you select in is applied to all the directly connected satellite devices that are part of the cluster.

- a. To select ports individually from the ports list, click **Identify Port > By List** or to select ports by specifying a range, click **Identify Port > By Range**. The Select Ports window opens.
- b. Select the device model of the satellite device.
- c. If you want to select ports individually, select the ports that you want to use as cluster ports on the selected satellite device.

- d. If you want to select ports using a port range, specify the starting and ending port ID.
 - e. Click **Add**. Network Director adds the selected ports that you specified or the ports that are available on the device from the range that you specified, to the list of cluster ports in the Create Template page.
16. Select the software image that you want to install on all the satellite devices from the Satellite Image list. The supported satellite image version is **satellite-3.1R1.3**. For more information about software images in Junos Fusion, see [“Software Requirements for Junos Fusion” on page 715](#).
 17. For a Fusion Enterprise Multiple Aggregation Device topology, click the Aggregation Device 2 tab.
Network Director copies the same settings that you specified for the first aggregation device for the second aggregation device. Review the settings.


If you want to modify the setting including the device model, click **Edit**. Select a device model and follow steps Step 7 through Step 15 to specify details for the second aggregation device.
 18. Click **Next** to review the template details.
 19. After you have reviewed the details, click **Finish**. Network Director creates the template and displays it in the Manage Fusion Configuration Templates page.

Create a Configuration Template for Junos Fusion Data Center

To create a Junos Fusion Data Center configuration template:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View**, or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.
3. From the Tasks menu, select **Network Builder > Manage Fusion Config. Templates**.
The Manage Fusion Configuration Templates page opens.
4. Click **Create**.
The Create Fusion Configuration Template wizard opens.

5. In the Template Type page of the Create Fusion Configuration Template wizard, select **Data Center** to create Junos Fusion Data Center config template..
6. Select **Fusion Data Center Multiple Aggregation Devices** as the Junos Fusion topology. The multiple aggregation topology indicates that each satellite device is connected to two aggregation devices forming an MC-LAG cluster at the aggregation layer.
7. Click + in the Available Satellite Images box to ensure that the appropriate satellite software image is available.

Network Director lists the satellite software image only if you have uploaded the software image in the Network Director Image Repository.

8. Click **Next**.

The Settings page opens.

9. Enter a name for the configuration template.

If you chose to configure with multiple aggregation devices, the Settings page displays two tabs—*Aggregation Device 1* and *Aggregation Device 2*.

10. In the Aggregation Device 1 tab, select the device model that you want to use as the aggregation device.

Network Director supports the following as aggregation devices:

- QFX10002-36Q and QFX10002-72Q in a Junos Fusion Data Center setup.

NOTE: In a Fusion Data Center Multiple Aggregation Device topology, when you select a device as the first aggregation device, Network Director interprets the second aggregation device also to be of the same device model.

11. Click **Select Cascade Ports** and perform the following steps to select cascade ports for the Junos Fusion Data Center:

- a. To select ports individually from the ports list, click **Identify Port > By List** or to select ports by specifying a range, click **Identify Port > By Range**.

The Select Ports window appears.

NOTE: In the Select Ports window, Network Director displays only the Gigabit Ethernet interfaces, 10-Gigabit Ethernet interfaces, and the 40-Gigabit Ethernet interfaces as these are the recommended interfaces for a cascade port.

- b. If you want to select ports individually, select the ports that you want to use as cascade ports on each FPC. Use the left navigation pane to view and select the ports on each FPC slot.
- c. If you want to select ports using a port range, specify the starting and ending port ID.
- d. Click **Add**.

Network Director adds the selected ports or the ports that are available on the FPC from the range that you specified, to the list of cascade ports in the Create Template page.

- e. Optionally, you can enter the FPC slot ID for the satellite device in the **Satellite Slot Id** for the selected cascade port.

Specify the FPC slot ID in the range 65-254.

NOTE: When you add a satellite-enabled device in a Junos Fusion setup, Network Director checks for the FPC slot ID that you specify and uses this ID while Network Director configures the satellite devices. If you have not provided an FPC slot ID, Network Director automatically generates the FPC slot ID for satellite device from the allowed range and provisions the same on the aggregation device.

12. For a Fusion Data Center Multiple Aggregation Device topology, you must specify the ports on the aggregation device that you want to use as the ICL port. However, for the ICCP, Junos Space platform automatically assigns one of the ports as an ICCP port and deploys all the necessary configurations to the port (when the aggregation devices are connected).

An ICL port is used to forward data traffic across the aggregation devices. This link provides redundancy when a link failure occurs in one of the active links.

An ICCP port is used to exchange the control information between two aggregation devices to ensure that data traffic is forwarded properly.

Click **Select ICL Ports** and perform the following steps to select the ports:

- a. To select ports individually from the ports list, click **Identify Port > By List** or to select ports by specifying a range, click **Identify Port > By Range**. The Select Ports window opens.
- b. If you want to select ports individually, select the ports that you want to use as ICL ports on each FPC. Use the left navigation pane to view and select the ports on each FPC slot.
- c. If you want to select ports using a port range, specify the starting and ending port ID.
- d. Click **Add**. Network Director adds the selected ports or the ports that are available on the FPC from the range that you specified, to the list of ICL ports in the Create Template page.

NOTE: The port type is selected as the ICL port.

- e. In Select ICL Ports window, you can select more than one ICL port as a LAG interface. The LAG interface provides redundancy and load balancing between the two aggregation devices.

When you select an ICL port as the LAG interface, there is no additional configuration required for the LAG on the selected ICL port. A default configuration is generated by Network Director for the aggregation devices. This configuration is pushed to the aggregation device at the time of link-up event when the connection between the aggregation devices on the ICL interfaces are made.

13. Select the software image that you want to install on all the satellite devices from the Satellite Image list. For more information about software images in Junos Fusion, see [“Software Requirements for Junos Fusion” on page 715](#). For a list of supported satellite image versions, see *Network Director 3.2 Quick Start Guide*.

14. Network Director copies the same settings that you specified for the first aggregation device for the second aggregation device. Review the settings.

If you want to modify the setting for the second aggregation device, click **Aggregation Device 2** and click **Edit**. Select a device model and follow steps Step 11 through Step 12 to specify details for the second aggregation device.

15. Click **Next** to review the template details.

16. After you have reviewed the details, click **Finish**.

Network Director creates the template and displays it in the Manage Fusion Configuration Templates page.

Clone a Configuration Template

You can make a copy of an existing template by cloning the template. When you clone a template, Network Director copies all the settings to the cloned template. However, you can modify the settings in the cloned template based on your requirements.

To clone a Junos Fusion Enterprise or Junos Fusion Data Center template:

1. Select the template that you want to clone and click **Clone**.

The Clone Fusion Configuration Template page opens.

2. Network Director copies all the settings in the original template to the cloned template. However, you can modify all the settings in the cloned template as per your requirement. Follow step 6 through 16 in the [“Create a Configuration Template for Junos Fusion Enterprise” on page 717](#) for instructions on specifying settings for the configuration template.

Apply Configuration Template to Devices

After you have created a configuration template, you can apply the template to aggregation devices. You can apply the template to aggregation devices that are already managed by Network Director or new aggregation devices.

If you plan to apply the template to a managed device, make sure that the device runs the Junos OS software image that is supported on Junos Fusion systems. For detailed steps on installing or upgrading software image using Network Director, see [“Managing Software Images” on page 1235](#). See Network Director release notes to know the Junos OS software version that is supported for Junos Fusion systems.

To apply a configuration template to Junos Fusion Enterprise or Junos Fusion Data Center devices:

1. From the Manage Fusion Configuration Template page, select the template that you want to apply to an aggregation device.
2. Do one of the following:
 - If the aggregation device is not managed by Network Director, click **Apply** and select **Unmanaged Devices**. The Apply Fusion Configuration Template page opens.
 - If the aggregation device is already managed by Network Director, but is not part of a fusion fabric, click **Apply** and select **Managed Devices**. The Apply Fusion Configuration Template page opens.

Skip to Step 4 and follow the instructions to apply the template to managed devices.

3. If you selected to apply the template to Unmanaged devices, perform the following steps to specify details about the aggregation devices:
 - a. In the Apply Fusion Configuration Template page, specify the DHCP server settings by following the descriptions given in [Table 154](#).

Table 154: DHCP Server Settings

Field	Description
DHCP Server	IP address or the hostname of the DHCP server.
DHCP Server Type	<p>The type of DHCP server that provides the necessary information to the aggregation devices. You can choose to use a CentOS DHCP server, an Ubuntu DHCP server, or any other DHCP server.</p> <p>NOTE: If you select Other, you must configure the DHCP server settings manually.</p>
Manually Configure Server	<p>Select to indicate that you want to manually configure the DHCP server. You can configure the CentOS and Ubuntu DHCP servers manually or from Network Director.</p> <p>If you want to use any other type of DHCP server, do the following:</p> <ol style="list-style-type: none"> a. Select the Manually Configure Server check box. Network Director hides all the other details except the DHCP Server Type. b. Follow the instructions displayed in this box to configure the DHCP server manually.
DHCP User	Username to log in to the DHCP server.
DHCP Password	Password for the specified username.
Confirm Password	Confirm the DHCP server password.

- b. Specify the File server settings described in [Table 155](#).

Table 155: File Server Settings

Field	Description
File Server Type	The type of file server where the software image to be installed on the aggregation device is to be stored. You can choose to use an FTP, HTTP, or an TFTP file server.
File Server	IP address or hostname of the file server.
File Server Root Directory	The root directory of the file server.

- c. Select the software image that you want to install on the aggregation device from AD Image.

NOTE: Ensure that the software image is uploaded to Network Director using the **Image Management > Manage Image Repository** in the Deploy mode. If the software image is not uploaded, Network Director does not display the software image in this field.

- d. ZTP process maps the management interface MAC address or the device chassis serial number of each aggregation device to the software image, IP address, hostname, and the configuration file stored on the file server. This mapping is stored in the DHCP server. When an aggregation device starts up, the device contacts the DHCP server to obtain the IP address and the software image location. The DHCP server looks up in its MAC address or serial number mapping database to identify the device and provide details about the file server that the device must contact to get the software image and configuration file. The device uses this information to contact the file server and obtain the software image and the configuration file for deploying on the device.

Click **Add** to add a row to the Device Details table. You specify the aggregation device details in the Device Details table. You can enter the device details manually or you can specify the details in a CSV file and import it.

- e. Do the following to import the device details from a CSV file:


NOTE: When you use the Import option, you must specify either the MAC address or the Serial number, but not both.

1. Click **Import > By Mgmt MAC addresses** to import the MAC addresses of aggregation devices in CSV format. You must enter the MAC addresses in the specified format. Click **CSV Sample** to download a sample CSV file that you can use to import MAC addresses.
 2. Click **Import > By Device Serial Numbers** to import the serial numbers of aggregation devices in CSV format. You must enter the serial numbers in the specified format. Click **CSV Sample** to download a sample CSV file that you can use to import serial numbers.
- f. If you want to specify a password for all the aggregation devices that you add to this fusion fabric, click **Set Default Password**. The Set Default Password window opens.

Enter the password, confirm the password and click **Set** to set the password as the default password for all the aggregation devices that you add to this fusion fabric.

You can skip this step if you wish to specify the password individually or if you want to use a different password for each of your aggregation device.

- g. Do the following to specify details about the aggregation device manually:

1. Network Director assigns a host name for the aggregation device. You can modify this name by clicking  .
2. Enter the **IP address of the management interface** or the **MAC address and the device chassis serial number** of the aggregation device.
3. If you are applying the template to a multihome Junos Fusion that has two aggregation devices, enter the ICCP IP address that this device uses to connect to the second aggregation device. If you are applying the template to a single-home Junos Fusion, you can leave this field blank.
4. Enter the MAC address of the management interface (for example, the em0 interface) or the device chassis serial number of the aggregation device in the MAC Addresses or Serial Number field.
5. Click **Import MAC Address** to import the MAC addresses of aggregation devices in CSV format. You must enter the MAC addresses in the specified format. Click **CSV Sample** to download a sample CSV file that you can use to import MAC addresses.

NOTE: When you use the Import option, you must specify either the MAC address or the Serial number, but not both.

6. Click **Import Serial Number** to import the serial numbers of the aggregation device in CSV format. You must enter the MAC addresses in the specified format. Click **CSV Sample** to download a sample CSV file that you can use to import serial numbers.
- h. Type the local IP address for the ICCP port. The IP address is used to communicate to the peers that hosts an ICCP port as a LAG interface.

NOTE: This field is available only when you choose to configure the ICCP port manually for a Junos Fusion Enterprise device.

- i. Click the text link in the Password column to set or modify the password for the aggregation device.
 - If you have not configured a default password for the aggregation device, then click **Set Password**. The Set Custom Password window opens. Specify the password, confirm the password, and click **Set**.
 - If you want to override the default password and specify a custom password for the aggregation device, then click **Edit**. The Edit Custom Password window opens. Select **Use Default Password** if you want to use the default password.

Select **Edit Custom** if you want to specify a different password for the aggregation device. Specify the password, confirm the password, and click **Set**.

- j. To view the configuration that Network Director deploys on the aggregation device, click **View** in the Config field.
 - l. Repeat above steps, Step d through Step j to add another aggregation device.
4. If you selected to apply the template to Managed devices, perform these steps:
- a. Select the devices that you want to add as aggregation device from the Select Managed Devices table.
 - b. Click on the link in the Backup Config column.
Network Director creates a ZIP archive file containing the existing configuration of the device. You save a copy of this onto your local machine.

NOTE: Before you perform this step, make sure that the device is running a Junos OS software image that is supported on Junos Fusion systems. If not, you must upgrade the software image before you proceed. For detailed steps on installing or upgrading software image using Network Director, see [“Managing Software Images” on page 1235](#).

5. Click **Apply**.

Network Director converts the device to a fusion fabric aggregation device and deploys all the necessary configurations on the device and opens the Manage Fusion Fabrics page.

Manage Fusion Fabrics page displays the status and details of all fusion fabrics that are created using Network Director.

After the template is applied successfully to a Junos Fusion fabric, there are quite a few tasks that Network Director performs internally that makes building your Junos Fusion fabric simple and error-free. Network Director converts the device that you specified in the Apply Template workflow to an aggregation device and applies the port settings on the various ports that you specified in the configuration template. When an EX4300 device is connected to one of the configured cascade ports, a *link up* event is triggered. The link up event initiates a syslog message to Network Director and Network Director initiates a *topology refresh* job. Network Director then installs the appropriate satellite software on the satellite device and performs the necessary configurations. If a second satellite device is connected to the first satellite device to form a satellite cluster, or another satellite device is connected to the aggregation device, another link up event is triggered and the same steps are repeated. This process continues for all additional satellite devices that are connected to the aggregation device.

NOTE: If ICL or ICCP physical connections are not made between aggregation devices when you apply the template to aggregation devices, Network Director displays the each aggregation device as separate single home Junos Fusion fabrics. When the ICL and ICCP connections between aggregation devices are established, a link up is triggered. Network Director deletes the device, rediscovers the device, and converts the Junos Fusion to a multihome fabric.

For Junos Fusion Enterprise, if you want to form a cluster setup and if the cascade port of the Aggregate device is connected to the base port (1 GE) of EX2300, EX3400, or EX4300 devices, you need to configure the cluster-policy manually by logging into the aggregation device. Execute the following commands to configure the cluster policy:

- If the satellite device is part of a cluster, log in to the CLI of the aggregation device and execute the following commands:

```
user@aggregation-device# set groups policy-number policy-options satellite-policies port-group-alias
port-group-alias-name pic pic-identifier port port-ID1
```

```
user@aggregation-device# set groups policy-number policy-options satellite-policies port-group-alias
port-group-alias-name pic pic-identifier port port-ID2
```

```
user@aggregation-device# set groups policy-number policy-options satellite-policies port-group-alias
port-group-alias-name pic pic-identifier port port-ID3
```

```
user@aggregation-device# set groups policy-number policy-options satellite-policies port-group-alias
port-group-alias-name pic pic-identifier port port-ID4
```

```
user@aggregation-device# set groups policy-number policy-options satellite-policies port-group-alias
port-group-alias-name set groups policy-number policy-options satellite-policies candidate-uplink-port-policy
cluster-policy-name uplink-port-group pic-identifier
```

```
user@aggregation-device#set apply-groups policy-number
```

```
user@aggregation-device# set chassis satellite-management cluster cluster-number cluster-policy
cluster-policy-name
```

View Details about a Configuration Template

To view details about a configuration template, select the template and click **Details**.

The Template Details for <template-name> window opens.

For a Junos Fusion Enterprise setup, the Template Details for <template-name> page displays the template name, deployment type, configuration type, software upgrade group details, and details of ports such as satellite cluster ports, cascade ports, ICL, and ICCP ports.

For a Junos Fusion Data Center setup, the Template Details for <template-name page> displays the template name, deployment type, configuration type, software upgrade group details, and details of ports such as cascade ports and ICL ports.

Delete a Configuration Template

To delete a Junos Fusion Enterprise or Junos Fusion Data Center configuration template, select the template and click **Delete**.

RELATED DOCUMENTATION

[Understanding Junos Fusion | 707](#)

[Managing Fusion Fabrics | 732](#)

Managing Fusion Fabrics

IN THIS SECTION

- [Modify the Fusion Fabric | 734](#)
- [View the Cabling Plan | 736](#)
- [View Fabric Connectivity | 736](#)
- [Replace Aggregation Device or Satellite Device in Junos Fusion | 737](#)

With Junos Fusion technology, network administrators can reduce network complexity and operational expenses by collapsing underlying network elements into a single, logical point of management using QFX Series and EX Series switches running the Junos operating system. You can create and provision Junos Fusion fabrics using Network Director. The Manage Fusion Fabrics page displays the status and details of all fusion fabrics that are created using Network Director.

To open the Manage Fusion Fabrics page:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View**, or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.

3. From the Tasks menu, select **Network Builder > Manage Fusion Fabrics**.

The Manage Fusion Fabrics page opens. [Table 156](#) describes the fields that are displayed in the Manage Fusion Fabrics page.

Table 156: Manage Fusion Fabrics Field Descriptions

Field	Description
Fusion Name	Name of the fusion fabric.
Fusion Type	Indicates the type of fusion fabric as Fusion Enterprise Single Home, Fusion Enterprise Multi Home, Fusion Data Center Single Home, or Fusion Data Center Multi Home.
Aggregation Devices	Hostname and IP address of the aggregation devices.
Summary	Summary of the fusion fabric. This field displays the number of cascade ports, number of satellite devices, and the status of the satellite devices.
Status	Status of the fusion fabric.
Action	Additional tasks that you can perform on the fusion fabric. Available actions are—View Topology, Download Cabling Plan, and View Connectivity.
Last Updated Time	Date and time when the fusion fabric was last updated.

You can perform the following tasks from the Manage Fusion Fabrics page:

Modify the Fusion Fabric

IN THIS SECTION

- [Edit Aggregation Device Details | 734](#)
- [Edit Satellite Device Details | 735](#)
- [Enable Uplink Failure Detection | 736](#)

From the Edit Fusion Fabric page:

1. Select a fabric from the table and click **Edit**.

The Edit Fusion Fabric page opens.

2. Edit the Aggregation device details. For procedure steps, see [“Edit Aggregation Device Details” on page 734](#).
3. Edit the satellite device details. For procedure steps, see [“Edit Satellite Device Details” on page 735](#).
4. Enable uplink failure detection for a satellite device. For procedure steps, see [“Enable Uplink Failure Detection” on page 736](#).
5. Click **Save** to save the changes to the fusion fabric and close the Edit Fusion Fabric page. The Configuration Deployment window opens displaying the status of the job the Network Director initiates to deploy the changes to the fusion fabric.
6. Click **Close** in the Configuration Deployment window to close the window and return to the Manage Fusion Fabrics page.

Edit Aggregation Device Details

To edit the aggregation device details:

1. Open the **Aggregation Devices** tab in the Edit Fusion Fabric page to edit the aggregation device details. You can add cascade ports or delete cascade ports from an aggregation device. To do this, select an aggregation device from the table. Network Director displays the cascade ports that are part of the selected aggregation device in the port details table.
2. To add new cascade ports to the aggregation device, click **Add Ports**. The Select Ports window opens. Click an FPC and select the ports that you want to include. When you have added the ports, click **Add**.

3. To delete a cascade port, select a port from the port details table and click **Remove**.

NOTE: In a multihome Junos Fusion fabric, if you remove a cascade port from one aggregation device, Network Director initiates a two-step commit process on both the aggregation devices. If one of the aggregation device is out-of-sync or is not reachable, the port is not deleted.

Edit Satellite Device Details

To edit the satellite device details:

1. Open the **Satellite Devices** tab in the Edit Fusion Fabric page to edit the satellite device details.

You can perform the following operations on the satellite devices in the Edit Fusion Fabric page:

- change the alias name for the satellite device. See Step 2.
- remove a satellite device. See Step 3.

2. In a multi-home setup, the configuration for a satellite device exists on both the aggregation devices. In such a topology, the alias name of the satellite device should be identical on both the aggregation devices. If the alias names of a satellite device does not match, you need to edit the alias name to be identical on both the aggregation devices. To edit the alias name of a satellite device:

- a. Select the FPC slot number.
- b. Click the alias name in the column **Alias Name**.
- c. Type the new alias name and then click **Save**.

Network Director displays a warning message **Alias name for AD1 and AD2 will be same after the edit**.

- d. Click **OK**.

The alias name of the satellite device is now identical on both the aggregation devices.

3. To remove a satellite device from the fusion fabric, select the device and click **Remove**.

NOTE: Replacing a satellite device that has no MAC address or serial number binding to the FPC slot on the aggregation device is plug-and-play. Replace the satellite device in your Junos Fusion topology and Network Director takes care of all the remaining configurations. This is irrespective of whether the device is a standalone device or a member of a satellite cluster. However, if there is a MAC address or serial number binding between the satellite device and the FPC slot on the aggregation device, you must run some commands from the CLI of the aggregation device to replace the satellite device. For more details, see [“Replace Aggregation Device or Satellite Device in Junos Fusion” on page 737](#).

Enable Uplink Failure Detection

To enable update failure detection:

1. Click the **Settings** tab in the Edit Fusion Fabric page.
2. Select **Uplink Failure Detection** check box corresponding to the device to enable uplink failure detection on a Junos Fusion.

Enabling uplink failure detection on a satellite devices detects link failures on the uplink interfaces used to connect to aggregation devices. When uplink failure detection detects uplink failure on a satellite device, all of the device's extended ports (which connect to host devices) are shut down.

You can also view the ICL and ICCP settings for a multihome Junos Fusion Enterprise fabric and ICL settings for a multihome Junos Fusion Data Center fabric.

View the Cabling Plan

To download the cabling plan for the Junos Fusion fabric, click **Download Cabling Plan** from the **Action** field corresponding to a fabric. Network Director generates and downloads the cabling plan as a PDF file.

View Fabric Connectivity

After you have set up and deployed Junos Fusion devices in Network Director, you can pictorially view the physical connectivity between the various devices in the fabric. This page displays the devices and their physical connectivity in the spine-and-leaf topology.

To view the connectivity between the Junos Fusion devices:

1. Do one of the following:
 - From the Manage Fusions page, click **View Topology** in the **Actions** field.

- While in the Logical, Location, Device, or Custom View, select the Junos Fusion fabric for which you want to view the connectivity details from the View pane and click **Connectivity > View Fabric Connectivity** from the Tasks pane.

The Fusion Connectivity page opens, displaying the connectivity between the aggregation devices and the satellite devices in the selected fusion fabric.

NOTE: If you change the position and arrangement of Junos Fusion devices in a topology and navigate to some other page in the user interface, Network Director preserves the position of the devices in the topology when you return to the topology.

2. From the Fabric Connectivity page, you can:

- Click **Device** and select **Provisioned SDs** to view all the satellite devices that are provisioned as part of the Junos Fusion system.
- Select Color Code Port Utilization to view the color coded port utilization in the graphical view by clicking **Links**.
- Mouse over each entity to view a window displaying the details about that entity and View Device Connectivity link. Clicking on View Device Connectivity link displays Device Connectivity page. For more information about this page, see [“Viewing Device Connectivity” on page 1139](#).
- Zoom in or zoom out of the connectivity view by using the + and - buttons.

Replace Aggregation Device or Satellite Device in Junos Fusion

Network Director enables you to replace faulty or non-responsive aggregation devices in your Junos Fusion fabric by using the Replace functionality. You can replace an aggregation device in a dual home Junos Fusion by selecting the device that you want to remove and by specifying the serial number or the MAC address of the replacement device.

Replacing a satellite device that has no MAC address or serial number binding to the FPC slot on the aggregation device is plug-and-play. Replace the satellite device in your Junos Fusion topology and Network Director takes care of all the remaining configurations. This is irrespective of whether the device is a standalone device or a member of a satellite cluster. However, if there is a MAC address or serial number binding between the satellite device and the FPC slot on the aggregation device, you must run some commands from the CLI of the aggregation device to replace the satellite device.

This topic describes the steps that you must perform to replace an aggregation device or a satellite device from the a Junos Fusion fabric.

To replace aggregation device or satellite device:

1. From the Manage Fusions page, select a fusion fabric for which you want to replace an aggregation and click **Replace**. The Replace Fusion Fabric page opens.
2. To replace an aggregation device, select the aggregation device that you want to replace and click **Replace**.

Enter the MAC address or the serial number of the new aggregation device.

3. Do one of the following to replace a satellite device that has a MAC address or serial number binding to the FPC slot of the aggregation device:

- If the satellite device is a standalone device, log in to the CLI of the aggregation device and execute the following commands:

```
user@aggregation-device# set chassis satellite-management fpc fpc-id system-id
MAC-address-of-the-device | Serial-number-of-the-device
user@aggregation-device# Commit
```

- If the satellite device is part of a cluster, log in to the CLI of the aggregation device and execute the following commands:

```
user@aggregation-device# set chassis satellite-management cluster cluster-name fpc fpc-id alias alias
user@aggregation-device# set chassis satellite-management cluster cluster-name fpc fpc-id description
10.204.248.62-member0
user@aggregation-device# set chassis satellite-management cluster cluster-name fpc fpc-id member-id
member-id
user@aggregation-device# set chassis satellite-management cluster cluster-name fpc fpc-id system-id
MAC-address-of-the-device | Serial-number-of-the-device
user@aggregation-device# Commit
```

RELATED DOCUMENTATION

[Understanding Junos Fusion | 707](#)

[Creating and Managing Fusion Configuration Templates | 716](#)

Creating and Managing Satellite Software Upgrade Groups

IN THIS SECTION

- [Create a Software Upgrade Group | 740](#)
- [Edit a Software Upgrade Group | 740](#)
- [View Details of a Software Upgrade Group | 741](#)
- [Delete a Software Upgrade Group | 741](#)

A satellite software upgrade group is a group of satellite devices that are designated to upgrade to the same satellite software version using the same satellite software package. One Junos Fusion can contain multiple software upgrade groups, and multiple software upgrade groups must be configured in most Junos Fusions to avoid network downtimes during satellite software installations.

In Network Director, you select a fusion fabric and create a software upgrade group. The software upgrade group can contain one or more satellite devices. When a satellite device is added to a Junos Fusion, the aggregation device checks whether the satellite device or the FPC ID that is used by the satellite device is included in the satellite software upgrade group that is assigned to the Fusion system. If it is, the device—unless it is already running the same version of satellite software—upgrades its satellite software using the satellite software associated with the satellite software upgrade group.

When the satellite software package associated with an existing satellite software group is changed, the satellite software for all member satellite devices is upgraded using a throttled upgrade. The throttled upgrade ensures that only a few satellite devices are updated at a time to minimize the effects of a traffic disruption caused by too many satellite devices upgrading software simultaneously.

You can create and manage software upgrade groups from the Manage Software Upgrade Groups page. After you create a software image group, you can open the Deploy Images to Devices page and select a satellite software image for each software upgrade group and use the Select Options tab to set the date and time when the upgrade must be performed.

To access the Manage Software Upgrade Groups page:

1. Select the fusion fabric in the View pane for which you want to create a software upgrade group.
2. Click **Deploy** in the Network Director banner.
3. In the Tasks pane, select **Image Management > Manage Software Upgrade Group**.

The Software Upgrade Groups page opens in the main window. The table lists the software upgrade groups that exist for the selected fabric, if any.

NOTE: This task is available only if you select a fabric in the View pane.

You can perform the following tasks from the Manage Software Upgrade Groups page:

Create a Software Upgrade Group

To create a software upgrade group:

1. Click **Add**. The Add Software Upgrade Group window opens.
2. Enter a name for the new software upgrade group.
3. Select one or more satellite devices from the Select Devices tab that you want to be part of the software upgrade group. The Select Devices table lists the satellite devices and the available FPC slots in each of these that are not yet part of any other software upgrade groups.
4. Click the **FPC Range** tab and select the FPCs that you want to add to the software upgrade group. To add a single FPC number, specify the number in **From** and click **Add**. To add a range of FPC numbers, enter the starting and ending FPC numbers in **From** and **To** respectively, and click **Add**. You can also add FPCs that are currently inactive to the software upgrade group.

Network Director displays the FPC numbers that you added in the Selected FPC Number/Range table.

5. Click the **Preview** tab to review the software upgrade group settings. You can click the Select Devices or FPC Range tabs to modify the configuration settings.
6. Click **Add** to create the software upgrade group.

Network Director lists the new software upgrade group in the Manage Software Upgrade Group page. You can now use the deploy image task to assign a software image to this software upgrade group and deploy the image to the devices that are part of the upgrade group. For details on selecting and deploying software images, see [“Deploying Software Images” on page 1238](#).

Edit a Software Upgrade Group

To edit a software upgrade group:

1. Select a software upgrade group and click **Edit**.

The Edit Software Upgrade group window opens.

2. You can edit the FPC number and the FPC range.
3. Click **Save** to save the changes.

View Details of a Software Upgrade Group

To view the details of a software upgrade group:

1. Select a software upgrade group and click **Details**. The Software Upgrade Group Summary window opens displaying the members that are part of the upgrade group and the associated satellite image name.
2. Click **OK** to close the summary window.

Delete a Software Upgrade Group

To delete a software upgrade group, select a software upgrade group and click **Delete**.

RELATED DOCUMENTATION

[Deploying Software Images | 1238](#)

[Understanding Junos Fusion | 707](#)


[Software Requirements for Junos Fusion | 715](#)

Understanding Layer 3 Fabrics

Most enterprises that host data centers are looking to increase resiliency and also support new technologies such as VMware NSX that allow them to deploy applications, servers, and virtual networks within seconds. Layer 3 Fabrics allow them to support better uptime, performance, and newer cloud infrastructures such as VMware NSX. In order to maintain the large scale required to host thousands of servers, the use of a multi-stage Clos architecture is required. Such an architecture allows the physical network to scale beyond the port density of a single switch. Layer 3 Fabrics use BGP as the control plane protocol to advertise prefixes, perform traffic engineering, and tag traffic. The most common designs in a multi-stage Clos architecture are a 3-stage and 5-stage networks that use the spine-and-leaf topology.

Spine-and-leaf topology is an alternate to the traditional three-layer network architecture, which consists of an access layer, aggregation layer, and a core. In the spine-and-leaf topology, all the leaf devices are connected to the spine devices in a mesh as shown in [Figure 25](#).

Figure 25: Layer 3 Fabric in a Spine and Leaf Topology



Typically, the spine devices are high-performance switches capable of Layer 3 switching and routing combined with high port density. Spine devices constitute the core and the leaf devices constitute the

access layer in Layer 3 Fabrics. Leaf devices enable servers to connect to the Layer 3 Fabric. They also provide uplinks to spine devices.

Network Director currently supports only the 3-stage design. The 3-stage design has two roles—the spine and the leaf. It is called a 3-stage design because the traffic must traverse three switches in the worst-case scenario.

The maximum number of spine devices that you can have in your Layer 3 Fabric depends on the number of 40-Gigabit Ethernet interfaces in your leaf devices. A Layer 3 Fabric that has 8 QFX5100-24Q spine devices and 32 QFX5100-96S leaf devices (each leaf supports 96 10-Gigabit Ethernet ports) can provide 3072 usable 10-Gigabit Ethernet ports.

RELATED DOCUMENTATION

[Managing Layer 3 Fabrics | 744](#)

[Creating Layer 3 Fabrics | 746](#)

[Network Director Documentation home page](#)

User Privileges Required for the DHCP and File Server While Using Zero Touch Provisioning

Ensure that you have the following user privileges on the DHCP server and the file server prior to configuring them for zero touch provisioning (ZTP).

- DHCP server—Ensure that the DHCP user has permissions to:
 - write to the **dhcpd.conf** file on the DHCP server.

NOTE: To fetch the `dhcp.conf` file, ensure that the DHCP server and the Layer 3 Fabric devices are in the same subnets. If you are not in the same subnet, you must specify the gateway IP address that these devices can use to reach Network Director and fetch the `dhcp.conf` file. For information about specifying the gateway IP address, see, “[Creating Layer 3 Fabrics](#)” on page 746.

- write to the `/etc/dhcp/ddns-keys` directory
- copy the file **dhcpd.conf** to the file **dhcpdbacknd.conf**
- start the **isc-dhcp-server** service

For more information about file permissions, refer DHCP server documentation.

- File server—Network Director uses the *anonymous* user to connect to the file server. You must modify certain configurations in the server configuration file to enable Network Director to access the file server. Change the configuration settings for the following file servers, depending on the file server type and the operating system that is running on the file server:

- For FTP server running CentOS (or any other FreeBSD-based servers)—Modify the configuration in the `/etc/vsftpd/vsftpd.conf` file as follows:

```
anonymous_enable=YES
anon_upload_enable=YES
anon_mkdir_write_enable=YES
file_open_mode=0644
anon_umask=033
```

- For TFTP running on a Linux server—Modify the configuration in the `/etc/xinetd.d/tftp` file as follows:

```
server_args = -c -s <dir>
disable = no
```

RELATED DOCUMENTATION

[Configuring and Monitoring Zero Touch Provisioning | 1261](#)

[Creating Layer 3 Fabrics | 746](#)


Managing Layer 3 Fabrics

You can view and manage Layer 3 Fabrics in your network, using the Manage Layer 3 Fabric page.

To manage Layer 3 Fabrics:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View**, or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  **Build** or  **Deploy** in the Network Director banner.

3. In the Tasks pane, click **Network Builder > Manage Layer 3 Fabrics**.

The Manage Layer 3 Fabrics page opens.

4. [Table 157](#) describes the information provided about Layer 3 Fabrics on the Manage Layer 3 Fabrics page. This page lists all Layer 3 Fabrics defined for your network, regardless of the scope you selected in the network view.

Table 157: Manage Layer 3 Fabrics Field Descriptions

Field	Description
Fabric Name	Name given to the Layer 3 Fabric when the fabric was created.
Description	Description given when the fabric was created.
Summary	<p>Displays the following details about the fabric:</p> <ul style="list-style-type: none"> • Type of fabric—3-stage • Number of deployed and active spine and leaf devices. <p>Click View Topology to view the topology of the Layer 3 Fabric.</p>
Status	<p>Displays the current status of the Layer 3 Fabric. Status can be—In Design, In Progress, Deployed, or Failed.</p> <ul style="list-style-type: none"> • In Design—The fabric creation is in progress. The user might have saved and exited the Create Layer 3 Fabric wizard without specifying all the details. • In Progress—All the details that Network Director requires to create the fabric was entered by the user. Network Director might be performing background actions such as copying software images or configurations. Click this field to view the status of the jobs that are running in the background. • Deployed—The fabric is deployed and provisioned in the network, but might not be connected. • Failed—Creation or deployment of the fabric failed.
Cabling	<p>You can download the cabling plan or run a connectivity check by clicking the respective buttons in this field. You can also view details about the last run connectivity check by clicking View Connectivity Results.</p> <p>Mouse over this field to view the date and time when the last connectivity check was run.</p>
Updated Time	Time when the fabric was last updated.

TIP: All columns might not be displayed. To show or hide fields listed in the table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

5. You can perform the following tasks from the Manage Layer 3 Fabrics page:
 - Click **Create** to create a new Layer 3 Fabric. For detailed steps, see [“Creating Layer 3 Fabrics” on page 746](#).
 - Select a Layer 3 Fabric and click **Edit** to modify the fabric details. For detailed steps, see [“Editing Layer 3 Fabrics” on page 761](#).
 - Select a Layer 3 Fabric and click **Delete** to delete the fabric.
 - Click **Download Cabling Plan** in the Cabling column to download the cabling plan of the fabric. For more details, see [“Performing Layer 3 Fabric Connectivity Checks” on page 765](#).
 - Click **Run Connectivity Check** or **Re-run Connectivity Check** in the Cabling column to check the cabling plan for the fabric. Network Director performs the cabling check. You can view the result of the cabling check by clicking **View Connectivity Results**.
 - Click **View Topology** in the Summary column to view the physical topology of the Layer 3 Fabric. For more details, see [“Viewing Layer 3 Fabric Connectivity” on page 764](#).

RELATED DOCUMENTATION

[Creating Layer 3 Fabrics | 746](#)

[Editing Layer 3 Fabrics | 761](#)

[Performing Layer 3 Fabric Connectivity Checks | 765](#)

[Understanding Layer 3 Fabrics | 742](#)

Creating Layer 3 Fabrics

IN THIS SECTION

- [Specifying the Fabric Requirements | 747](#)
- [Specifying the Device Details | 752](#)
- [Specifying Configuration Details | 753](#)

- [Viewing the Cabling Plan | 754](#)
- [Specifying Zero Touch Provisioning Details | 757](#)
- [Reviewing the Layer 3 Fabric Settings | 760](#)

You can create and manage 3-stage Layer 3 Fabrics in Network Director by using the Create Layer 3 Fabrics wizard. Use the various pages of the wizard to specify the requirements and configurations for a Layer 3 Fabric. You can save the data that you have entered in one or more wizard pages, and come back later to specify the remaining details and complete the fabric creation.



CAUTION: Ensure that you always create the Layer 3 Fabric using this wizard and perform the physical connections based on the cabling plan that Network Director generates for your fabric. Not following this set order might render your Layer 3 Fabric defunct.

Before you begin, ensure that you have the necessary privileges on the FTP and the file server that Network Director uses for Zero Touch Provisioning. For more details, see [“User Privileges Required for the DHCP and File Server While Using Zero Touch Provisioning”](#) on page 743.

You can do the following tasks from the Create Layer 3 Fabric wizard pages:

Specifying the Fabric Requirements

To specify the fabric requirements:

1. Enter a name for the Layer 3 Fabric. The fabric name must be unique and can contain alphanumerals, hyphens, and underscores.
2. Enter a description for the Layer 3 Fabric.

NOTE: Network Director currently enables you to create 3-stage Layer 3 Fabrics and hence this is the default selection. You cannot modify the Fabric Type.

3. Select **QFX10008**, **QFX10002-36Q**, **QFX10002-72Q**, **QFX5100-24Q-2P** or **QFX5200-32C-32Q** as the device model for spine devices. All spine devices will be of the model that you select.

NOTE: If you select QFX10008 as spine, all the line cards must be homogenous across spines. For example, if you are building an IP fabric with four QFX10008 spines containing L1, L2, and L3 line cards, all the four spines must have L1, L2, and L3 line cards only and in the same slots.

4. Enter the number of spine devices that you plan to have initially and the maximum number of devices that you plan to have in this fabric, in the Initial Capacity and Max Capacity boxes respectively. You can have a minimum of 2 and a maximum of 8 spine devices.

NOTE: Initial capacity must be less than or equal to the maximum capacity. Maximum capacity must be greater than or equal to the initial capacity and must not be more than 8.

5. If you selected **QFX10008** as the spine device you must build the device chassis using the chassis builder.

To do this:

- a. Click **Build New Chassis** to create a new chassis.

The Build Chassis window opens. The Build Chassis window has two panes— the *Available line cards* pane and *the Chassis: FPC slots* pane.

The Available line cards pane lists the line cards that are supported on the selected aggregation device and the Chassis: FPC slots pane lists the available FPC slots on the device.

- b. Drag and drop the line cards that you want to add to the chassis from the Available line cards pane to the appropriate FPC slots in the Chassis: FPS slots pane.
 - c. Click **Set** after you have added all the required line cards to the FPC slots. The Build Chassis window closes.
 - d. Mouse over **Preview** to preview the chassis with the line cards that you added.
6. In the Fabric leaves section, click a row in the table to select a leaf device model and specify the capacity of the selected model that you plan to have in the fabric. Click **Add** to add subsequent rows.

NOTE: This is an optional step, however, it is mandatory to specify the maximum number of leaf devices you plan to have in this fabric. If you do not add any leaf devices, Network Director considers these devices as unknown and creates a cabling plan accordingly. After the fabric is deployed, you can plug and play any of the supported leaf device models to the fabric. After reaching the initial capacity for the spine devices, Network Director regenerates the cabling plan. Follow this plan to connect additional spine devices.

You can add one or more of the following device models as leaf devices in your Layer 3 Fabric:

- QFX5100-48S-6Q
- QFX5100-96S-8Q
- QFX5100-48T-6Q
- QFX5200-32C-32Q
- QFX5100-24Q-2P
- EX4300-32F
- EX4300-48P
- EX4300-24P
- EX4300-48T
- EX4300-24T

NOTE: QFX5100-48T-6Q can be standalone or Virtual Chassis leaf devices. QFX10008, QFX10000-36Q, QFX10000-30C, QFX10000-60S-6Q, and QFX5200 are supported only as standalone devices.

Network Director supports a maximum of two members in a Virtual Chassis.

If you want to delete a device entry, select a row and click **Remove**.

7. If you want to include Virtual Chassis as a leaf device, select **Include Virtual Chassis (VC) as a leaf**.

Enter the number of Virtual Chassis that you want to deploy immediately in **Initial Capacity** and the total number of Virtual Chassis that will be part of the Layer 3 Fabric in **Max. Capacity**. The minimum number of devices you can specify in **Initial Capacity** is 0. The **Max. Capacity** is the maximum number of devices you can specify, which depends on the spine device that you have selected. See [Table 158](#).

Table 158: Maximum Virtual Chassis Supported on Spine Devices

If you choose the spine device as...	then, the maximum number of virtual chassis leaf devices supported is...
QFX5100-24Q-2P	16
QFX10002-36Q	18
QFX5200-32C-32Q	16
QFX10002-72Q	36
QFX10008	~144 NOTE: Depends upon the type of line card connected.

NOTE: Plug and play is not supported for Virtual Chassis leaf members. Therefore, before you physically connect the Virtual Chassis members, make sure that you add the Virtual Chassis leaf members by using this Layer 3 Fabric wizard.

Initial capacity must be less than or equal to the maximum capacity.

For example, if your selected spine model is QFX5100-24Q-2P and if all of the leaf device members are Virtual Chassis, each containing 2 members, then the maximum number of Virtual Chassis leaf devices is restricted to 16, as there is a connection from both the master and backup member of the Virtual Chassis to each spine device. See [Table 159](#) for the maximum number of devices supported on various spine devices.

NOTE: You cannot modify the number of Virtual Chassis after the Layer 3 Fabric is created.

Network Director helps in creating the access link aggregation group (LAG) between Virtual Chassis members and host access devices. Network Director creates the access LAG in either of the two ways.

- **Dynamic LAG creation**—As the access devices are connected to the Virtual Chassis members, Network Director creates the LAG (if there are more than one connection between the access device and the Virtual Chassis members) dynamically. To identify the connected links for LAG creation, Network Director uses the Topology Discovery, which requires LLDP to be enabled in both the host and leaf members. For Network Director to create the LAG dynamically, ensure that LLDP is enabled in both the host and leaf (Virtual Chassis) devices.

- Preprovisioning LAG configuration on Virtual Chassis members—If LLDP is not enabled in the access or host devices, Network Director generates the LAG configuration on the Virtual Chassis member devices during the workflow creation and pushes the configuration to the Virtual Chassis members when they are connected to the network. The LAG interfaces are depicted in the cabling plan graph and in the grid view generated by Network Director. You must connect the host devices to the Virtual Chassis member devices according to the cabling plan.

You can enable Network Director to create the LAG as the physical connections are established. Select the **Dynamically create LAG when hosts are connected** check box.

8. Enter the maximum number of leaf devices, which includes standalone and Virtual Chassis devices, that the fabric can accommodate in **Max Capacity**. The minimum value you can enter is 1 and the maximum value depends on the spine device that you choose. See [Table 159](#).

Table 159: Maximum Number of Leaves

If you choose the spine device as...	then, the maximum number of leaves is...
QFX5100-24Q-2P	32
QFX10002-36Q	36
QFX10002-72Q	72
QFX10008 <ul style="list-style-type: none"> • QFX10000-36Q • QFX10000-60S-6Q • QFX10000-30C 	<ul style="list-style-type: none"> • < 288 • < 48 • < 240
QFX5200-32C	32

9. Do one of the following:
 - Click **Next** to open the Devices page where you can view and modify details of the spine and the leaf devices.
 - Click **Save & Exit** to save the changes and exit the Create Layer 3 Fabric wizard. You can resume this task later without losing any information that you entered.

Specifying the Device Details

The Devices page displays the number of spine and leaf devices that you are provisioning as part of the initial capacity, enables you to edit the hostname for all the spine and leaf devices. Select a model for each member of Virtual Chassis if you have opted for Virtual Chassis leaves, and search for a specific device in the fabric.

Network Director prefixes the name of the fabric that you specified in the Fabric Requirements page to the name of all the spine and leaf devices. If required, you can modify this prefix in the Devices page. You can also use the search box to search for specific devices in the fabric.

To specify the device details:

1. Click **Edit Host Name Prefix** if you want to change the device name prefix to something other than the name of the fabric. The Edit Host Name window opens.
2. Enter the name that you want to use as the device name and click **OK**.
Network Director replaces the device name prefix with the name that you entered.
3. The Devices page displays the details of the hostname and the devices associated with it. See [Table 160](#).

NOTE: The details of the device in each row, which is colored blue are to be provisioned now, and those colored orange are reserved for future allocation.

Table 160: Devices Page Description

Column	Description
Host Name	Displays the hostname with the name of the fabric, which you specified in the Fabric Requirement page.
Model	<p>Displays the model of the switch.</p> <p>If you have selected Virtual Chassis to be included in your Layer 3 Fabric in the Fabric Requirements page, the <i>Type</i> of the model will be <i>Virtual Chassis</i> and the <i>Model</i> is not displayed. You can select the switch model for the Virtual Chassis member from the drop-down list, which lists all supported Virtual Chassis members.</p> <p>NOTE: It is mandatory to select the switch model for Virtual Chassis member that you are provisioning now. For the Virtual Chassis members that are <i>Reserved for future</i> you may select the model later.</p>
Type	Displays the type of switch—standalone, virtual chassis, virtual chassis member, or FPC.

Table 160: Devices Page Description (continued)

Column	Description
Role	Displays the role being played by the switch model.

4. Do one of the following:

- Click **Next** to open the Configuration page, where you can specify the configuration details of the Layer 3 Fabric.
- Click **Save & Exit** to save the changes and exit the Create Layer 3 Fabric wizard. You can resume this task later without losing any information that you entered.

Specifying Configuration Details

To specify the configuration details for the Layer 3 Fabric:

1. Enter details in the Configuration page by following the descriptions given in [Table 161](#).

Table 161: Layer 3 Fabric Configuration Details

Field	Description
Loopback Network Address	Specify the IP address block that you want to use for configuring the loopback interface in each member. Each device in the fabric is assigned one IP address from the block. This IP address can be used for troubleshooting and for checking connectivity between switches.
Interconnect Network Address	Specify the IP address block that you want to use for configuring the IP addresses for interconnect links between leaves and spines. Each interconnect link is assigned two IP addresses from this block.
VLAN Network Address	Specify the IP address block to be reserved for the virtual machines or hosts that you want to connect to the leaves. Network Director allocates each leaf device with a subnet from the given IP address block.
Start Management IP	Specify the management IP address that Network Director will use to manage each switch. NOTE: If you have provisioned for Virtual Chassis members in the Layer 3 Fabric, each Virtual Chassis member is initially treated as a standalone device and it goes through the ZTP process. The Management IP address block is sufficient to provide individual unique IP address for each of the Virtual Chassis member in the fabric.

Table 161: Layer 3 Fabric Configuration Details (*continued*)

Field	Description
Max Hosts/VMs per leaf	Specify the maximum IP addresses that are required in the subnet to be allocated from the VLAN Network Address.
Spine-BGP Autonomous System Number	Specify the starting autonomous system (AS) number to be assigned to the first spine device. Subsequent spine devices are assigned incremental AS numbers starting from the number you specified. Network Director updates the last AS number based on the number of spine devices that you plan to have in the fabric. You cannot modify the last AS number.
Leaf-BGP Autonomous System Number	Specify the starting autonomous system (AS) number to be assigned to the first leaf device. Subsequent leaf devices are assigned incremental AS numbers starting from the number you specified. Network Director updates the last AS number based on the number of leaf devices that you plan to have in the fabric. You cannot modify the last AS number.
Device Password	Specify the default password that you want to set for all the devices in the fabric.
Management Gateway	If Network Director and the Layer 3 Fabric devices are in different subnets, specify the gateway IP address that these devices can use to reach Network Director. NOTE: This is an optional field if the Layer 3 Fabric and Network Director are in the same subnet.

2. Do one of the following:

- Click **Next** to open the Cabling page where you can view the cabling plan for your Layer 3 Fabric. This might take some time depending on the fabric capacity.
- Click **Save & Exit** to save the changes and exit the Create Layer 3 Fabric wizard. You can resume this task later without losing any information that you entered.

Viewing the Cabling Plan

The Cabling Plan page displays the recommended cabling plan for the device that you select in the left pane. If you specify all the spine and leaf devices, the cabling plan displays the exact port numbers that you must use to connect your spine and leaf devices. However, if you have not specified any leaf devices and have only specified the maximum leaf count, the plan displays all the leaf devices as unknown. The leaf devices in this case are plug-and-play and you can use any of the uplink ports on your plug-and-play leaf device.

This holds good until you have reached the initial capacity of the spine devices. If you are adding an additional spine device, beyond the initial capacity, Network Director regenerates the cabling plan and you must follow the recommended cabling plan for all subsequent spine to leaf connections. Note that the connections to the existing devices need not be changed as part of this change.



Network Director regenerates the cabling plan, if one of the following occurs:

- A spine device is added
- A spine device is deleted
- A leaf device is added
- A leaf device is deleted

If the selected spine device model in the Fabric Requirements page is QFX10002-72Q, the cabling plan is represented as two chassis images. The first chassis image displays the connections for the ports in the first and second rows, and the second chassis image displays the connections for the ports in the third and fourth rows.

If the selected spine device model in the Fabric Requirements page is QFX10008, and selected line card model is QFX10000-60S-6Q in the Build New Chassis section, cabling plan is represented in two chassis images. The first image displays connections for the ports in first and third rows, and the second chassis image displays the connections for the ports in the middle row.

From the Cabling page, you can:

1. View the cable connectivity that you must follow for each device in your fabric. The device table is color coded to identify the devices that are provisioned now (identified by  color), the devices reserved for future (identified by  color), the Virtual Chassis connections (identified by green color), and access LAG ports (identified by pink color).
2. Click **Grid View** to view the cabling plan in a grid. Select a device in the left pane to view the cabling details and access LAG connections of the selected device.

NOTE: The Access LAG ports are displayed if you have not selected **Dynamically create LAG when hosts are connected** in the Fabric Requirement page. Network Director preprovisions the LAG configuration in the Virtual Chassis members.

3. Click **Graph View** to view the graphical representation of the cabling plan.
4. Do one of the following:
 - Click **Next** to open the ZTP page where you can specify the Zero Touch Provisioning (ZTP) details for the Layer 3 Fabric.

- Click **Save & Exit** to save the changes and exit the Create Layer 3 Fabric wizard. You can resume this task later without losing any information that you entered.

Specifying Zero Touch Provisioning Details

Zero Touch Provisioning (ZTP) enables you to provision devices in your network automatically, without manual intervention. When a device is physically connected, it boots up with factory-default configuration and auto installs a configuration file from the network. In Network Director, the ZTP is used to provision Layer 3 fabric and all the configurations are pushed through OpenClos. To specify the ZTP details:

NOTE: When you select QFX10008 as the spine model, only the leaf models are provisioned with ZTP configuration. For the spine model QFX10008, you must either copy the config file from Network Director or manually download it from the file server. To copy the config file from the file server, SSH or Telnet must be enabled on the device (QFX10008).

1. Specify the DHCP server settings by following the descriptions given in [Table 162](#).

Table 162: DHCP Server Details

Field	Description
DHCP Server	IP address or the hostname of the DHCP server.
DHCP Server Type	The type of DHCP server that provides the necessary information to the switch. You can choose to use a CentOS DHCP server, an Ubuntu DHCP server, or any other DHCP server. NOTE: If you select Other, you must configure the DHCP server settings manually.
Manually Configure Server	Select to indicate that you want to manually configure the DHCP server. You can configure the CentOS and Ubuntu DHCP servers manually or from Network Director. If you want to use any other type of DHCP server, do the following: <ol style="list-style-type: none"> a. Select the Manually Configure Server check box. Network Director hides all the other details except the DHCP Server Type. b. Follow the instructions displayed in this box to configure the DHCP server manually.
DHCP User	Username to log in to the DHCP server.
DHCP Password	Password for the specified username.
Confirm Password	Confirm the DHCP server password.

NOTE: *When you are replacing a member device*—If the member that is replaced is up, Network Director obtains the latest configuration from the replaced device and maps this configuration to the corresponding MAC or serial number in the DHCP server.

However, if the member that is replaced is down, Network Director is not able to reach the device to get its latest configuration. In such case, Network Director maps the configuration that is generated from OpenClos (Stage-2 for leaf devices) for the replaced device to the MAC or serial number of the new device in the DHCP server. Note that the mapped configuration in the DHCP server does not have any configuration that is pushed from Network Director to the device.

NOTE: The DHCP server configuration file does not contain entries related to the spine device QFX10008 as the device does not go through ZTP.

2. Specify the File server settings by following the descriptions given in [Table 163](#).

Table 163: File Server and Software Details

Field	Description
File Server Type	The type of file server where the software images are to be stored. You can choose to use an FTP, HTTP, or an TFTP file server.
File Server	IP address or hostname of the file server.
File Server Root Dir	The root directory of the file server.
Spine Image	<p>The software image file that you want to use for your spine devices.</p> <p>NOTE: Ensure that the software image is uploaded to Network Director using the Image Management > Manage Image Repository in the Deploy mode. Else Network Director does not display the software image.</p>
Leaf Image	<p>The software image file that you want to use for your leaf devices.</p> <p>NOTE: Ensure that the software image is uploaded to Network Director using the Image Management > Manage Image Repository in the Deploy mode. Else Network Director does not display the software image.</p> <p>As Network Director supports two device models—EX4300 and QFX5100—and their variants as leaf devices, you can specify a software image for each of these leaf devices irrespective of the variant that you have selected for your fabric. The same software image applies to all the variants of a device.</p>

3. ZTP process maps the management interface MAC address or the device chassis serial number of each spine device to the device-specific software image, IP address, hostname, and the configuration file stored on the file server. This mapping is stored in the DHCP server. When a spine device starts up, the device contacts the DHCP server to obtain the IP address and the software image location. The DHCP server looks up in its MAC address or serial number mapping database to identify the device and provide details about the file server that the device must contact to get the software image and configuration file. The device uses this information to contact the file server and obtain the software image and the configuration file for deploying on the device.

Do one of the following to specify the MAC address or the serial number of your spine devices:

- Enter the MAC address of the management interface (for example, the em0 interface) or the device chassis serial number of the spine devices in **MAC Addresses** or **Serial Number** in the table.
- Click Import MAC Address to import the MAC addresses of spine device in CSV format. You must enter the MAC addresses in the specified format. Click **Download CSV format** to download a sample CSV file that you can use to import MAC addresses.

NOTE: When you use the Import option, you must specify either the MAC address or the Serial number, but not both.

- Click Import Serial Number to import the serial numbers of spine device in CSV format. You must enter the MAC addresses in the specified format. Click **Download CSV format** to download a sample CSV file that you can use to import serial numbers.

NOTE: You can specify serial number only for spine devices running Junos OS Release 14.1X53D15 or later.

NOTE: Entering serial number or MAC address of the spine device is not applicable for the device model QFX10008 as it is not provisioned through ZTP.

4. To view the configuration that is deployed on a spine device, click **Actions > View Config**.
5. To view the configuration that is deployed initially on the leaf devices, click **View Leaf Config**.
6. Do one of the following:
 - Click **Next** to open the Review page where you can review the Layer 3 Fabric settings.
 - Click **Save & Exit** to save the changes and exit the Create Layer 3 Fabric wizard. You can resume this task later without losing any information that you entered.

Reviewing the Layer 3 Fabric Settings

From the Review page you can:

- View the DHCP configuration that will be deployed on to the DHCP server by clicking **View DHCP Config** in the **ZTP Settings** sub-tab.

The DHCP configuration opens in a new window.

- Review the Layer 3 Fabric settings in the Review page, Devices sub-tab and the Configuration sub-tab.

- Click **Deploy** to deploy the Layer 3 Fabric.
- Click **Save & Exit** to save the changes and exit the Create Layer 3 Fabric wizard. You can resume this task later without losing any information that you entered.

RELATED DOCUMENTATION

[Understanding Layer 3 Fabrics | 742](#)

[Managing Layer 3 Fabrics | 744](#)

[Editing Layer 3 Fabrics | 761](#)

[Network Director Documentation home page](#)

Editing Layer 3 Fabrics

Network Director enables you to edit some of the settings for a Layer 3 Fabric after you deploy the fabric successfully.

To edit a Layer 3 Fabric:

1. From the Manage Layer 3 Fabric page, select the fabric that you want to edit and click **Edit**.

The Edit Layer 3 Fabric page opens. The Edit Layer 3 Fabric page has similar wizard pages as in the Create Layer 3 Fabric page. However, you will not be able to modify all the fields in these wizard pages while you are editing the settings of a fabric.

2. [Table 164](#) lists the settings that you can edit and the wizard pages to which these settings belong.

Table 164: Layer 3 Fabric Settings that can be edited

Wizard page	Field Name	Action
Fabric Requirement	Description	<p>Modify the description to a description of your choice.</p> <p>NOTE: If you add a plug-and-play leaf device to the fabric, make sure that you modify the Description field. If you do not do this, the Cabling page might not update the cabling plan for that leaf in the graph and grid views.</p>

Table 164: Layer 3 Fabric Settings that can be edited (continued)

Wizard page	Field Name	Action
Devices	Add Spine	<p>To add a spine device:</p> <ol style="list-style-type: none"> Click Add Spine. The Select Devices window opens. Select a device that you want to add and click OK. Network Director adds the selected device to the device table and updates the status of the device as Added. In the ZTP wizard page, specify the MAC address or serial number of the added device.
Devices	Add Leaf (VC)	<p>To add a virtual chassis device as a leaf member:</p> <ol style="list-style-type: none"> Click Add Leaf (VC). The Select Devices window opens. Select a device that you want to add and click OK. Network Director adds the selected device to the device table and updates the status of the device as Added. You can specify the model for Virtual Chassis members. Once you have specified the model for the Virtual Chassis member, then the cabling plan changes, and Network Director displays a new cabling plan in the Cabling Plan page. Follow the cabling plan to physically connect the devices.
Devices	Remove	Select a device from the list and click Remove .

Table 164: Layer 3 Fabric Settings that can be edited (*continued*)

Wizard page	Field Name	Action
Devices	Replace	<p>To replace device:</p> <ol style="list-style-type: none"> a. Select the check box corresponding to device that you want to replace from the list and click Replace. <p>Network Director updates the status of the device as Replaced.</p> <p>NOTE: You cannot replace an inactive leaf device.</p> b. For spine devices and standalone leaf devices, in the ZTP wizard page, specify the MAC address or serial number of the replaced device. <p>NOTE: While replacing a spine device, you can specify either the MAC address or the serial number. However, while replacing a leaf device, you can specify only the MAC address.</p> <p>NOTE: When you are replacing a spine devices or a standalone leaf devices, if the device that is replaced is up, Network Director obtains the latest configuration from the replaced device and maps this configuration to the corresponding MAC address or serial number of the new device in the DHCP server.</p> <p>However, if the device is down, Network Director is unable to reach the device to get its latest configuration. In such case, Network Director maps the configuration that is generated from OpenClos (Stage-2 for leaf devices) for the new device to the MAC address or serial number of the device in the DHCP server. Note that the mapped configuration in the DHCP server will not have any configuration that is pushed from Network Director to the device.</p> <p>NOTE: You can replace only one member of a Virtual Chassis at a time. While replacing a Virtual Chassis member, you need not specify the MAC address or serial number of the device.</p>

Table 164: Layer 3 Fabric Settings that can be edited (continued)

Wizard page	Field Name	Action
ZTP	<i>Image Details and the Device Details</i>	<p>You can modify the software image and the Device details in the ZTP wizard page.</p> <p>New and replaced devices are listed in the Device Details section. You must specify the MAC address or serial number for these devices.</p> <p>NOTE: While replacing a spine device, you can specify either the MAC address or the serial number. However, while replacing a leaf device, you can specify only the MAC address.</p> <p>For more details about the fields in the ZTP wizard page, see “Configuring and Monitoring Zero Touch Provisioning” on page 1261.</p>

3. Click **Deploy** to save and deploy your edits.

RELATED DOCUMENTATION

[Understanding Layer 3 Fabrics | 742](#)

[Managing Layer 3 Fabrics | 744](#)

[Network Director Documentation home page](#)

Viewing Layer 3 Fabric Connectivity


After you have set up and deployed a Layer 3 Fabric in Network Director, you can pictorially view the physical connectivity between the various devices in the fabric. This page displays the devices and their physical connectivity in the spine-and-leaf topology.

To view the connectivity between the devices:

1. Do one of the following:
 - While in the Logical, Location, Device, or Custom View, select the Layer 3 Fabric for which you want to view the connectivity details from the View pane and click **Connectivity > View Layer 3 Fabric Connectivity** from the Tasks pane.
 - Click **View Topology** in the Manage Layer 3 Fabric page.

- While in the topology view, zoom in to a rack that has a Layer 3 Fabric member device, select the device and click **View Layer 3 Fabric Connectivity** from the Task pane.

The Layer 3 Fabric Connectivity page opens.

2. You can perform the following tasks from the Layer 3 Fabric Connectivity page.
 - Mouse over each entity to know more details about that entity.
 - You can Zoom in or zoom out of the connectivity view by using the + and - buttons.
 - View the faults and alarms on an entity by zooming in to the entity. Double-click the alarm or fault count to view more details about it in the Fault mode.
 - Click  on a leaf device to expand and view the host machines that are connected to the device.

RELATED DOCUMENTATION

[Understanding Layer 3 Fabrics | 742](#)

[Network Director Documentation home page](#)

Performing Layer 3 Fabric Connectivity Checks

After you have deployed a Layer 3 Fabric, you can run connectivity checks to troubleshoot any connectivity issues in the fabric. You can initiate a connectivity check from the Manage Layer 3 Fabrics page. Network Director uses LLDP to check the connectivity of each device with the device's neighbor and compares it with the recommended cabling plan and reports the results in the Cabling Check Results page.

NOTE: Network Director performs connectivity check only for the devices that are scheduled for immediate deployment and not for devices that are reserved for the future.

The Cabling Check Results page contains three tabs—Unknown Devices, Devices with Cabling Faults, and Devices Connected Properly. Each of these tabs displays a graphical representation and a grid view of each leaf device with the corresponding details. To switch between the Graphical view and the Grid view, toggle the **Switch to Grid View** and **Switch to Graph View** buttons.

From the Cabling Check Results page, you can:

1. Click the **Unknown Devices** tab to view the devices that are part of the Layer 3 Fabric, but are not physically connected to the network and have not undergone zero touch provisioning. Click each leaf device listed in the Devices box to see a detailed view of the device with the connections as per the cabling plan.
2. Click the **Cabling Faults** tab to view the devices which did not adhere to the recommended cabling plan. This might be because of missing connections or additional connections. You can view the details section to identify the faulty connections for each device.
3. Click **Devices Connected Properly** tab to view details of devices that are connected as per the cabling plan. Click each device to see the connections.
4. Click **Export Results to PDF** to export the connectivity check results as a PDF file.
5. Click **OK** to close the Cabling Check Results page.

RELATED DOCUMENTATION

[Managing Layer 3 Fabrics | 744](#)

[Network Director Documentation home page](#)

Setting Up Virtual Chassis Fabrics


IN THIS SECTION

- [Selecting a Provisioning Method | 767](#)
- [Adding Devices | 768](#)
- [Specifying Software Packages | 772](#)
- [Reviewing the Virtual Chassis Fabric Summary and Deploying Changes | 773](#)

The Juniper Networks Virtual Chassis Fabric (VCF) provides a low-latency, high-performance fabric architecture that can be managed as a single device. The VCF architecture is optimized to support small and medium-sized data centers that contain a mix of 1-Gbps, 10-Gbps, and 40-Gbps Ethernet interfaces.

Network Director supports VCFs that are constructed using the spine-and-leaf architecture. In the spine-and-leaf architecture, each spine device is connected to one or more leaf devices. You can configure up to 20 QFX5100 switches in a VCF configuration, with two to four QFX5100s in the spine and up to 18 QFX5100s as leaf nodes. See [Figure 26](#) for an illustration of the VCF spine-and-leaf architecture:

Figure 26: VCF Spine-and-Leaf Architecture



This topic describes how to set up a VCF using Network Director.

Before you setup a VCF, make sure at least one spine device is managed by Network Director. You can use this spine device to designate as a master device in the following configuration.

To start setting up a VCF:

1. Select **Build** or **Deploy** in the Network Director banner.
2. Select **Network Builder** > **Set Up Virtual Chassis Fabric** in the Tasks pane.

The Select Provisioning Method page of the Setup Virtual Chassis Fabric wizard opens.

This topic describes:

Selecting a Provisioning Method

Use the Select Provisioning Method page to select the VCF provisioning method.

1. Select a provisioning method.

For a description of the fields on the Select Provisioning Method page, see [Table 165](#).

Table 165: Select Provisioning Method Page

Field	Description
Select Provisioning Method Section	
Select the provisioning method.	
Auto Provision (Plug n Play)	Selects automatic provisioning. You need to provision only the spine nodes when you use this method. All eligible nodes connected to spine nodes that are running the factory-default configuration will be added automatically to the VCF.
Manual Provision	Selects manual provisioning. You must manually provision all VCF nodes when you use this method.
Enable Mixed Switch Mode	Selects a mixed mode VCF, which can contain a mix of device models. If a VCF is not in mixed mode, it can contain only QFX 5100 devices. In a spine-and-leaf mixed mode VCF, the spine nodes must be QFX 5100 devices.

2. Click **Next** to continue the wizard. The Add Devices page opens.

Adding Devices

Use the Add Devices page to assign devices to the VCF.

The following sections describe the methods of assigning devices to a VCF:

- [Assigning Devices to a Virtual Chassis Fabric by Using a Graph | 768](#)
- [Assigning Devices to a Virtual Chassis Fabric by Using a Grid | 771](#)

Assigning Devices to a Virtual Chassis Fabric by Using a Graph

On the graph, the VCF is represented by two sets of rectangles that represent devices. The top set represents spine devices, and the bottom set represents leaf devices.

Follow these steps to assign devices to a VCF by using a graph:

1. If the grid view is open, click **Switch to Graph View** on the Assign Devices page.
2. Click **Add Spines** to add spine devices. The Add Spine Switches to Virtual Chassis Fabric window opens.
3. To add devices that are managed by Network Director as spine nodes, click **Add Managed Devices as Spine Switches**.

If the VCF is in non-mixed mode, you can select only QFX 5100 devices to be members. Do the following for each device that you want to add in the Add Managed Switches to Virtual Chassis Fabric window that opens:

- a. Select the check box next to the device.
- b. Select the node's role in the **Role** column.
- c. Select the star icon in the Designated Master column if you want to make it the Master Routing Engine node.

Only nodes with the role Routing Engine can be designated master.
- d. When you finish configuring nodes, click **OK**. The window closes and the added devices appear in the node list on the Add Devices as Spine Node(s) window.

4. To add devices that are not managed by Network Director as spine or Routing Engine nodes, click **New Devices**. Do the following in the Add New Devices window that opens:

- a. Enter the device serial number in the **Serial Number** field.
- b. Select the node's role from the **Role** list.
- c. To add another device, click **Add More** and repeat the steps to add more devices.
- d. When you are finished adding nodes, click **Add**. The window closes and the added devices appear in the node list on the Add Devices as Spine Node(s) window.

5. When you finish adding spine nodes, do the following in the Add Devices as Spine Node(s) window:

- To remove a node from the list, click the trash can icon at the end of the node's table row.
- When you finish configuring the list of spine nodes, click **Done**. The window closes and the spine nodes you configured appear in the Graph on the Setup Virtual Chassis Fabric page. The Designated Master Routing Engine node is identified with a star.

6. (Optional) If you selected the Auto Provisioning (Plug and Play) provisioning method, you can assign devices that are managed by Network Director to leaf nodes. Manually assigning a leaf node device causes the device to reset to the default configuration. When a device with the default configuration is connected to a VCF spine node, the device is added as a leaf node:

- a. Select the **Assign Leaf nodes manually** check box.
- b. Click a rectangle representing a leaf node to assign one or more devices to that role. The Add Device as Leaf Node(s) window opens.

- c. Click **Managed Devices**. The Select Managed Device(s) window opens. If the VCF is in non-mixed mode, you can select only QFX 5100 devices to be members.
 - d. Select the check box for the device(s) you want to add as leaf nodes.
 - e. Click **OK**. The window closes and the selected nodes are listed on the Add Device as Leaf Node(s) window.
 - f. In the Add Device as Leaf Node(s) window, do the following:
 - To remove a node from the list, click the trash can icon at the end of the node's table row.
 - When you finish configuring the list of leaf nodes, click **Done**. The window closes and the leaf nodes you configured appear in the graph on the Setup Virtual Chassis Fabric page.
7. If you selected the Manual Provision provisioning method, click **Add Leaves** to add one or more devices as leaf nodes. The Add Device to Leaf Node(s) window opens.
8. To add devices that are managed by Network Director as leaf nodes:
- a. Click **Managed Devices**. The Select Managed Devices window opens. If the VCF is in non-mixed mode, you can select only QFX 5100 devices to be members.
 - b. Select the check box next to the device(s) to add.
 - c. Click **OK**. The window closes and the added devices appear in the node list on the Add Devices as Leaf Node(s) window.
9. To add devices that are not managed by Network Director as leaf nodes:
- a. Click **New Devices**. The Add New Devices window opens.
 - b. Enter the device serial number in the **Serial Number** field.
 - c. To add another device, click **Add More** and repeat the steps to add more devices.
 - d. When you are finished adding nodes, click **Add**. The window closes and the added devices appear in the node list on the Add Devices as Leaf Node(s) window.
10. When you finish adding leaf nodes, do the following in the Add Devices as Leaf Node(s) window:
- To remove a node from the list, click the trash can icon at the end of the node's table row.

- When you finish configuring the list of leaf nodes, click **Done**. The window closes and the leaf nodes you configured appear in the graph on the Setup Virtual Chassis Fabric page.

11. Click **Next** to continue the wizard. The Specify Software Package page opens.

Continue with [“Specifying Software Packages” on page 772](#).

Assigning Devices to a Virtual Chassis Fabric by Using a Grid

This section describes how to assign devices to a VCF by using a grid.

To assign devices to a VCF by using a grid:

1. If the graph view is open, click **Switch to Grid View** on the Assign Devices page.
2. To add devices that are managed by Network Director as nodes, click **Managed Devices**. If the VCF is in non-mixed mode, you can select only QFX 5100 devices to be members. Do the following for each device that you want to add in the Select Managed Devices window that opens:
 - a. Select the check box next to the device.
 - b. Select the node's role in the **Role** column.
The options are Routing Engine or Line Card.
 - c. Select the node's type in the **Type** column.
The options are Spine or Leaf. Nodes with the role of Routing Engine must have the type Spine.
 - d. Select the star icon in the **Designated Master** column if you want to make the node the master Routing Engine node.
Only nodes with the role of Routing Engine can be Master. If you select the Designated Master star for a node with the role of Line Card, the role is changed to Routing Engine, and if necessary, the type is changed to Spine.
 - e. When you are finished configuring nodes, click **OK**. The window closes and the added devices appear in the node list on the Add Devices page.
3. To add devices that are not managed by Network Director as nodes, click **New Devices**. Do the following in the Add New Devices window that opens:
 - a. Enter the device serial number in the **Serial Number** field.
 - b. Select the node's type from the **Type** list.
The type Leaf is available only if you selected the Manual provisioning method.

- c. Select the node's role from the **Role** list.

The role Routing Engine is applicable only to nodes of the type Spine. If you select Routing Engine, the type changes to Spine automatically.

- d. To add another device, click **Add More** and repeat the steps to add more devices.

- e. When you are finished adding nodes, click **Add**. The window closes and the added devices appear in the node list on the Add Devices page.

4. When you finish adding nodes, you can remove a node from the list on the Add Devices page by clicking the trash can icon at the end of the node's table row.

5. Click **Next** to continue the wizard. The Specify Software Package page opens.

Continue with ["Specifying Software Packages" on page 772](#).

Specifying Software Packages

You can configure a VCF to automatically upgrade new devices that are added to the VCF to a specified software package. For this functionality to work, you must specify the location of the software upgrade image files. Use the Specify Software Package page to specify the location of the image files for each device type that is included in the VCF. If the VCF is in non-mixed mode, only the QFX 5100 device type will appear on this page because it is the only device type allowed in the VCF.

To configure software packages on the VCF, do the following for each device type listed on the page:

1. Select the location of the software image by selecting a radio button:
 - Select **FTP Server** if the software image file is located on an FTP server.
 - Select **HTTP Server** if the software image file is located on an HTTP server.
 - Select **Device** if the software image file is located on the device.
2. Enter the URL where the software image file is located in the text box:
 - If you selected **FTP Server** or **HTTP Server**, the text box is labeled Software Package URL. Enter the URL where the software image file is located on the server. Begin the URL with **ftp://** for an FTP server or **http://** for an HTTP server.
 - If you selected **Device**, the text box is labeled Software Package Pathname. Enter the pathname where the software image file is located on the device. Begin the pathname with a slash (/) to indicate the root directory.
3. Click **Next** to continue the wizard.

The Review page opens.

Reviewing the Virtual Chassis Fabric Summary and Deploying Changes

In the Review page, review the VCF setup. To make any changes, click **Back** to return to a previous step, or click the step in the wizard flowchart at the top of the page.

To deploy the VCF configuration, click **Deploy**.

When you deploy the VCF configuration, Network Director performs these major steps:

1. Configures the VCF on the Master device, including the mixed/non-mixed mode setting, then reboots the Master device.
2. Configures the VCF management interface (vme) based on the Master device's management interface (for example, the em0 interface).
3. Provisions the VCF member devices on the Master device.
4. Restarts all designated VCF devices, other than the Master device, in factory-default mode.

The Master device will remain a standalone switch in network Director until more than one VCF member device becomes operational, when it will become a VCF. Until this happens, you cannot edit the VCF using the Manage Virtual Chassis Fabric task. You can see the status of the deployment job. See "[Managing Configuration Deployment Jobs](#)" on page 1194 for information. If deployment to the VCF Master device fails, the entire VCF deployment is cancelled.

Note these considerations when you create or manage a VCF using Network Director:

- If the management interface of a designated member is configured using DHCP, Network Director might lose connectivity with the member after successfully forming a VCF.
- If the em1 or em2 interface is used to manage a designated member, Network Director might lose connectivity with the member after successfully forming a VCF.

RELATED DOCUMENTATION

[Understanding Build Mode in Network Director | 184](#)

[Network Director Documentation home page](#)

Managing Virtual Chassis Fabrics

IN THIS SECTION

- [Adding Spine Device\(s\) to the Virtual Chassis Fabric | 774](#)
- [Adding Leaf Node\(s\) to the Virtual Chassis Fabric | 775](#)
- [Replacing Spine or Leaf Devices | 776](#)
- [Removing Spine or Leaf Devices | 777](#)
- [Configuring Software Package for Upgrade | 778](#)
- [Deploying VCF Configuration Changes | 778](#)

Network Director supports Virtual Chassis Fabrics (VCFs) that are constructed using the spine-and-leaf architecture. In the spine-and-leaf architecture, each spine device is connected to one or more leaf devices.

Network Director enables you to modify an existing VCF by using the Manage Virtual Chassis Fabric task. You can add, remove, or replace VCF members, and configure software packages for automatically upgrading new member devices. If the VCF is in non-mixed mode, you can add only QFX 5100 devices.

To modify a VCF:

1. In Build mode, select the VCF that you want to modify from the View pane and select **Device Management > Manage Virtual Chassis Fabric** from the Tasks pane.

The Virtual Chassis Topology page opens displaying the connectivity of the VCF in the spine-and-leaf architecture. The Designated Master node is marked with a gold star. The Backup Master node is marked with a silver star. Nodes that are not present, inactive, unprovisioned, or pre-provisioned appear in gray color. You cannot replace or remove such nodes.

2. Click **Edit VCF**. Network Director displays the Add Spines, Add Leaves, and Specify Software Package buttons.

You can perform the following tasks from the Virtual Chassis Topology page:

Adding Spine Device(s) to the Virtual Chassis Fabric

A VCF can contain up to 4 spine nodes and a total of 20 nodes.

To add a spine device:

1. Click **Add Spine**. The Add Device as Spine Node(s) window opens. If the VCF is in non-mixed mode, you can add only QFX 5100 devices, so only those devices are listed.
2. To add devices that are managed by Network Director as spine nodes, click **Managed Devices**. Do the following for each device that you want to add in the Select Managed Devices window that opens:
 - a. Select the check box next to the device.
 - b. Select the node's role in the **Role** column.
 - c. When you have selected all the nodes you want, click **OK**. The window closes and the added device appears in the node list on the Add Devices as Spine Node(s) window.
3. To add devices that are not managed by Network Director as spine nodes, click **New Devices**. Do the following in the Add New Devices window that opens:
 - a. Enter the device serial number in the **Serial Number** field.
 - b. Select the node's role from the **Role** list.
 - c. To add more devices, click **Add More** and repeat the steps to add more devices.
 - d. When you finish adding devices, click **Add**. The window closes and the added devices appear in the node list on the Add Devices as Spine Node(s) window.
4. When you finish adding spine nodes, do the following in the Add Devices as Spine Node(s) window:
 - To remove a node from the list, click the trash can icon at the end of the node's table row.
 - When you finish configuring the list of spine nodes, click **Done**. The window closes and the spine nodes you configured appear as spine devices in the Virtual Chassis Topology page. The master Routing Engine node is identified with a gold star. The backup Routing Engine node is identified with a silver star.
5. Click **Deploy** to deploy the changes that you made to the VCF.

Adding Leaf Node(s) to the Virtual Chassis Fabric

A VCF can contain up to 4 spine nodes and a total of 20 nodes.


To add leaf nodes to the VCF:

1. Click **Add Leaves**. The Add Device as Leaf Node(s) window opens. If the VCF is in non-mixed mode, you can only add QFX 5100 devices, so only those devices are listed.
2. To add managed devices as leaf nodes, click **Managed Devices**:
 - a. Select the check box next to the devices you want to add.
 - b. Click **OK**. The window closes and the added devices appear in the node list on the Add Devices as Leaf Node(s) window.
3. To add a device that is not managed by Network Director as a leaf node, click **New Devices**. This option is not available if the VCF is configured for automatic provisioning. Do the following in the window that opens:
 - a. Enter the device serial number in the **Serial Number** field.
 - b. To add another device, click **Add More** and repeat the steps to add more devices.
 - c. When you finish adding devices, click **Add**. The window closes and the added devices appear in the node list on the Add Devices as Leaf Node(s) window.
4. When you finish adding leaf nodes, do the following in the Add Devices as Leaf Node(s) window:
 - To remove a node from the list, click the trash can icon at the end of the node's table row.
 - When you finish configuring the list of leaf nodes, click **Done**. The window closes and the leaf nodes you configured appear as leaf devices in the Virtual Chassis Topology page.
5. Click **Deploy** to deploy the changes that you made to the VCF.

Replacing Spine or Leaf Devices


Network Director enables you to replace a spine or leaf device with another VCF-compatible device. You cannot replace the Designated Master node.

To replace a spine or leaf device:

1. Mouse over the leaf or spine device that you want to replace and click . The Replace Spine Node or the Replace Leaf Node window opens depending on the type of device that you want to replace.
2. To replace a spine or the leaf node with a device that Network Director is managing, click **Managed Devices**. Do the following for the device that you want to replace in the Select Managed Devices window that opens:
 - a. Select the check box next to the device.
 - b. Select the node's role in the **Role** column.
 - c. Click **OK**. The window closes and the selected devices appear in the node list on the Replace Spine Node or the Replace Leaf Node window.
3. To replace a device that is not managed by Network Director as a spine or the leaf node, click **New Devices**. This option is not available if the VCF is configured for automatic provisioning. Do the following in the Add New Devices window that opens:
 - a. Enter the device serial number in the **Serial Number** field.
 - b. Click **Add**. The window closes and the added devices appear in the node list on the Replace Spine Node or the Replace Leaf Node window.
4. Click **Done**. The window closes and the nodes that you replaced appear in the Virtual Chassis Topology page, marked with a pencil icon.
5. Click **Deploy** to deploy the changes that you made to the VCF.

Removing Spine or Leaf Devices

Follow these steps to remove a spine or leaf device from the VCF. You cannot remove the Designated Master device.

1. Mouse over the leaf or spine device that you want to replace and click . Network Director prompts you to confirm the deletion.
2. Click **Yes** to confirm or **No** to cancel the removal.
3. Click **Deploy** to deploy the changes that you made to the VCF.

Configuring Software Package for Upgrade

You can configure a VCF to automatically upgrade new devices that are added to the VCF to a specified software package. For this functionality to work, you must specify the location of the software upgrade image files. Use the Specify Software Package page to specify the location of the image files for each device type that is included in the VCF.

To configure software packages for upgrading VCF devices:

1. Click **Specify Software Package**. The Specify Software Package window opens.
2. To configure software image upgrades on the VCF, do the following for each device type listed on the page:
 - a. Select the location of the software image by selecting a radio button:
 - Select **FTP Server** if the software image file is located on an FTP server.
 - Select **HTTP Server** if the software image file is located on an HTTP server.
 - Select **Device** if the software image file is located on the device.
 - b. Enter the URL where the software image file is located in the text box:
 - If you selected **FTP Server** or **HTTP Server**, the text box is labeled Software Package URL. Enter the URL where the software image file is located on the server. Begin the URL with **ftp://** for an FTP server or **http://** for an HTTP server.
 - If you selected **Device**, the text box is labeled Software Package Pathname. Enter the pathname where the software image file is located on the device. Begin the pathname with a slash (/) to indicate the root directory.
3. Click **Done** to save the changes.
4. Click **Deploy** to deploy the changes that you made to the VCF.

Deploying VCF Configuration Changes

When you deploy the modified VCF configuration, Network Director performs these major steps:

1. Deletes the VCP ports that are connected to the members that are selected for deletion.
2. Recycles the deleted member.
3. Deploys the configuration to the VCF.
4. Restarts all newly added managed VCF devices in factory-default mode.

You can see the status of the deployment job. See [“Managing Configuration Deployment Jobs” on page 1194](#) for information. If deployment to the VCF Master device fails, the entire VCF deployment is cancelled.

RELATED DOCUMENTATION

[Setting Up Virtual Chassis Fabrics | 766](#)

[Understanding Build Mode in Network Director | 184](#)

[Network Director Documentation home page](#)

Understanding QFabric System Setup in Network Director

The QFabric Setup task allows you to set up these components of a QFabric system:

- Node aliases—Aliases replace the hardware serial numbers of components, making it easier to identify system devices and simplify configuration tasks.
- Node groups—Node groups help you combine multiple Node devices into a single virtual entity within the QFabric system to enable redundancy and scalability at the edge of the data center. There are three types of Node groups in a QFabric system:
 - Automatically generated server Node groups—By default, every Node device that joins the QFabric system is placed within an automatically generated server Node group that contains one Node device (the device itself). Server Node groups connect to servers and storage devices.
 - Network Node groups—You can assign up to eight Node devices to a network Node group. When grouped together, the Node devices within a network Node group connect to other routers running routing protocols such as OSPF and BGP.
 - Redundant server Node groups—You can assign two Node devices to a redundant server Node group. When grouped together, you can create link aggregation groups (LAGs) that span the interfaces on both Node devices to provide resiliency and redundancy.
- Control plane network—The control plane network is an out-of-band Gigabit Ethernet management network that connects all QFabric system components. The EX Series switches and Virtual Chassis that comprise this network are called Control Plane Ethernet (CPE) switches. The control plane network connects the Director group to the management ports of the Node and Interconnect devices.

In Network Director, you discover the control plane switches or Virtual Chassis as separate devices. Then you use QFabric setup to map the discovered switches or Virtual Chassis to the QFabric to which they are acting as CPE switches.

RELATED DOCUMENTATION

[Setting Up QFabric Systems | 780](#)

[Network Director Documentation home page](#)

Setting Up QFabric Systems

IN THIS SECTION

- [Managing Node Aliases | 780](#)
- [Managing Node Groups | 781](#)
- [Identifying CPE Switches | 782](#)
- [Reviewing the QFabric Summary and Deploying Changes | 782](#)

This topic describes how to set up a QFabric using Network Director.

To start setting up a QFabric:

1. Select **Build** from the Network Director banner.
2. Select the QFabric to set up in the View pane.
3. Select **Setup QFabric** from the Device Management section of the Task pane.

The Node Alias page of the Setup QFabric wizard opens.

This topic describes:

Managing Node Aliases

Use the Node Alias page to manage aliases for QFabric nodes.

For a description of the information shown in the Manage Node Alias table, see [Table 166](#).

To assign aliases to nodes:

1. To change one node's alias, click its **Alias (To edit, click individual cell)** table cell, then enter the new alias.

2. To assign automatic aliases to a group of nodes:

a. Select the nodes from the list by selecting their check boxes.

b. Click **Auto-alias selected devices**.

The Auto-alias selected devices window opens.

c. Enter auto-aliasing parameters in the window:

The Selected Devices field shows the number of selected devices. All selected devices will be auto-aliased.

- **Node Name**—Enter a string to use as a prefix in the automatically generated node names.
- **Starts With**—Enter a number. This number is appended to the alias prefix of the first node that is automatically aliased. The number is incremented in the alias of the remaining selected nodes.

The New Aliases field shows a preview of the new aliases as you change the parameters.

d. Click **OK**.

Aliases are automatically assigned to the selected nodes.

3. Click **Next** to continue the wizard.

The Node Groups page opens.

Table 166: Manage Node Alias Table

Column	Description
Device Serial Number	Serial number of the node.
Alias (To edit, click individual cell)	Alias of the node. Click the alias name to edit it.
Device Type	Device type of the node.
Connection State	Connection state of the node.

Managing Node Groups

Use the Node Groups page to manage QFabric node groups.

This step of the Setup QFabric wizard is the same as the separate Manage Node Groups task. You can manage node groups by using either method. For information about managing node groups, see [“Creating and Managing Node Groups for a QFabric System” on page 1252](#).

Click **Next** to continue the wizard.

The Selected CPE Switch(es) page opens.

Identifying CPE Switches

Use the Identify CPE Switch(es) page to identify the switches that comprise the QFabric system control plane network.

To identify the control plane switches:

1. To add a control plane switch
 - a. Click **Add**. The Please Select Devices window opens.
 - b. Select the control plane switches from the list.
 - c. Click **OK**.
2. To remove a control plane switch, select it from the list, then click **Remove**.
3. Click **Next** to continue the wizard.

The Review page opens.

Reviewing the QFabric Summary and Deploying Changes

In the Review page, review the QFabric setup with the changes you made. Expand the folders under the QFabric icon to see the fabric's components. To make any changes, click **Back** to return to a previous step, or click the step in the wizard flowchart at the top of the page.

To deploy the changes you made to the QFabric system click **Deploy**.

RELATED DOCUMENTATION

[Understanding QFabric System Setup in Network Director | 779](#)
[Network Director Documentation home page](#)

Configuring Cloud-Based Datacenter Networks

IN THIS CHAPTER

- Understanding Cloud Networking | 784
- Understanding Virtual Network Management | 785
- Using OpenStack with VMware NSX | 787
- Understanding the Build Mode Tasks Pane for Datacenter View | 788
- Cloud Infrastructure Requirements | 791
- Creating Data Centers Using Network Director | 793
- Managing Cloud Infrastructure | 798
- Viewing the Virtual Machine Inventory in a Cloud Infrastructure | 803
- Viewing Overlay Networks | 807
- Viewing Virtual Tunnel End Point (VTEP) Details | 808
- Managing IP Connectivity | 809
- Viewing Data Center Connectivity | 812
- Viewing Bare Metal Server Details | 818
- Managing the Virtual Switch Inventory | 819
- Viewing the Hypervisor Servers in a Data Center | 821
- Managing Network Adapter Associations | 822

Understanding Cloud Networking

Cloud computing is a popular mode of computing in which large groups of remote servers are interconnected to enable centralized data storage, and online access to network resources and services. The primary focus of cloud computing is resource sharing. Cloud resources shared by multiple users and are also allocated dynamically on demand. For example, using cloud computing you can serve a specific application, say e-mail, during the Asia Pacific business hours and later allocate the same application to serve the North American business hours. This mode of networking can thus reduce the load on the servers, save electricity, rack space, expenses, and most importantly, reduce damage to the environment.

Cloud computing can be classified into public clouds, private clouds, and hybrids clouds; depending on the type and sensitivity of data that each cloud handles.

Public clouds provide data services over a network that is open for public use. Public cloud usage can be free or be on a pay-per-usage model.

Private clouds provide data services for a single organization. As the name implies, the data in a private cloud is accessible only to authorized users of an organization.

Hybrid clouds are a combination of private and public clouds. For example, your organization might opt for a hybrid cloud if you have business-critical data that needs to be on a private cloud while some of this data must be used by a business intelligence application that is hosted on a public cloud.

There are various computing components that work together to implement a cloud computing network. Networking is one of the important components that enable cloud computing. With the advent of server virtualization, virtual and physical networks work in tandem to provide best-in-class virtualization and cloud services. To leverage the benefits of cloud computing, data centers are increasingly using cloud computing. Cloud data center focuses on providing data center connectivity from various external networks to services located within a multitenant data center. A cloud data center typically has more challenging business requirements than a traditional data center, which services only a single entity. Multitenancy requires high security, scalability, and performance. You can discover, build, manage, and monitor your cloud data center from Network Director, using the Datacenter View. Network Director supports four types of cloud infrastructure providers—VMware vCenter, VMware vCenter with NSX, OpenStack, and OpenStack with NSX plug-in.

VMware vCenter servers provide centralized management for ESXi based virtual machines. These networks typically consist of distributed and standardized virtual switches that enables virtual networking between the virtual machines. You can also add VMware NSX to a vCenter based network to extend virtualization technologies across your physical data center network. VMware NSX uses VMware's Software Defined Data Center (SDDC) architecture that enables you to create and manage software-based virtual networks. You can use Network Director to discover, manage, visualize, and monitor VMware vCenter networks and VMware vCenter with NSX networks. To know more about the VMware vCenter and the VMware NSX, refer to the VMware documentation.

OpenStack is an open-source platform used for cloud computing. OpenStack deals with all the aspects of the cloud including computing, storage, networking, and security. You can use the networking component of OpenStack to manage networks and thereby ensure that the network never becomes a bottleneck in a cloud deployment. You can create your own networks and control traffic flow, using OpenStack.

If OpenStack is your infrastructure provider, you can use the networking capabilities of OpenStack or can have another application to do this. One approach is to use VMware NSX as the networking provider for OpenStack. VMware NSX is a network virtualization product from VMware. NSX creates an abstraction layer consisting of VXLANs over your physical network, thereby virtualizing the entire physical network. When VMware NSX is used as the networking provider, tasks such as creating the network, attaching the network, creating the router, and so on, that OpenStack initiates, are redirected to VMware NSX. VMware NSX creates equivalent logical entities that can then be managed from Network Director. NSX does this without making modifications to the underlying physical network.

Network Director unifies the physical and virtual networks, that use one of the above cloud service providers, to provide network operators with a comprehensive view of the complete end-to-end network infrastructure.

You can use the Datacenter View in Network Director to view and manage data centers that are deployed by using cloud infrastructure, bare metal server (BMS), or both.

RELATED DOCUMENTATION

[Understanding Virtual Network Management | 785](#)

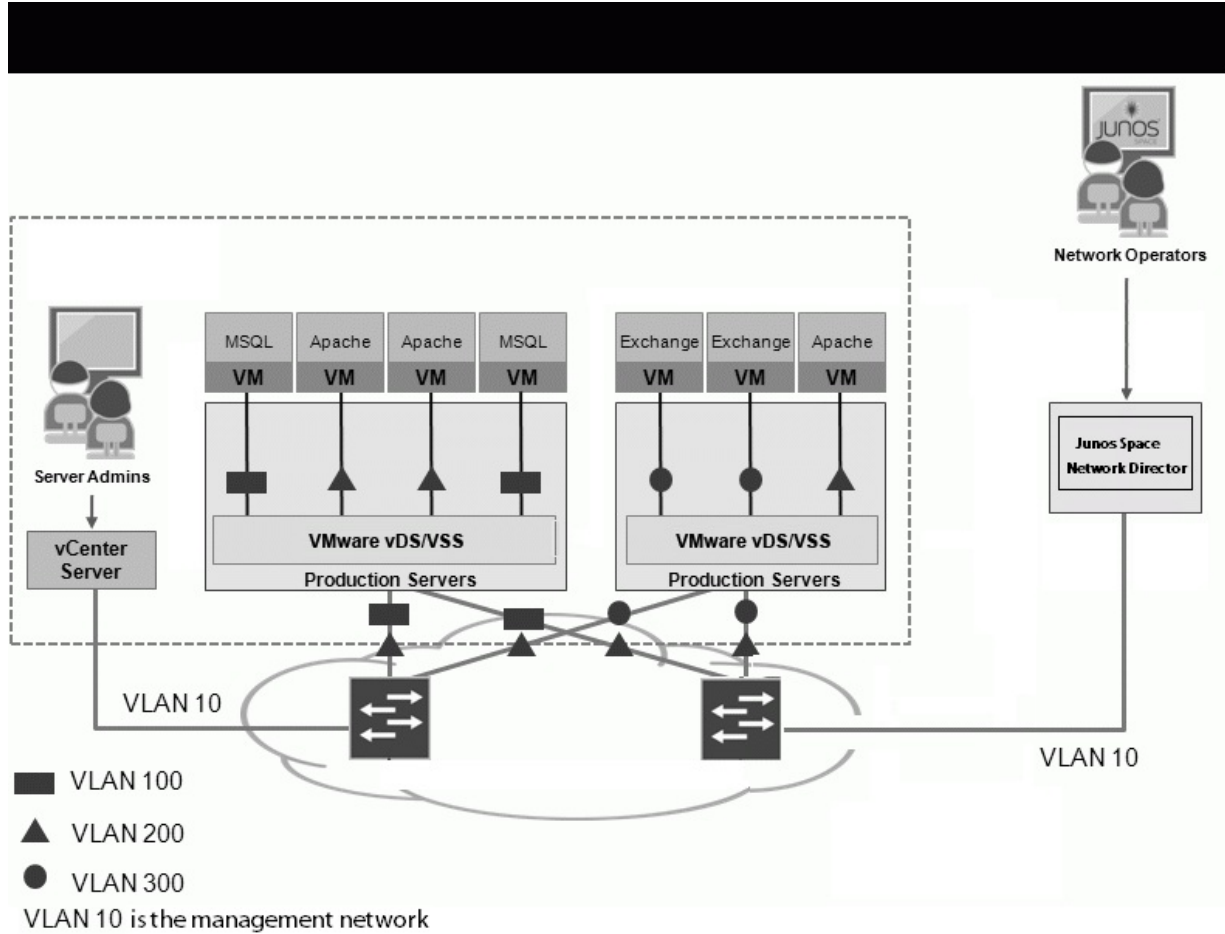
[Using OpenStack with VMware NSX | 787](#)

[Network Director Documentation home page](#)

Understanding Virtual Network Management

Network Director unifies the physical, wireless, and virtual networks and provides network operators with a comprehensive view of the complete end-to-end network infrastructure. You can use the Datacenter View perspective in Network Director to view and manage virtual networks deployed in virtualized environments in data centers. This ensures that network policies are consistently and automatically applied across physical, wireless, and virtual networks. [Figure 27](#) illustrates the virtual machines (VMs) and physical interfaces in a virtual network.

Figure 27: Network View of Virtual Machines



VMs residing on a host connect to ports in a virtual switch. Virtual switches, in turn, are associated with port groups that have a fixed number of ports. Ethernet adapters (also called *uplink adapters*) connect the virtual environment to the physical network. These elements constitute the inventory of the virtual network.

You can use Network Director to discover and monitor the virtual resources in your virtual environment. Network Director also provides consistent orchestration and operation of the physical and virtual components of the environment.

Network Director enables you to use the services and functionality of other Junos Space applications to manage and monitor the virtual network just as you would do with a physical network.

RELATED DOCUMENTATION

[Creating Data Centers Using Network Director | 793](#)

[Managing Cloud Infrastructure | 798](#)

[Managing the Virtual Switch Inventory | 819](#)

[Viewing the Hypervisor Servers in a Data Center | 821](#)

[Managing Network Adapter Associations | 822](#)

[Viewing the Virtual Machine Inventory in a Cloud Infrastructure | 803](#)

[Network Director Documentation home page](#)

Using OpenStack with VMware NSX

You can have your cloud data center use OpenStack as the cloud infrastructure provider and VMware NSX as the networking provider. The physical network that is being used by the network virtualization software, in this case, the one comprising the Juniper Network devices, is called *underlay* and the logical networking entities created by VMware NSX are called *overlay*. Although the physical network is abstracted out by VMware NSX, there are several networking scenarios where the underlay management might be required.

Following are some such scenarios:

- Configuring the underlay to make it ready for overlay. For example, Layer 3 Fabric configuration and MTU configuration.
- Hosts or virtual machines (VMs) in overlay networks need to communicate with hosts or VMs in a non-overlay network.
- Troubleshooting networking issues in an overlay network.
- Monitoring the load of overlay traffic on physical network devices.
- Connect overlay networks across multiple data centers or locations through traditional physical networking devices as a gateway.
- Hardware or software maintenance on physical network devices.

You can create your cloud data center that uses OpenStack and VMware NSX, in Network Director and discover the networking components by using the discover data center task. Once you have discovered the OpenStack and NSX controller, you can view and manage your cloud data center and the various components from the Datacenter View in Network Director.

RELATED DOCUMENTATION

[Understanding Cloud Networking | 784](#)

[Network Director Documentation home page](#)

Understanding the Build Mode Tasks Pane for Datacenter View

Network Director enables you to perform the following tasks for devices in your data center network:

- **Datacenter Management**—Create and manage data centers by using the tasks listed in this section. [Table 167](#) describes the Datacenter Management tasks.
- **Connectivity**—Manage NIC associations and view the connectivity of a data center. [Table 169](#) describes the Connectivity tasks.
- **Inventory**—View and manage the inventory of virtual machines (VMs), overlay networks, hosts, VTEPs, and virtual switches that are part of your data center network. [Table 168](#) describes the Inventory tasks.
- **Discovery**—Refresh the topology discovery and view the status of device discovery job by using the tasks described in this section. [Table 170](#) describes the Discovery tasks.
- **Key Tasks**—Group the tasks that you perform frequently and create a list of key tasks. You can add any task from the Tasks pane to the Key Tasks list by selecting a task and clicking the plus (+) sign that appears adjacent to the task. This feature is available in Task pane irrespective of your current mode, scope, or view.

For more information about Build mode features, see “[Understanding Build Mode in Network Director](#)” on page 184.

Tables [Table 167](#) through [Table 170](#) describe the tasks that you can perform in the virtual network category, including the scope in the View pane that you must select to access the task.

Table 167: Datacenter Management Tasks

Task	Description	Scope
Setup Datacenter	Enables you to create a data center.	View: Datacenter Object: My Datacenters > Datacenter Management
Manage Network Infrastructure	Enables you to select, add, and remove network devices to a data center.	View: Datacenter Object: <i>Data center name</i> > Datacenter Management
Manage Cloud Infrastructure	Enables you to add a cloud infrastructure to a data center. You can also modify, delete, or resynchronize the cloud infrastructure. For a VMware vCenter-based cloud infrastructure, you can set the orchestration mode for the data center network.	View: Datacenter Object: <i>Data center name</i> > Datacenter Management

Table 167: Datacenter Management Tasks (continued)

Task	Description	Scope
View Cloud Discovery Status	Displays the status of cloud infrastructure discovery jobs.	View: Datacenter Object: Virtual Networks > Virtual Network Management
Rename Datacenter	Enables you to rename the selected data center.	View: Datacenter Object: <i>Data center name</i> > Datacenter Management
Delete Datacenter	Enables you to delete the selected data center. This task deletes the virtual networks that are discovered as part of the data center creation and removes the association of devices that are part of the data center. The devices itself are not deleted from Network Director.	View: Datacenter Object: <i>Data center name</i> > Datacenter Management
IP Connectivity	Enables you to configure the L3 routing protocol for a data center.	View: Datacenter Object: My Datacenters > Datacenter Management
View Datacenters	Enables you to view details about the data centers that are created in Network Director.	View: Datacenter Object: My Datacenters > Datacenter Management

Table 168: Inventory Tasks

Task	Description	Scope
View Virtual Machines	Displays details of all the VMs that are part of the selected data center. Select a VM and click Show Connectivity to view the network connectivity of the virtual machine to the physical network.	View: Datacenter Object: <i>Data center name</i> > Inventory
View Overlay Networks	Displays details of all the VXLANs that are part of the selected data center. NOTE: This task is available only for Openstack+NSX-based cloud infrastructure.	View: Datacenter Object: <i>Data center name</i> > Inventory

Table 168: Inventory Tasks (continued)

Task	Description	Scope
View Virtual Switches	Displays details of all the virtual switches that are part of the selected virtual network. NOTE: This task is available only for vCenter-based cloud infrastructure.	View: Datacenter > vCenter hostname Object: Virtual Networks > Inventory
View Hypervisor Servers	Displays details of all the hyperhosts that are part of the selected data center.	View: Datacenter Object: Data center name > Inventory
View Baremetal Servers	Displays details of the bare metal servers that are part of a data center. Network Director automatically discovers bare metal servers that are part of your data center network.	View: Datacenter Object: Baremetal Servers > Inventory
View Inventory	Displays information about all the devices in the selected object and all its child objects.	View: Datacenter Object: Network Devices > Inventory
View VTEP	Displays the virtual tunnel end points for the selected data center.	View: Datacenter Object: Data center name > Inventory

Table 169: Connectivity Tasks

Task	Description	Scope
Manage NIC Associations	Enables you to view and manage the connectivity between a virtual host and the physical network objects (network adapter on a physical switch). Manual associations that you create using this task is only used for performing orchestration. NOTE: This task is available only for vCenter-based cloud infrastructure.	View: Datacenter Object: My Datacenters > Connectivity
View Connectivity	Displays the graphical connectivity of various devices within the selected data center.	View: Datacenter Object: Data center name Connectivity
Manage IP Connectivity	Displays the autonomous systems configured on the different devices in a data center.	View: Datacenter Object: Data center name Connectivity

Table 170: Discovery Tasks

Task	Description	Scope
Refresh Topology	Refreshes the topology discovery. This task also refreshes the device connectivity.	View: Datacenter Object: My Datacenters > Discovery View: Datacenter Object: <i>Data center name</i> > Discovery
View Topology Discovery Job	Displays detailed status of the discovery jobs.	View: Datacenter Object: My Datacenters > Discovery View: Datacenter Object: <i>Data center name</i> > Discovery

RELATED DOCUMENTATION

[Creating Data Centers Using Network Director | 793](#)

[Managing Cloud Infrastructure | 798](#)

[Managing the Virtual Switch Inventory | 819](#)

[Viewing the Hypervisor Servers in a Data Center | 821](#)

[Managing Network Adapter Associations | 822](#)

[Network Director Documentation home page](#)

Cloud Infrastructure Requirements

You can create data centers in Network Director, using Juniper Network devices, cloud infrastructure, or a combination of both. Network Director supports three types of cloud infrastructure providers—VMware vCenter, OpenStack, and a combination of OpenStack and VMware NSX. Each of these cloud infrastructure providers must meet the following requirements, else Network Director will not be able to discover these devices and create the data center:

- For bare metal servers, ensure that you enable LLDP on the servers and the physical switches.
- Ensure that your cloud infrastructure meets the requirements given in [Table 171](#).

Table 171: Cloud Infrastructure Requirements

Cloud Infrastructure	Requirements
OpenStack	Supported Release—Icehouse Ensure that the APIs listed in Table 172 are running.
VMware NSX	Version 4.1 or 4.2
VMware ESXi	VMware ESX versions 4.0 or 4.1 VMware ESXi versions 5.0, 5.1, 5.5, or 6.0

Table 172: API Requirements

Name of the API	OpenStack Based data center	OpenStack + NSX Based data center
Keystone API v2	Yes	Yes
Nova API v2	Yes	Yes
Neutron API v2	No	Yes
Ceilometer API v2	Yes, if you want to run the VM stats monitoring feature	Yes, if you want to run the VM stats monitoring feature

NOTE: Network Director does not support OpenStack with Neutron-based VXLANs.

RELATED DOCUMENTATION

[Understanding Cloud Networking | 784](#)

[Network Director Documentation home page](#)

Creating Data Centers Using Network Director

IN THIS SECTION

- [Create a Data Center | 793](#)
- [View the Data Center Creation Status | 796](#)
- [Assigning Network Devices to a Data Center | 797](#)

The cloud data center focuses on providing data center connectivity from various external networks to services located within a multitenant data center. You can discover, build, manage, and monitor your cloud data center from Network Director, using the Datacenter View. Network Director supports four types of cloud infrastructure providers—VMware vCenter, VMware vCenter with NSX plug-in, OpenStack, and OpenStack with NSX plug-in.

This topic describes the following:

Create a Data Center

To create a data center:

1. While in Build mode with Datacenter View selected, from the Tasks pane, click **Setup Datacenter** from the Datacenter Management menu. The Create Datacenter page opens.
2. In the Setup Datacenter wizard page, enter a name for the data center.
3. Click **Next** to specify details of the cloud infrastructure that the data center uses.
4. In the Cloud Infrastructure wizard page, do one of the following:
 - Click **No** if you do not want to specify a cloud infrastructure now or if your data center is a controller-less data center that uses only Juniper Networks devices. Skip to step [11](#).
 - Click **Yes** if you want to specify the cloud infrastructure for your data center.
5. Select the type of cloud infrastructure provider that the data center uses from the Cloud Infrastructure Provided box. You can choose VMware vCenter, VMware vCenter with NSX, OpenStack, and OpenStack with NSX as the provider.
6. Do one of the following:
 - Enter the IP address or the hostname of the vCenter or OpenStack server.

- If you want to add a VMware vCenter that uses Platform Services Controller (PSC) location service, do the following:

NOTE: PSC, a component of the VMware Cloud Infrastructure Suite, is available in VMware vSphere Version 6.0 or later. If your server is running an older version of vSphere, this step does not apply to you.

- a. Click **Search vCenter Using PSC Location Service**. The Search vCenter Using PSC Location Service window opens.
 - b. Enter the IP address or the host name of the platform services controller.
 - c. Click **Search vCenters**. Network Director displays all the vCenters that are managed by the PSC, in the list.
 - d. Select a vCenter from the list and click **Select**. Network Director adds the vCenter that you selected to the Cloud Infrastructure wizard page.
7. Specify the port that Network Director uses to connect to the server. The default port used to connect to a vCenter server is 443 and that for an OpenStack Keystone service is 35357.

NOTE: You can modify this and specify a port of your choice. If you do so, make sure to manually change the Junos Space firewall settings and apply to this port.

8. Specify the administrator username and password for the server you selected. The username and password must match the name and password configured on the server.

NOTE: The administrator username that you specify for discovering the OpenStack server must have admin privileges and must belong to an admin tenant in the OpenStack server.

9. If you specified a vCenter server that is running VMware vSphere version 6.0 or later and would like Network Director to discover the virtual machines tags from vSphere, you must specify some additional settings. To do this, click **Additional Settings**. The vCenter Discovery Additional Settings window opens displaying the IP address of the vAPI endpoint (default value is the IP address or host name of the vCenter server) and the port (default is 443) that Network Director uses to connect to the vAPI endpoint. If you plan to use an external PSC, you must specify the IP address of the vAPI endpoint and click **Update**.

Network Director closes the vCenter Discovery Additional Settings window.

NOTE: If you specified the vCenter details by using the PSC location service as mentioned in step 6, Network Director obtains the vAPI details from the settings you specify. You need not specify the vAPI details as mentioned in this step.

vAPI endpoint is a part of vCenter (running vSphere version 6.0 or later) that exposes APIs to new functionalities such as the tagging objects. Network Director leverages this functionality to obtain and display the virtual machine tag using the vAPI.

10. If you selected VMware vCenter with NSX or OpenStack with NSX as the provider, specify the following details about the NSX controller:
 - a. Enter the IP address or the hostname of the NSX controller.
 - b. Specify the port that Network Director uses to connect to the NSX plug-in. The default port used to connect to an NSX plug-in is 443.

NOTE: You can modify this and specify a port of your choice. If you do so, make sure to manually change the Junos Space firewall settings and apply to this port.

- c. Specify the administrator username and password for the NSX server. The username and password must match the name and password configured on the server.
11. Click **Next** to specify the method that you want Network Director to use to build the data center network.
12. Select the Juniper Network devices that you want to add to the data center from the Available Devices table and click >> to add it to the Selected Devices table. You can also add data center fabrics such as a Layer 3 Fabric to your data center.

If you want to remove a device or fabric from the Selected Devices table, select the device or fabric and click <<.
13. To refresh the devices discovered and managed by Network Director, select the Refresh Topology check box. If you have not specified the SNMP details while discovering the devices or if you want to modify the SNMP details for the devices, click **Configure**. The Refresh Topology window appears.

Specify the SNMP details by following the instructions given in the *Refreshing the Topology* section of [“Managing the Topology View” on page 251](#).

14. Click **Done** to save the data center details.

A message window opens, displaying the status of the cloud infrastructure discovery job name and job ID. Click **OK**.

You can view the status of the discovery job in the Cloud Infrastructure Discovery Jobs page.

Network Director tries to discover the servers that you specified in the Cloud Infrastructure wizard page, if you specified the cloud infrastructure details. Network Director then adds the remaining network infrastructure based on the details you specified in the Network Infrastructure wizard page. Once these two steps are complete, Network Director lists that data center along with the devices that are part of the network infrastructure under My Datacenters in the View pane. You can perform various tasks on the data center by selecting the data center in the View pane.

View the Data Center Creation Status

After you have specified all the details in the Create Datacenter wizard and submitted it, you can view the discovery status from the View Cloud Discovery Status page in the Datacenter Management menu.

The Cloud Infrastructure Discovery Jobs page displays all the discovery jobs. From this page, you can view the details described in [Table 173](#).

Table 173: Cloud Infrastructure Discovery Jobs page Field Descriptions

Field	Description
Job ID	An identifier assigned to the job.
Job Name	The name of the job (user-created).
Percent	Percentage of the job that is complete.
Status	<p>Job status. The possible states are:</p> <ul style="list-style-type: none"> ● CANCELLED—The job was canceled by a user. ● FAILURE—The job failed. This state is displayed if any of the devices in the job failed; even if some of the devices might have completed successfully. View the job details for the status of each device. ● INPROGRESS—The job is running. ● SCHEDULED—The job is scheduled but has not run yet. ● SUCCESS—The job completed successfully. This state is displayed if all of the devices in the job completed successfully.

Table 173: Cloud Infrastructure Discovery Jobs page Field Descriptions (*continued*)

Field	Description
Summary	Summary of the job scheduled and executed with status. NOTE: If Network Director encounters an error while trying to retrieve the virtual machine tag, details about the error is displayed in this field.
Scheduled Start Time	The UTC time on the client computer when the job is scheduled to start.
Actual Start Time	The actual time when the job started.
End Time	The time when the job was completed.
User	The login ID of the user that initiated the job.
Recurrence	The recurrent time when the job will be restarted.

To view the details of a job, select the check box against **Job ID** or **Job Name** and click **Show Details**. The Discover Network Elements window displays details of the virtual network discovery job.

NOTE: During cloud infrastructure discovery, if Network Director is unable to read the device configurations, then the status is displayed as Failed. You can check the reason for the failure from the Manage Jobs page in System mode. You must make the required changes to the device configuration by using the CLI so that Network Director can read the configuration. Network Director automatically resynchronizes once you enable a cloud infrastructure discovery job. If Network Director cannot discover the device even after resynchronization, then you must rediscover the device after making the appropriate changes in the device configuration by using the CLI.

Assigning Network Devices to a Data Center

You can add network devices to a data center while you create a data center. Alternatively, you can skip this step while creating the data center and add network devices later by using the Assign Devices to Datacenter task.

To add network devices after you have created the data center:

1. While in Build mode with Datacenter View selected, select the data center to which you want to add devices in the View pane and click **Datacenter Management > Manage Network Infrastructure** from

the Tasks pane. The Assign Devices to Datacenter page opens, listing the network devices that are discovered in Network Director.

2. Select one or more network devices that you want to add to the data center and click **Done**.

Network Director adds the selected devices and updates the data center in the View pane.

RELATED DOCUMENTATION

[Understanding Cloud Networking | 784](#)

[Managing Cloud Infrastructure | 798](#)

[Managing Network Adapter Associations | 822](#)

[Network Director Documentation home page](#)

Managing Cloud Infrastructure

IN THIS SECTION

- [View the Cloud Infrastructure for a Data Center | 799](#)
- [Open VMware vSphere Web Client | 800](#)
- [Modify Cloud Infrastructure Details | 800](#)
- [Configure Orchestration Mode | 800](#)
- [Delete Cloud Infrastructure | 802](#)
- [Resynchronize Cloud Infrastructure | 802](#)

You can view information about all the cloud infrastructures managed by Network Director in the Manage Cloud Infrastructure page. Network Director uses information from each cloud infrastructure to create an inventory view with a detailed list of all the elements in each cloud network.

While in Build mode with Datacenter View selected, from the Tasks pane, click **Manage Cloud Infrastructure** from the Cloud Infrastructure Management menu to open the Manage Cloud Infrastructure page.

You can do the following from the Manage Cloud Infrastructure page:

View the Cloud Infrastructure for a Data Center

To view the cloud infrastructure details:

1. While in Build mode with Datacenter View selected, select a data center from the Datacenter View pane and from the Tasks pane, click **Manage Cloud Infrastructure** from the Datacenter Management menu. The Manage Cloud Infrastructure page opens.
2. [Table 174](#) describes the fields that are displayed in the Manage Cloud Infrastructure page.

Table 174: Manage Cloud Infrastructure page Field Descriptions

Field	Description
Name	Name of the cloud infrastructure provider.
IP Address/Host Name	Hostname or the IP address of the cloud infrastructure.
Port	Port that Network Director uses to connect to the cloud infrastructure.
Type	Type of cloud infrastructure. Type can be VMware vCenter, OpenStack, or OpenStack NSX.
Version	Version of software that is running on the cloud infrastructure server,
Connection Status	Current status of the connection between Network Director and the cloud infrastructure.
Sync Status	Status of the previous synchronization job. Click this field to view details about the previous synchronization job in the Synchronization Status window.
Orchestration Mode	The orchestration mode that is set for the vCenter network. Orchestration mode can be Enable Orchestration, Disable Orchestration and Remove Existing, or Disable Orchestration. Orchestration mode is applicable only for cloud infrastructures of type vCenter. For any other type of cloud infrastructure, this field displays <i>Not Applicable</i> .
Last Orchestration	Status of the previous orchestration job. Click this field to view details about the previous orchestration job in the Orchestration Status window.

Open VMware vSphere Web Client

You can open the VMware vSphere Web Client from Network Director. The vSphere Web Client enables you to connect to a vCenter Server to manage one or more ESXi hosts through your browser. The vSphere Web Client opens in a new tab or window depending on your browser settings.

To open the Web Client for a vCenter, select the vCenter server from the View pane and click **Inventory >Launch vSphere Web View** from the Tasks pane.

Network Director opens the vSphere Web Client in a separate tab or window.

Modify Cloud Infrastructure Details

You can modify some of the details of an existing cloud infrastructure.

To modify credentials for a cloud infrastructure:

1. Select the cloud infrastructure that you want to modify and click **Edit**.

The Edit Cloud Infrastructure dialog box is displayed.

2. Modify the necessary fields. You can modify the port, username, and password for each cloud service provider.

3. Click **Modify** to submit the changes that you made.

After you click Modify, Network Director verifies if the server is accessible using the modified details. If the connection is not successful, then Network Director does not save the details.

Configure Orchestration Mode

Orchestration applies aggregated VLAN configurations of the required port groups to the appropriate ports of the physical switch including the link aggregation group (LAG) ports. Using Network Director, you can seamlessly orchestrate across physical, virtual, and cloud infrastructure elements.

Each virtual switch, depending on the switch type, can span across one or more hosts and is configured to have a minimum of one uplink port for each host. A distributed virtual switch can span multiple hosts, whereas a standalone virtual switch is associated with only one host. A virtual switch comprises one or more port groups, each of which is a collection of ports. Each of these ports can, in turn, be associated with a virtual machine (VM) or host kernel. Traffic coming from these ports is sent out of one of the uplink ports of the virtual switch, which is configured on the same host on which the virtual machine (VM) or kernel resides.

Each physical NIC in a host is connected to an access port in the physical switch. Any configuration or restrictions that needs to be applied to the physical NIC to manage traffic is applied to the access port of the physical switch.

NOTE: You cannot configure orchestration for cloud infrastructures that use OpenStack or a combination of OpenStack and NSX.

To set the orchestration mode operable on a cloud infrastructure:

1. Select a cloud infrastructure from the Manage Cloud Infrastructure page and click **Configure Orchestration**.

The Orchestration Mode dialog box is displayed.

2. Select one of the following orchestration mode for the selected cloud infrastructure:
 - **Enable Orchestration**—Network Director performs VLAN orchestration on the physical switch ports automatically.
 - **Disable Orchestration**—Disables orchestration, but retains the current orchestrations. No new orchestrations will be done.
 - **Disable Orchestration and Remove Existing**—Network Director removes all the existing orchestrations and no new orchestrations will be done.

NOTE: While modifying the orchestration mode, Network Director enables you to change the status from *Enable Orchestration* to *Disable Orchestration and Remove Existing*, and not from *Disable Orchestration* to *Disable Orchestration and Remove Existing*. However, you can change the status from *Disable Orchestration and Remove Existing* to either *Enable Orchestration* or to *Disable Orchestration*.

3. Click **OK** to set the cloud infrastructure to the orchestration mode you prefer.

Network Director displays a message to indicate the successful modification of the orchestration mode.

Network Director tracks the status of orchestration by using an orchestration job. After an orchestration job is completed successfully, you must manually resynchronize the physical switch's configuration with Network Director. If the system of record (SOR) mode set for the Junos Space Network Management Platform is:

- Network as system of record (NSOR), then performing a resynchronization ensures that Junos Space automatically resynchronizes its configuration record to match the device configuration and sets the

device configuration state to In Sync when the synchronization completes. For more details, see [“Resynchronizing Devices When Junos Space Is in NSOR Mode” on page 1222](#).

- Junos Space as system of record (SSOR), then you must perform a resynchronization and accept the out-of-band changes. Both the Junos Space configuration record and the Network Director Build mode configuration are resynchronized to reflect the out-of-band configuration changes. For more details, see [“Resynchronizing Devices When Junos Space Is in SSOR Mode” on page 1223](#)

Delete Cloud Infrastructure

You can dissociate a virtual network from Network Director by deleting it from the system.

To delete a virtual network:

1. Select a cloud infrastructure from the Manage Cloud Infrastructure page and click **Delete**.
Network Director displays the Delete Cloud Infrastructure dialog box.
2. In the Delete Cloud Infrastructure dialog box, click **Delete** to delete the cloud infrastructure from Network Director.

NOTE: You can add a new cloud infrastructure to a data center only after you delete the existing cloud infrastructure.

Resynchronize Cloud Infrastructure

You can resynchronize a cloud infrastructure registered with Network Director. During resynchronization, the cloud infrastructure configuration replaces the corresponding configuration available in Network Director. Any changes that you might have made to the cloud infrastructure by using Network Director might be lost after a resynchronization.

Synchronization is run as a job and is managed by the Job Manager in Network Director. Changes that you make to a vCenter network are dynamically updated in Network Director. However, changes that are made on an Openstack network requires you to wait for the periodic synchronization job to run or you must perform a manual resynchronization for the changes to be updated in Network Director.

To initiate a manual resynchronization:

1. Select a cloud infrastructure from the Manage Cloud Infrastructure page and click **Resynchronize**.
The Status dialog box opens, showing the corresponding Job ID.

2. In the **Job Information** dialog box, click the job ID link.

The **Job Details** page displays details about the resynchronization job.

RELATED DOCUMENTATION

[Understanding Cloud Networking | 784](#)

[Understanding Virtual Network Management | 785](#)

[Creating Data Centers Using Network Director | 793](#)

[Network Director Documentation home page](#)

Viewing the Virtual Machine Inventory in a Cloud Infrastructure

IN THIS SECTION

- [View the Virtual Machine Inventory | 803](#)
- [View Connections Between VMs and the Physical Network | 804](#)

You might have many virtual machines (VMs) configured on a single host machine. Once you have discovered a cloud infrastructure in Network Director, the hosts, VMs, and virtual switches also come under the Network Director management. You can use the View Virtual Machine inventory task to view details of VMs.

You can also view the network connectivity of a VM with the host, virtual switch, and physical switch from the View Virtual Machine page.

View the Virtual Machine Inventory

To view the VM Inventory:

1. While in Build mode with Datacenter View selected, do one of the following:
 - Click **Inventory > View Virtual Machines** from the Tasks pane to view details of all the VMs in your virtual network.
 - Select a hypervisor server in the View pane and click **Inventory > View Virtual Machines** from the Tasks pane to view details of VMs that belong to the selected server.

- Select an overlay network in the View Overlay Networks page and click **Show Virtual Machines**.

The View Virtual Machine page opens displaying the virtual machine inventory table.

2. The View Virtual Machine page displays the details for all the virtual machines. [Table 175](#) describes the fields in this table.

Table 175: View Virtual Machines Page Field Descriptions

Field	Description
Name	Name assigned to the VM.
IP Address	IP address of the VM.
Datacenter Name	Name of the data center to which the VM belongs.
MAC Address	The MAC address of the VM.
Host Name	Name of the host on which the VM is running.
Power Status	Indicates if the VM is powered on or off.
Tenant	Name of the tenant that uses VM.
Guest Operating System	The operating system that is installed and running on the VM.
Tag Names	A label that is assigned to a virtual machine for easy identification. You can specify tags for virtual machines by using the vSphere user interface. NOTE: Tags are available only if your vCenter is running vSphere Version 6.0 or later.

3. To view the network connectivity of a VM, select the VM and click **View Connectivity**.

The Virtual Network Connectivity page opens, displaying the network connectivity and VLAN associations for the selected VM. For more details, see [“View Connections Between VMs and the Physical Network” on page 804](#).

View Connections Between VMs and the Physical Network

For a given virtual machine, you can pictorially view the connections between the VM and the physical switch, using the Virtual Network Connectivity page. This view provides you a host of details that can be useful for monitoring the health and utilization of your virtual network and viewing the connection status, IP addresses and host names of the various virtual and physical devices in your network. You can mouse

over the link from each virtual machine to view details of the association, including the virtual adapter name, port group name, virtual switch IP address, and so on.

To view the connections between the VMs and the physical network:

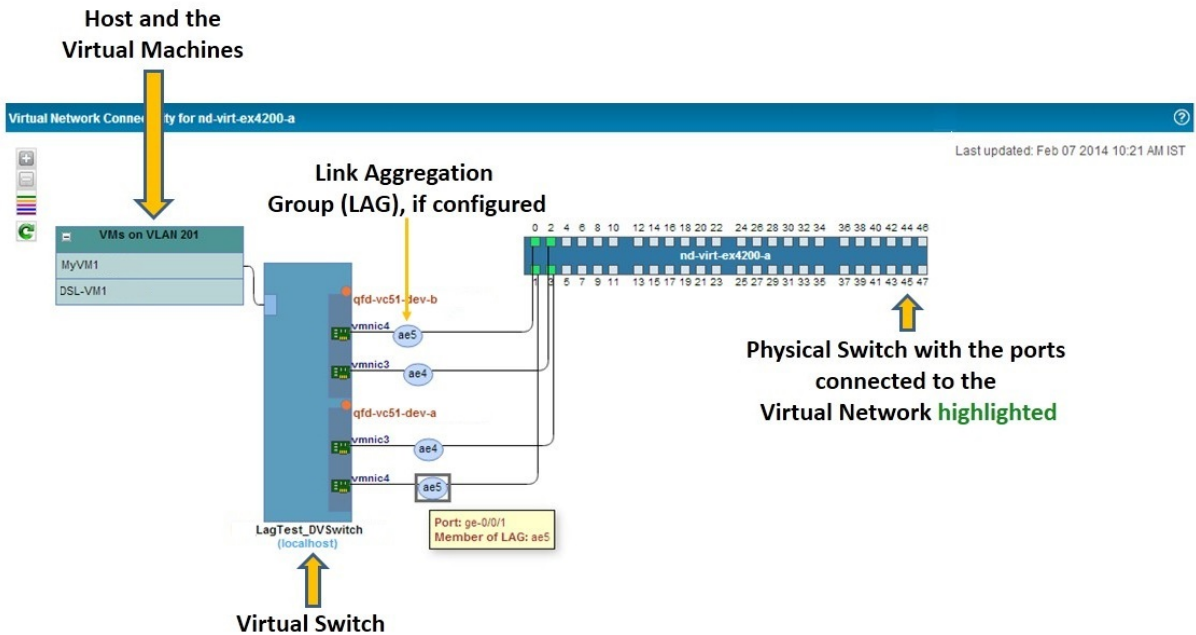
1. Do one of the following:

- From the View Virtual Machines page, select the virtual machine for which you want to view the connectivity and click **View Connectivity**.
- While in Build mode with Logical or Datacenter View selected, select a network device or a switch from the View pane and click **Connectivity > View Virtual Network Connectivity** from the Tasks pane.

Network Director displays the virtual network connectivity from the virtual machine to the network device or the switch in a new window.

Network Director displays the connectivity between the virtual machines to the physical network as shown in [Figure 28](#).





Figure 28: View Cloud Infrastructure Connectivity




From the Connectivity page, you can:

- Click the virtual NIC on a virtual machine to highlight the connection between the virtual machine and the physical switch. This connection is shown as traversing the connected virtual switch or a virtual distributed switch and the host before reaching the physical switch.
- Click the physical NIC on the switch to highlight all the connections between the selected physical NIC and the virtual machines. The enroute host and virtual switch or a virtual distributed switch are also displayed.
- View details about link aggregation groups, if configured. If you mouse over a LAG, Network Director displays the port and the LAG interface.
- Mouse over the physical NIC, host, or a virtual machine to view the host name and the bandwidth utilization for the selected device. In addition to these details, the physical NIC and the host details also include the number of active alarms on that device. A colored dot that appears on the upper right corner of the physical switch or the host indicates that there are alarms on the device—the color of the dot identifies the severity level of the alarm. See [Table 176](#) to know more about the alarm severity indicator for each alarm severity.

Table 176: Alarm Severity Indicator

Alarm Severity	Indicator
Critical	
Major	
Minor	
Informational	

NOTE: Devices might have alarms of different severities. Network Director displays the indicator based on the most severe alarm on a device. For example, if a device has 3 informational alarms and one major alarm, Network Director displays the indicator for the major alarm ().

RELATED DOCUMENTATION

[Understanding Cloud Networking | 784](#)

Viewing Overlay Networks

Virtual Extensible Local Area Network (VXLAN) represents a technology that enables you to segment your networks (as VLANs do) but that also solves the scaling limitation of VLANs and provides benefits that VLANs cannot. VXLAN is often described as an overlay technology because it enables you to stretch Layer 2 connections over an intervening Layer 3 network by encapsulating (tunneling) Ethernet frames in a VXLAN packet that includes IP addresses.

The video *Why Use an Overlay Network in a Data Center?* presents a brief overview of the advantages of using VXLANs.



Video: [Why Use an Overlay Network in a Data Center?](#)

You can view the overlay networks in a data center by using the view overlay networks task. You can also view all the virtual machines (VMs) that are part of an overlay network using this task.

To view the overlay networks in a data center:

1. While in Build mode with Datacenter View selected, select a data center from the View pane and click **Inventory > View Overlay Networks** from the Tasks pane.

The View Overlay Networks page opens.

2. [Table 177](#) describes the fields that are displayed in the View Overlay Networks page.

Table 177: View Overlay Network page Field Descriptions

Field	Description
Tenant	Name of the tenant that uses the given overlay network.
Network Name	Name of the overlay network.
Tunnel Type	Type of overlay tunnel. Network Director supports only VXLAN tunnels.
Tunnel ID	Unique identification number of the overlay network.
Subnet	Subnet assigned to the tenant.
Gateway	Gateway used by the overlay network

Table 177: View Overlay Network page Field Descriptions (continued)

Field	Description
DHCP	Indicates whether DHCP is enabled or disabled for the overlay network.

- To view the VMs that are part of an overlay network, select an overlay network and click **Show Virtual Machines**.

The View Virtual Machines page opens displaying all the VMs. For detailed descriptions on the fields in this page, see “[Viewing the Virtual Machine Inventory in a Cloud Infrastructure](#)” on page 803.

RELATED DOCUMENTATION

[Understanding Cloud Networking | 784](#)

[Network Director Documentation home page](#)

Viewing Virtual Tunnel End Point (VTEP) Details

Overlay networks based on VXLANs transport frames after encapsulating them as VXLAN packets. Encapsulation and de-encapsulation in these networks are done by an entity called the Virtual Tunnel End Point (VTEP). VTEPs can be implemented in your overlay network as—virtual bridges in a hypervisor server, VXLAN-specific virtual applications, or switching hardware that is capable of handling VXLANs.

After you set up a data center, you can view details about the VTEPs in your network, using the view VTEP task.

To view the VTEPs in a data center:

- While in Build mode with Datacenter View selected, select a data center from the View pane and click **Inventory > View VTEP** from the Tasks pane.

The View VTEP page opens.

- [Table 178](#) describes the fields that are displayed in the VTEP table.

Table 178: VTEP page Field Descriptions

Field	Description
VTEP IP	Name of the tenant that uses the given overlay network.

Table 178: VTEP page Field Descriptions (*continued*)

Field	Description
VTEP Hostname	Name of the overlay network.
VTEP Host Type	Type of overlay tunnel. Network Director supports only VXLAN tunnels.
Switch Ports	Unique identification number of the overlay network.

- To view details about the overlay network, select a VTEP entry from the table and click **View Overlay Networks**.

The View Overlay Networks page opens displaying all the overlay networks that are using the selected VTEP. For detailed descriptions of the fields in the View Overlay Networks page, see “[Viewing Overlay Networks](#)” on page 807.

RELATED DOCUMENTATION

[Understanding Cloud Networking | 784](#)

[Network Director Documentation home page](#)

Managing IP Connectivity

IN THIS SECTION

- [Adding an Autonomous System | 810](#)
- [Adding Devices | 810](#)
- [Creating Links | 811](#)

Network Director enables you to configure Layer 3 (L3) routing protocol, BGP.

The IP Connectivity task displays all the autonomous systems configured in the different devices in a data center. When you mouse over the device, the autonomous system in which the device is added is displayed. You can also view the details of the device such as the IP address, platform, connection state, serial number, and action that can be performed on the device (double-click to add or modify the router reflector). When

you double-click an autonomous system, the Configure Autonomous System and Policy window is displayed, where you can configure the import and export policy for the autonomous system. When you mouse over the autonomous system, a cross mark (✖) is displayed. Clicking on this deletes the autonomous system.

All the devices within an autonomous system are Internal BGP (IBGP) peers and a logical full mesh formed between these IBGP peers are displayed. These logical links are represented by a line. Similarly, all the devices across different autonomous system are External BGP (EBGP) peers. When you mouse a device, a cross mark (✖) is displayed. Clicking on this deletes the device from the autonomous system.

Adding an Autonomous System

To add an autonomous system:

1. Select a data center from the Datacenter View pane.
2. Click **Manage IP Connectivity** under Tasks.
3. Click **Add Autonomous System**.

The Configure Autonomous System and Policy window is displayed.

4. Enter a number to identify the autonomous system.

A circle appears in the viewport showing the number you specified.

5. Double-click within the circle.

The Configure Autonomous System and Policy window is displayed.

6. Click **Import** and click **Export** to import or to export the policy for the IBGP peers within the autonomous system.
7. Click **OK**.

Adding Devices

To add devices to the autonomous systems:

1. Click **Add Devices**.
2. All the devices in a data center are listed in a **Select Devices** device panel on the left pane.

3. Drag and drop the devices from this panel to autonomous system or to the canvas directly (outside of all autonomous systems).
4. Click **Done**.

Creating Links

To create a link between two devices:

1. Click the **Add EBGP/IP Links** button, and then click on one of the devices within the autonomous system between which you want to create the link.

The pointer changes to a drawing cursor that enables you to draw a link between devices.

2. Draw a link from the device to the device to which you want to create the link.

The Link Details window is displayed.

3. You can create links between EBGP peers, IBGP peers, or between a device inside an autonomous system and a device outside of autonomous system:

- Enter the interface settings such as interface type, logical unit, and IP address, and click **Done**.
- For EBGP peers, configure both the physical link and the routing policies and click **Done**. When you draw a link between two devices across different autonomous systems, the devices are treated as EBGP peers.
- For IBGP peers, double-click the link and configure the interface settings for the link and the routing policies for each device in the link end points, and click **Done**. IBGP requires a full logical mesh linking all the peers. When you drag and drop a device to add it to the autonomous system, logical links are automatically drawn between the newly added device and all the other devices in the autonomous system.

In case of large networks, you cannot form full logical mesh between IBGP peers. Therefore, you can configure some of the devices as router reflectors. These devices have the capability to readvertise the IBGP packets to their clients. In an autonomous system, a logical full mesh forms between the router reflector and non-clients. When you mouse over a device in an autonomous system and if the device can be configured as a router reflector, router reflection configuration options are displayed for that device.

To add a route reflector:

1. Double-click a device that has router reflection configuration options.

The Configure Route Reflector window appears.

2. Select **Route Reflector**.

3. Select the clients.
4. Click **OK**.

You can click **Deploy** and then click **Run Now** to immediately deploy the configuration. Or, you can click **Deploy** and then click **Run Later** to schedule the deployment.

NOTE: You can discard changes only for the entire data center and not for each device.

RELATED DOCUMENTATION

| [Understanding the Build Mode Tasks Pane for Datacenter View | 788](#)

Viewing Data Center Connectivity

After you have set up a data center in Network Director, you can pictorially view the physical connectivity between the network devices and the bare metal server, the hypervisor server, or both, using the View Connectivity task. This view also displays all the virtual machines (VMs) that are part of each data center.

To view the connectivity between the physical and virtual devices in a data center:

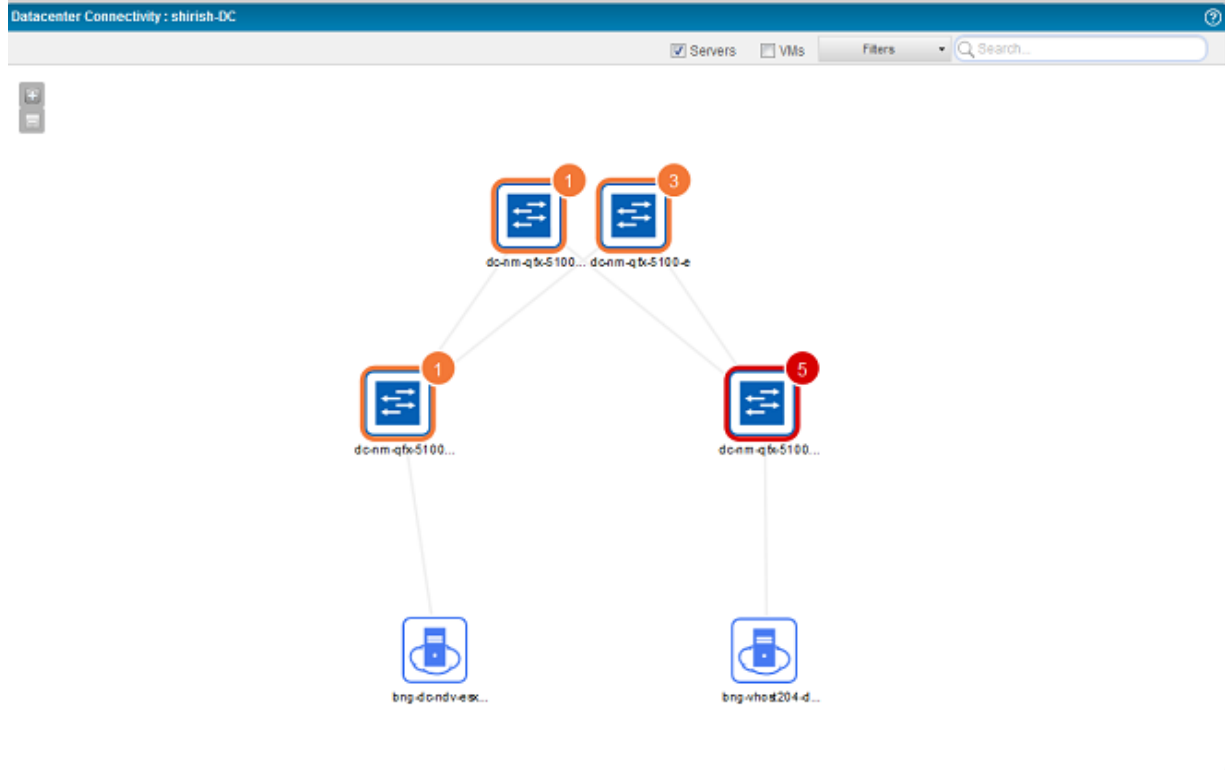
1. While in Build mode with Datacenter View selected, select a data center from the View pane and click **Connectivity > View Connectivity** from the Tasks pane.

The Datacenter Connectivity page opens displaying the connections between the network devices and the data center servers (hypervisor servers and bare metal servers).

2. You can perform the following tasks from the Datacenter Connectivity page.

- View the connectivity between switches, between switches and hosts, and between hosts and virtual machines. When you open the Datacenter Connectivity page, Network Director displays the connectivity between switches and switches to hosts as shown in [Figure 29](#).

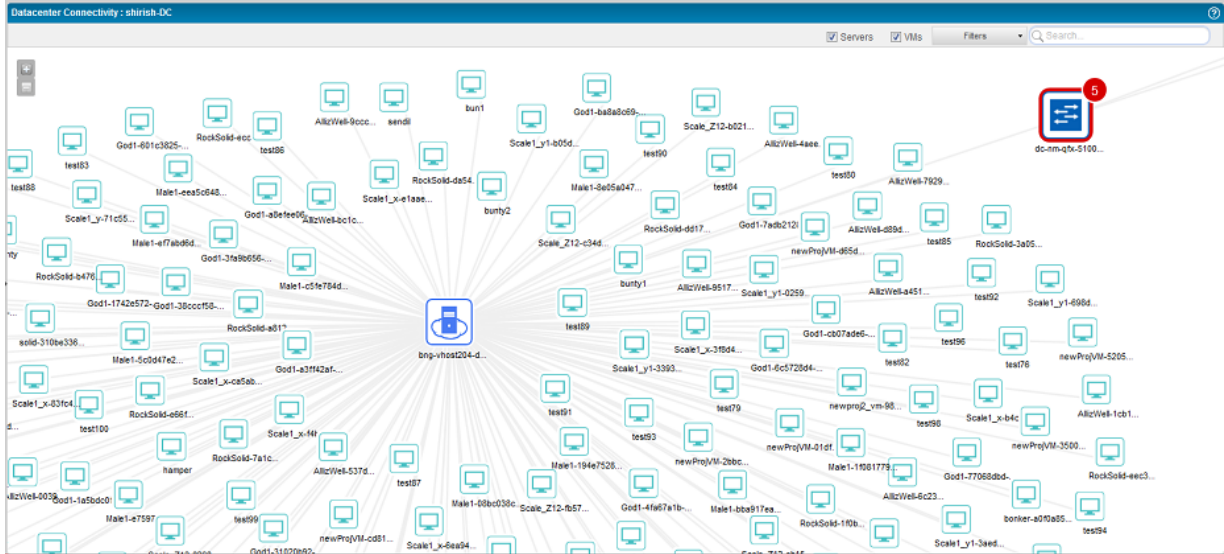
Figure 29: Datacenter Connectivity View—Switch-to-Host Connectivity



You can view port utilization between switches, and switches and host servers by selecting the **Color Coded Port Utilization** check box. Use the legend displayed on this page to infer the port utilization on each link. You can also mouse over each link to know the exact port utilization on that link. [Figure 29](#) displays the use of color coded port utilization legend.

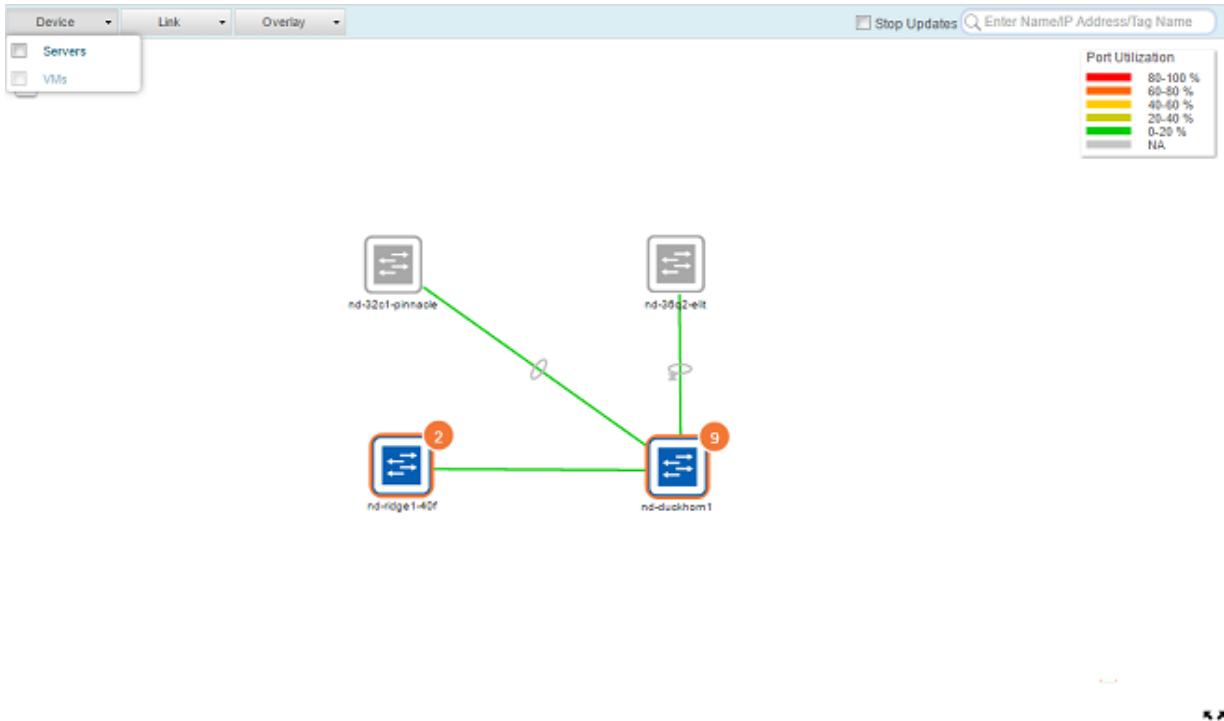
To view the connectivity between a host and virtual machines, click the host or click the **Device** button and select **VMs** check box as shown in [Figure 30](#).

Figure 30: Datacenter Connectivity View—Host-to-Virtual Machine Connectivity



To view only the connectivity between the switches, clear the **Servers** check box and the **VMs** check box as shown in [Figure 31](#).








Figure 31: Datacenter Connectivity View—Switch-to-Switch Connectivity



NOTE: Manual NIC associations are used for orchestration purpose only, and are not visible in the Datacenter Connectivity page.

- Disable the topology updates to the current view by clicking **Disable Updates**.
- Mouse over each entity to know more details about that entity.
- You can zoom in or zoom out of the connectivity view by using the + and - buttons.
- View the type and number of faults and alarms on Juniper Networks devices. The color that outlines each device indicates whether there are alarms on that device. Network Director also displays a count of the number of alarms of the highest severity on that device in this page.
- [Table 179](#) describes what each device icon that appears in the Datacenter Connectivity page indicates.


Table 179: Datacenter Connectivity View—Device Icon Legends

Device Icon	Indicates that
	The device is a switch that has no faults or alarms.
	The device is a bare metal server.
	The device is a hypervisor server (KVM or ESXi host).
	The device is not managed by Network Director.
	The device has one critical alarm.
	The device has one major alarm.
	The device has one informational alarm.

- View more details about the alarms present on a device. Zoom in to the device and click the device. Details about the device and the alarms are displayed. Click the alarm count to view more details about that alarm. Network Director opens the Fault mode and displays details about the alarm.

- To view the virtual tunnel end points (VTEPs) in your data center network, click the down arrow in the **Overlay** button and select **Show VTEPs** as shown in [Figure 32](#).

Figure 32: Filter Button

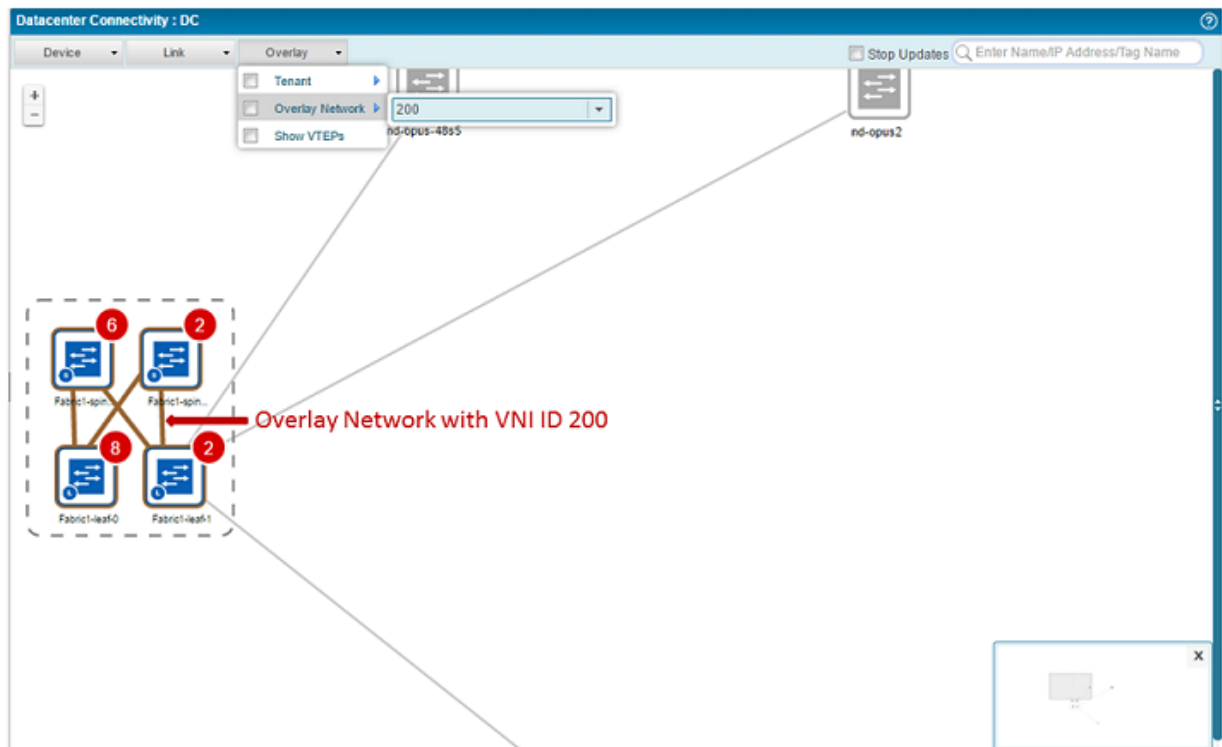


Network Director highlights and displays the VTEPs in your data center in the Datacenter Connectivity page.

- To view the overlay connectivity, do one of the following:
 - To view the overlay network connectivity based on the VNI number, click the down arrow in the **Overlay** button, select the **Overlay Network** check box, and select a VNI number from the list.

Network Director highlights the VXLAN connectivity between the various entities for the selected VNI as shown in [Figure 33](#).

Figure 33: VXLAN Connectivity



- To view the overlay networks for a tenant, click the down arrow in the **Overlay** button, select the **Tenant** check box, and select a tenant name from the list.

Network Director highlights the overlay network between the various entities for the selected tenant.

NOTE: This option is not available for data centers that use NSX as the cloud service provider.

5. Select **Stop Updates**, to stop the link status changes in real time in the Datacenter Connectivity page that might occur while the user is performing some tasks in this page.

RELATED DOCUMENTATION

[Understanding Cloud Networking | 784](#)

[Network Director Documentation home page](#)

Viewing Bare Metal Server Details

Bare metal servers balances the scalability and automation of the virtualized cloud data center. They are not virtualized and do not run a hypervisor. The bare metal server hardware is fully dedicated to the tenant. Bare metal servers are ideal in data centers that require to perform short-term, data-intensive functions without any kind of latency or overhead delays.

In Network Director, you can have data centers that are a combination of bare metal servers, network devices, and optionally a cloud infrastructure. You can view details about the bare metal servers in your data center using the View Baremetal Server task.

To view details about bare metal servers in your data center:

1. While in Build mode with Datacenter View selected, select Baremetal Servers under any data center from the View pane and click **Inventory** > **View Baremetal Servers** from the Tasks pane.

The View Baremetal Servers page displays details about the baremetal servers that are part of the selected data center, as shown in [Table 180](#).

Table 180: Baremetal Server Details

Field	Description
Name	Host name of the server.
IP Address	IP address of the server.
Vendor Name	Name of the server hardware vendor.
Operating System	The operating system that is running on the server.

RELATED DOCUMENTATION

[Understanding Cloud Networking | 784](#)

[Network Director Documentation home page](#)

Managing the Virtual Switch Inventory

IN THIS SECTION

- [View the Virtual Switch Inventory | 819](#)
- [Enable LLDP on Virtual Switches | 820](#)

Virtual switches enable virtual machines (VMs) on different hosts to communicate with each other by using the same protocols that are used over physical switches, without the need for any additional networking hardware. Once you have discovered a VMware vCenter in Network Director, the hosts, VMs, and virtual switches also come under Network Director management. You can use the View Virtual Switch Inventory task to manage and view the virtual switches in your virtual network infrastructure.

You can use the View Virtual Switches page to:

View the Virtual Switch Inventory

To view the virtual switch Inventory:

1. While in Build mode with Datacenter View selected, click the Virtual Center for which you want to view the virtual switch inventory details. From the Tasks pane, click **View Virtual Switches** from the Inventory menu to open the View Virtual Switches page.
2. The View Virtual Switches page displays the details for all the virtual switches in the Network. [Table 181](#) describes the fields in this table.

Table 181: View Virtual Switches page Field Descriptions

Field	Description
Name	Name assigned to the virtual switch. If it is a standalone virtual switch, the hostname on which the virtual switch resides is also displayed.
Version	Indicates the vSphere version.

Table 181: View Virtual Switches page Field Descriptions (*continued*)

Field	Description
Type	<p>The type of virtual switch. The switch type can be:</p> <ul style="list-style-type: none"> • Standalone vSwitch—The traditional virtual switch where the administrator configures and maintains a virtual switch for each host. • Distributed vSwitch— A distributed switch spans multiple hosts at the virtual network level. It is created by abstracting individual host-level virtual switches into a single, large, distributed switch. This considerably simplifies the VM networking by enabling you to set up VM access switching for your entire data center from a centralized interface.
MTU	<p>Maximum Transmission Unit—The maximum size of a protocol data unit that can be transmitted.</p>
Number of Ports	<p>Total number of ports held in the port groups configured for this switch.</p>
LLDP Status	<p>Indicates whether Link Layer Discovery Protocol (LLDP) is enabled on the specific virtual switch.</p>
LLDP Operation	<p>If LLDP is enabled, this field indicates the LLDP operation mode. The operation mode can be one of the following:</p> <ul style="list-style-type: none"> • Listen—Detects and displays information about the associated physical switch port, but information about the distributed vSwitch is not available to the physical switch. • Advertise—Makes information about the distributed vSwitch available to the physical switch, but does not detect or display information about the physical switch. • Both—Detects and displays information about the associated physical switch and makes information about the distributed vSwitch available to the physical switch. • Not applicable—Indicates that LLDP is not enabled. <p>NOTE: If you plan to use the Orchestration and Topology features extensively, then we recommend that you set the LLDP operation mode to Both.</p>

Enable LLDP on Virtual Switches

The Link Layer Discovery Protocol (LLDP) is a link layer protocol that is used by network devices to advertise their identity, capabilities, and neighbors. You can enable LLDP for a virtual switch in your virtual network infrastructure. Once you enable LLDP, you can view properties of the physical switch such as chassis ID, system name and description, and device capabilities from Network Director. LLDP is available only for vSphere virtual switch versions 5.0 and later.

To enable LLDP on a virtual switch:

1. While in Build mode with Datacenter View selected, click the Virtual Center for which you want to view the virtual switch inventory details. From the Tasks pane, click **View Virtual Switches** from the Inventory menu to open the View Virtual Switches page.
2. Select the virtual switch for which you want to enable LLDP and click **Enable LLDP**.
Network Director enables LLDP on the selected switch and changes the status field to Enabled.

RELATED DOCUMENTATION

[Understanding Virtual Network Management | 785](#)

[Network Director Documentation home page](#)

Viewing the Hypervisor Servers in a Data Center

The hypervisor server inventory displays details of all the hypervisor servers in a data center.

To view details about the hypervisor servers:

1. While in Build mode with Datacenter View selected, click the data center for which you want to view the host inventory details. From the Tasks pane, click **View Hypervisor Servers** from the Inventory menu to open the View Hypervisor Servers page.

The Hypervisor Servers page displays details of all the hypervisor servers that are part of the selected data center. [Table 182](#) describes the fields in this table.

Table 182: View Hypervisor Server page Field Descriptions

Field	Description
Host	Name or the IP address of the hypervisor server.
Hypervisor Type	Type of hypervisor. For example, ESX/ESXi, which is the hypervisor for the VMware.
Hypervisor Version	The hypervisor version that is running on the host.
Hardware Vendor	Name of the hardware vendor of the hypervisor server.

Table 182: View Hypervisor Server page Field Descriptions (*continued*)

Field	Description
Hardware Model	Name of the host hardware model.
Connection State	One of the following: <ul style="list-style-type: none"> • Connected—The host is connected to the cloud infrastructure. • Disconnected—The user has explicitly taken down the host and the host is not connected to the cloud infrastructure. • Not responding—The host is connected to the cloud infrastructure, but the server is not receiving heartbeats from the host. The connection state automatically changes to Connected when heartbeats are received again.
Number of VMs	Total number of VMs running on the hypervisor server.
Number of NICs	Total number of physical NICs on the hypervisor server.

RELATED DOCUMENTATION

[Understanding Virtual Network Management | 785](#)

[Network Director Documentation home page](#)

Managing Network Adapter Associations

IN THIS SECTION

- [Manage Network Adapter Associations | 824](#)
- [Configure Network Adapter Associations | 824](#)
- [Set Up the Devices for LLDP-Based Automatic Link Discovery | 826](#)
- [Export Network Adapter Associations | 827](#)
- [Delete Network Adapter Associations | 827](#)

Network Director enables you to configure and view the connectivity or association between a host network adapter and the corresponding physical network adapter on the physical switch. If LLDP is enabled

on the physical switch and the virtual switch, Network Director automatically discovers and displays the host network adapter to physical network adapter connection. However, if LLDP is not enabled, you must manually configure these connections by importing the connection details in to Network Director.

You can automate orchestration by using Network Director only if the host to physical switch connectivity is established. Using the Manage Physical Switch Association table, you can view the connectivity between a host network adapter and a physical switch port. You can also view the virtual switch that is bound to the given host network adapter. If any of the host network adapters is not connected to a corresponding the physical switch port, then follow the steps mentioned in [“Configure Network Adapter Associations” on page 824](#) to manually configure network adapter associations.

After you make any change to the network adapter association, Network Director initiates an orchestration job, if orchestration is enabled for the given virtual network.

NOTE: Manual NIC associations are used for orchestration purpose only. They will not be visible in the Datacenter Connectivity page.

You can do the following tasks to manage network adapter associations:

Manage Network Adapter Associations

To view and manage the host associations:

1. Select the virtual network or the host for which you want to view and manage the associations from the View pane and click **Connectivity > Manage NIC Associations** from the Tasks pane.

The Manage NIC Associations page opens. [Table 183](#) describes the fields in this page.

Table 183: Manage NIC Associations page Field Descriptions

Field	Description
Network Adapter Name	Network adapter of the host to which the virtual switch is connected.
MAC Address	MAC address of the host.
Physical Switch	IP address or the name of the physical switch to which the virtual switch is connected.
Physical Switch Port	Number of the physical switch port to which the virtual switch is connected.
Discovered Through	Displays how the association was discovered: <ul style="list-style-type: none"> • LLDP—Indicates that the association was automatically discovered by Network Director, using LLDP. • Manual—Indicates that the association was manually configured, using the Import Associations feature. For detailed steps for importing network adapter associations, see “Configure Network Adapter Associations” on page 824. • Blank—Indicates that the association was not discovered or was deleted.
Virtual Switch	Name assigned to the virtual switch. If it is a standalone virtual switch, then the hostname on which the virtual switch resides is also displayed.

Configure Network Adapter Associations

Network Director does not automatically discover or display virtual network adapter to physical network adapter associations if LLDP is not enabled on the physical switch and the virtual switch. In this scenario, you can identify the connection details outside of Network Director by doing one of the following:

- Use the Edit Association window from the Manage Network Adapter Associations page to edit individual associations.

- Use the Import Network Adapter Associations page to import one or more associations into Network Director to establish these associations manually.

To manually configure individual network associations:

1. Select **Cloud Infrastructure** in the View pane.
2. To edit individual associations, from the Manage Network Associations page, select a row and click **Edit Association**.
3. in the Edit Associations window, select the hostname of the physical switch and the switch port to which you want the selected virtual network adapter to be connected to.
4. Click **Update** to save the details and update the Manage Network Adapter Associations table.

To manually configure multiple network associations:

1. Select **Cloud Infrastructure** in the View pane.
2. In the Manage Network Associations page, click **Import Network Adapter Associations**. The Import Network Adapter Associations page opens.
3. Click **Browse** and select the CSV file in which you have specified the network adapter associations that you want to import.

You can download a sample CSV file by clicking **Download Sample CSV**.

Network Director imports the associations after checking the connectivity and changes the status of the successful associations to Manual in the Manage Network Adapter Associations page.

Set Up the Devices for LLDP-Based Automatic Link Discovery

Network Director enables you to configure and view the connectivity or association between a host network adapter and the corresponding physical network adapter on the physical switch. If LLDP is enabled on the physical switch and the virtual switch, Network Director automatically discovers and displays connectivity between the host network adapter and the physical network adapter and enables the system to perform orchestrations. However, for devices that are directly connected to the ESXi server in your vCenter based data center network, you must perform the following additional steps to enable this autodiscovery:

To set up automatic link discovery and orchestration for devices that are directly connected to the vCenter based data center network:

1. Log in to the switch, that is directly connected to the ESXi server in your vCenter based data center network, by using the CLI.
2. Configure the port ID type, length, and value (TLV) by running the following command:

```
[edit]
user@switch# set protocols lldp port-id-subtype interface-name
```

3. Only for *QFabric Systems*, *Virtual Chassis*, and *Virtual Chassis Fabrics* that are connected to the vCenter based data center network, set the LLDP management address by running the following command in the configuration mode:

```
[edit]
user@switch# set protocols lldp management-address IP-address
```

IP-address is the management address of the QFabric system, Virtual Chassis, or the Virtual Chassis Fabric system .

4. Commit your changes.

The physical device that is directly connected to the ESXi server starts advertising their respective management address through LLDP.

5. Resynchronize the virtual network from Network Director. For more information about resynchronizing virtual networks, see [“Managing Cloud Infrastructure” on page 798](#).

Network Director uses the information advertised by the physical device for autodiscovery of links between the physical device and the hosts.

Export Network Adapter Associations

You can export network associations to Microsoft Excel and make modifications outside Network Director and later import the modified associations into Network Director. Exporting network adapter associations are also beneficial if you were to reinstall Network Director. In case of a reinstallation, you would have to configure all the manual associations after reinstalling Network Director. To save this trouble, you can export the network adapter associations before uninstalling Network Director and import the associations once Network Director is reinstalled.

To export network adapter associations:

1. Select the virtual network or the host for which you want to view and manage the associations from the View pane and click **Connectivity > Manage Network Adapter Associations** from the Tasks pane.
2. Click **Export Associations**.

Network Director exports the network adapter associations. Save the file to a location of your choice. Modify the file if required, and import the associations back into Network Director. For detailed steps for importing network adapter associations, see [“Configure Network Adapter Associations” on page 824](#).

Delete Network Adapter Associations

You can delete the network adapter associations that you no longer want to use, from the Manage Network Adapter Associations page.

NOTE: You cannot delete LLDP associations that are automatically created by Network Director.

To delete a network adapter association:

1. Select the row corresponding to the network adapter association that you want to delete and click **Delete Association**.
2. Confirm the deletion.

Network Director deletes the selected association.

RELATED DOCUMENTATION

[Understanding Virtual Network Management | 785](#)

[Network Director Documentation home page](#)

Configuring Overlay Networks and Tenants

IN THIS CHAPTER

- [VXLAN–EVPN Overlay Overview | 828](#)
- [Create a Layer 3 Fabric based Underlay Network | 829](#)
- [Creating and Managing Overlay Fabrics | 832](#)
- [Setting Up a VXLAN–EVPN-Based Data Center | 836](#)
- [Creating and Managing Tenants | 837](#)

VXLAN–EVPN Overlay Overview

Spanning Tree Protocol (STP), multichassis link aggregation group (MC-LAG), and Transparent Interconnection of Lots of Links (TRILL) were some of the commonly used technologies in traditional data centers. However, as the data centers started to grow exponentially, these technologies were not able to scale to meet the requirements of data centers. To cater to these requirements, data center administrators started using orchestration tools such as VMware vCenter, VMware vCenter with NSX, OpenStack, and OpenStack with NSX plug-in to orchestrate the networking needs of the tenants that a data center serves. This approach meant that the data center might require additional plug-ins to configure VLANs and gateways or that you might need to make changes to the physical network topology, to accommodate a new tenant. Most of the current day data centers handle multiple customer groups, organizations, or tenants that require a new data center architecture that decouples the underlay network from tenant overlay networks. A Layer 3 Fabric underlay coupled with a Virtual Extensible LAN (VXLAN)–Ethernet VPN (EVPN) overlay solution that uses bare metal servers and/or virtual servers, or both Network Director for management enables data center and cloud operators to deploy much larger networks than that are otherwise possible with traditional Layer 2 Ethernet-based architectures.

Some of the major advantages that VXLAN–EVPN overlay networks provide are:

- **Scalability**—Most enterprises accommodate their growth by increasing the use of cloud services, while others choose to deploy their own private and hybrid clouds. Service providers also must be able to grow rapidly to have sufficient capacity to meet the demands of the enterprises. Today's networks are often too rigid and difficult to change for scaling to meet the needs of large enterprises and service providers. But by using VXLAN–EVPN, data centers and cloud operators can have up to 16 million overlay networks in a cloud data center.

- **Operational efficiency**—As enterprises expand geographically, the physical distance between the data centers and users also increases, which makes timely maintenance and application mobility a challenge. The VXLAN—EVPN solution enables network administrators to easily migrate applications within the data center and between data centers for business continuity so that they can maintain the data center without downtime, for effective load balancing.
- **High Performance**—End users often experience poor response times and even outages of business-critical applications caused by bandwidth limitations and latency problems. Multi-pathing and control plane learning features, that are part of the VXLAN—EVPN solution, can optimize network traffic flows, rein in network faults, and ensure maximum utilization of bandwidth.

RELATED DOCUMENTATION

[Setting Up a VXLAN—EVPN-Based Data Center | 836](#)

[Creating and Managing Tenants | 837](#)

Create a Layer 3 Fabric based Underlay Network

Layer 3 Fabrics based on multi-stage Clos architecture increases resiliency and also supports new technologies such as VMware NSX that allow enterprises to deploy applications, servers, and virtual networks within seconds. The Layer 3 Fabrics can also function as an underlay network for a EVPN—VXLAN based overlay networks. Layer 3 Fabrics use BGP as the control plane protocol to advertise prefixes, perform traffic engineering, and tag traffic. Layer 3 Fabric uses spine-and-leaf topology. In the spine-and-leaf topology, all the leaf devices are connected to the spine devices in a mesh.

You can use one of the following device models as the spine device in a Layer 3 Fabric:

- QFX10008
- QFX10002-36Q
- QFX10002-72Q
- QFX5100-24Q-2P
- QFX5200-32C-32Q

You can use the following device models as leaf devices:

- QFX5100-48S-6Q
- QFX5100-96S-8Q
- QFX5100-48T-6Q

- QFX5200-32C-32Q
- QFX5100-24Q
- QFX5100-24Q-2P

NOTE: QFX5100-48T-6Q, QFX5100-48S-6Q and QFX5100-96S-8Q can be standalone or Virtual Chassis leaf devices. Whereas, QFX10008 and QFX5100-24Q are supported only as standalone devices. QFX5200-32C-32Q can be supported as both spine and leaf device, however, it cannot be a Virtual Chassis member.

You can quickly create and deploy Layer 3 Fabrics by using Network Director.

Before you begin, ensure that you have the necessary privileges on the FTP and the file server that Network Director uses for Zero Touch Provisioning. For more details, see [“User Privileges Required for the DHCP and File Server While Using Zero Touch Provisioning” on page 743](#).

To create and deploy a Layer 3 Fabric using Network Director:

1. With Build mode or Deploy mode selected, select one of these options under Views: **Logical View**, **Location View**, **Device View**, or **Custom Group View**.
2. In the Tasks pane, click **Network Builder > Manage Layer 3 Fabrics**.

The Manage Layer 3 Fabrics page opens.

3. Click **Create** to create a new Layer 3 Fabric. [Table 184](#) describes the various pages that are available in the Layer 3 Fabric wizard and the actions that you must perform on each page.

Table 184: Layer 3 Fabric Wizard Pages

Name of the wizard page	Action
Fabric Requirements	Specify the fabric name, spine and leaf device models, number of spine and leaf devices, and settings for dynamic host link aggregation group (LAG) creation.
Device Details	Displays the number of spine and leaf devices that you are provisioning as part of the initial capacity, enables you to edit the hostname prefix for all the spine and leaf devices.
Configuration	Specify the configuration details such as the Loopback Network Address, Interconnect Network Address, VLAN Network Address, Start Management IP, Maximum Hosts or VMs per leaf, Spine-BGP Autonomous System Number, Leaf-BGP Autonomous System Number, Device Password, and Management Gateway IP address.

Table 184: Layer 3 Fabric Wizard Pages (continued)

Name of the wizard page	Action
Cabling	<p>Displays the recommended cabling plan for the device that you select in the left pane. If you specify all the spine and leaf devices, the cabling plan displays the exact port numbers that you must use to connect your spine and leaf devices. However, if you have not specified any leaf devices and have specified only the maximum leaf count, the plan displays all the leaf devices as unknown. The leaf devices in this case are plug-and-play and you can use any of the uplink ports on the leaf devices.</p> <p><i>You must follow the cabling plan for all the physical connections between the spine and leaf devices in your fabric.</i></p>
ZTP	<p>Zero Touch Provisioning (ZTP) enables you to provision devices in your network automatically, without manual intervention. You can use the ZTP wizard page that is part of the Layer 3 Fabric wizard. When a device is physically connected, it boots up with the factory-default configuration and auto installs a configuration file from the network. In Network Director, ZTP is used to provision Layer 3 Fabrics and all the configurations are pushed through OpenClos.</p> <p>NOTE: When you select QFX10008 as the spine device, only the leaf devices in the fabric are provisioned using ZTP. For the QFX10008 spine device, you must either copy the configuration file from Network Director or manually download it from the file server. To copy the configuration file from the file server, SSH or Telnet must be enabled on the device (QFX10008).</p>
Review	<p>Review the Layer 3 Fabric configuration, deploy the fabric, or save and close the Layer 3 Fabric wizard.</p>

For detailed steps for creating the Layer 3 Fabric, see [“Creating Layer 3 Fabrics” on page 746](#).

4. After you have deployed the Layer 3 Fabric, you can run connectivity checks to troubleshoot any connectivity issues in the fabric. You can initiate a connectivity check by clicking **Run Connectivity Check** or **Re-run Connectivity Check** in the Cabling column in the Manage Layer 3 Fabrics page.

Network Director uses LLDP to check the connectivity of each device with the device’s neighbor and compares it with the recommended cabling plan and reports the results in the Cabling Check Results page. For more details, see [“Performing Layer 3 Fabric Connectivity Checks” on page 765](#).

5. Click **View Topology** in the Summary column of the Manage Layer 3 Fabric page to view the physical topology of the Layer 3 Fabric.
6. After you have successfully deployed a Layer 3 Fabric, you can monitor the performance, identify errors or bottlenecks in your underlay network by using the Monitoring mode of Network Director. For details, see [“Status Monitor for Layer 3 Fabrics” on page 1422](#) and [“Port Status for IP Fabric Monitor” on page 1399](#).

RELATED DOCUMENTATION

[Understanding Layer 3 Fabrics | 742](#)

[Creating Layer 3 Fabrics | 746](#)

Creating and Managing Overlay Fabrics

IN THIS SECTION

- [Creating an Overlay Fabric | 834](#)
- [Modifying an Overlay Fabric | 835](#)
- [Viewing Overlay Fabric Configuration Details | 835](#)

You can create and manage overlay fabrics from the Manage Overlay Fabrics page.

Before you begin, make sure that you have at least one Layer 3 Fabric based data center in Network Director.

NOTE: Layer 3 Fabric configured using CLI and discovered in Network Director is not used to create an overlay fabric. You can use the Layer 3 fabric created in Network Director to create an overlay fabric.

To open the Manage Overlay Fabrics page:

1. Open the Datacenter View in Network Director.

2. Select the Layer 3 Fabric based data center from the View pane and select **Overlay Networks Builder** > **Manage Overlay Fabrics** from the Tasks pane.

The Manage Overlay Fabrics page opens. [Table 185](#) describes the fields that are displayed in the Manage Overlay Fabrics page.

Table 185: Manage Overlay Fabrics Field Descriptions

Field	Description
Name	Name of the overlay fabric.
Description	Description for the overlay fabric.
Autonomous System No.	The autonomous system number given for the overlay fabric. Autonomous system (AS) is a collection of devices that act as a single administrative entity, such as an overlay fabric. Each autonomous system is assigned a unique autonomous system number. In BGP routing, AS numbers are used to uniquely identify a network on the Internet.
Route Reflector Network Address	IP address of the route reflector device. Route reflector devices have the ability to readvertise routes learned from an internal peer to other internal peers. Therefore, rather than all internal peers of the fabric to be fully meshed with each other, route reflection requires that only the route reflector be fully meshed with all internal peers. The route reflector and all of its internal peers form a cluster.
Layer 3 Fabrics	The Layer 3 Fabrics that is part of the overlay fabric.
Deployment State	Deployment status of the overlay fabric. Deployment state can be one of the following: <ul style="list-style-type: none"> • Pending Deployment—The fabric details have been saved, but not deployed. • Deployment in progress—The fabric is being deployed on the Layer 3 Fabric. • Partially deployed—The overlay fabric is successfully deployed on some devices, and the deployment failed on the other devices. • Deployed—The overlay fabric is successfully deployed on all the devices. • Deploy failed—Deployment failed on all devices. • Pending decommission—The overlay fabric is deleted. There might be a slight delay for Network Director to apply the changes to the devices; the system displays this status during this time.

From the Manage Overlay Fabrics page, you can:

Creating an Overlay Fabric

You can create an overlay fabric by following the steps given in this section. Alternatively, you can create an overlay fabric inline while you create a tenant.

To create an overlay fabric:

1. Click **Add** in the Manage Overlay Fabrics page.

The Create Overlay Fabric window appears.

2. Enter a name and description for the overlay fabric.
3. Specify the autonomous system (AS) number for the overlay fabric. You can specify the AS number in the range 1 to 4294967295.
4. Specify the IP address and the subnet mask of the route reflector device in the overlay fabric.
5. Select the Layer 3 Fabric that you want to include in the overlay fabric from the **Available** list and click the arrow to add it to the overlay network. Network Director lists all the Layer 3 Fabrics that are part of the data center and available for assignment in the Available list.

NOTE:

- You must have at least one Layer 3 Fabric available when you create an overlay fabric. To create a Layer 3 Fabric, follow the procedure given in [“Creating Layer 3 Fabrics” on page 746](#).
- Layer 3 Fabric configured using CLI and discovered in Network Director is not used to create an overlay fabric. You can use the Layer 3 fabric created in Network Director to create an overlay fabric.

6. Click **OK** to create the overlay fabric.
7. Click **Preview & Deploy** to deploy the overlay configuration.

After you have created or deployed the overlay configuration, you can create a tenant. See [“Creating and Managing Tenants” on page 837](#) for the procedure to create a tenant.

Modifying an Overlay Fabric

To modify an overlay fabric:

1. Select the overlay fabric in the Manage Overlay Fabrics page and click **Edit**. The Edit Overlay Fabric page appears.
2. Modify the overlay fabric details. You can modify the Description, Autonomous System Number, Route Reflector Network Address, and the Layer 3 Fabric. By modifying Layer 3 Fabric, you can assign a new Layer 3 Fabric to the overlay fabric or delete a Layer 3 Fabric from the overlay fabric.

NOTE: You cannot delete a Layer 3 Fabric that has tenants created in the overlay network. You must remove all the tenants from the Overlay Fabric before you can delete the Layer 3 Fabric.

Viewing Overlay Fabric Configuration Details

After you have successfully deployed an overlay fabric, you can view the configuration details of the overlay fabric from the View Overlay Fabric Configuration page. Click the **Deployed**, **Partially Deployed**, or the **Deploy Failed** link corresponding to the overlay fabric under the Deployment State column in the Manage Overlay Fabrics page to open this page.

The View Overlay Fabric Configuration page displays configuration details as described in [Table 186](#).

Table 186: View Overlay Fabric Configuration Details

Field	Description
Overlay Fabric Name	Name of the overlay fabric for which the configuration details are being viewed.
Description	Description of the overlay fabric.
Devices	Device list mapped to the overlay fabric.
Element Name	Name of the instance mapped to the overlay fabric.
Configurations	Configuration details of the device for the element (for example, Overlay Fabric).
Status	Deployment status of the overlay fabric.

Table 186: View Overlay Fabric Configuration Details (*continued*)

Field	Description
Error Message	Error message for a failed deployment. NOTE: The error message is visible when the user clicks the hyperlink under the status column.

RELATED DOCUMENTATION

[VXLAN–EVPN Overlay Overview | 828](#)

[Creating and Managing Tenants | 837](#)

Setting Up a VXLAN–EVPN-Based Data Center

Junos Space Network Director enables you to set up a VXLAN–EVPN based data center network by performing a few easy tasks. After you have created a VXLAN–EVPN based data center, you can monitor the performance of your data center, find bottlenecks in your data center, and troubleshoot these issues.

Setting Up a VXLAN–EVPN based data center using Network Director involves the following steps:

1. Create a Layer 3 Fabric (s) using Network Director. To know more about creating Layer 3 Fabrics, see [“Create a Layer 3 Fabric based Underlay Network” on page 829](#).

NOTE: Layer 3 Fabric configured using CLI and discovered in Network Director is not used to create an overlay fabric.

2. Make sure that you have created at least one *controller-less* data center in Network Director. A controller-less data center does not use any cloud infrastructure providers. If you have not created a controller-less data center, create one by following the steps given in [“Creating Data Centers Using Network Director” on page 793](#).
3. Assign the Layer 3 Fabric or Fabrics to the data center. For detailed steps on assigning a fabric to a data center, see Assigning Network Devices to a Data Center section in [“Creating Data Centers Using Network Director” on page 793](#).

4. Select the data center and create an overlay fabric from the Manage Overlay Fabrics page. You can add one or more Layer 3 Fabrics to an overlay fabric. For detailed steps, see [“Creating and Managing Overlay Fabrics” on page 832](#).

NOTE: You can skip this step and choose to create an overlay fabric when you create the tenant.

5. Create a tenant using the Create Tenant task under Manage Tenants and assign the tenant to an overlay fabric. You can either create an overlay fabric inline while creating a tenant or choose an existing overlay fabric. As part of the create tenant task, you also create one or more overlay networks, and select ports and multi-homed LAG to be part of this overlay network. You can also create new LAGs as part of this task. For detailed steps, see [“Creating and Managing Tenants” on page 837](#).
6. Preview and deploy the tenant to the data center.
7. After you have successfully deployed a tenant, you can monitor the performance and identify errors or bottlenecks in your overlay network by using the tenant monitor. For detailed steps on monitoring tenant details, see [“Monitoring Tenant Details” on page 1304](#).

RELATED DOCUMENTATION

[VXLAN–EVPN Overlay Overview | 828](#)

[Creating and Managing Tenants | 837](#)

Creating and Managing Tenants

IN THIS SECTION

- [Creating a Tenant | 838](#)
- [Modifying a Tenant | 842](#)
- [Viewing Tenant Configuration Details | 843](#)

You can create and manage tenants from the Manage Tenants page.

Before you begin, make sure that you have at least one Layer 3 Fabric based data center in Network Director.

To open the Manage Tenants page:

1. Open the Datacenter View in Network Director.
2. Select a data center from the **View** pane for which you want to create a tenant and select **Overlay Networks Builder > Manage Tenants** from the **Tasks** pane.

The Manage Tenants page opens. [Table 187](#) describes the fields that are displayed in the Manage Tenants page.

Table 187: Manage Tenants Field Descriptions

Field	Description
Tenant Name	Name of the tenant.
Description	Description given for the tenant.
Overlay Name	Name of the overlay fabric.
Deployment State	<p>Deployment status of the tenant. Deployment state can be one of the following:</p> <ul style="list-style-type: none"> • Pending Deployment—The tenant details have been saved, but not deployed. • Deployment in progress—The tenant is being deployed on the Layer 3 Fabric. • Partially deployed—The tenant is successfully deployed on some devices and the deployment failed on the other devices. • Deployed—The tenant is successfully deployed on all devices. • Deploy failed—Deployment failed on all the devices. • Pending decommission—The tenant is deleted. There might be a slight delay for Network Director to apply the changes to the devices; the system displays this status during this time.

You can perform the following tasks from the Manage Tenants page:

Creating a Tenant

To create a tenant:

1. Click **Add** in the Manage Tenants page.
The Create Tenant page opens.
2. Enter a name and description for the tenant.

You can enter a maximum of 255 characters. Underscore is the only special character that is allowed in the name field.

3. Click **Overlay Fabric Name** to select an overlay fabric for the tenant.

You can also create an overlay fabric inline by clicking **Create New**.

To create the overlay fabric:

- a. Enter a name and description for the overlay fabric.
 - b. Specify the autonomous system (AS) number for the overlay fabric.
 - c. Specify the IP address and the subnet mask of the route reflector network address in the overlay fabric.
 - d. Select the Layer 3 Fabric that you want to include in the overlay network from the Available list and click the arrow to add it to the overlay fabric. Network Director lists all the Layer 3 Fabrics that are part of the data center and available for assignment in the Available list.
 - e. Click **OK** to create the overlay fabric. The overlay fabric is selected to be used by the tenant.
4. Click **Add** in the Overlay Networks table to create an overlay network for the tenant.

The Create Overlay Network for Tenant window opens.

5. Enter the Virtual Extensible LAN Network Identifier (VNI) for the VXLAN overlay network. You can also select a VLAN from the list of existing VLANs.

VNI is a numeric value to identify a VXLAN. You can enter a number from 1 through 16777214.

6. Create a VRF instance that is to be assigned to the VXLAN overlay network or use the default VRF instance. You can also select a VRF instance from the list of existing VRF instances.

The default VRF instance is named as **default_tenant-name**.

Each tenant in an overlay fabric has its own tenant-specific routing table that contains the routing information for that tenant. To separate the route of each tenant from the other network traffic and routes, the spine device in the Layer 3 Fabric creates a separate routing table for each tenant called the Virtual Routing and Forwarding (VRF) instance. A can tenant can have multiple VRF instances.

To create a new VRF instance, click **Create**. Specify a name, description, and loopback network address for the VRF instance and click **OK**.

NOTE: The VRF loopback network address is a mandatory field for routing devices. You can modify the value in this field, but cannot delete it once it is deployed.

NOTE: If you have mapped a VRF instance to an overlay network, for example, overlay-network1 is mapped to VRF1, you cannot modify the VRF instance for this overlay network (after it is deployed) to another VRF instance, for example, to VRF2. To map the new instance, you must delete the overlay network and re-create the overlay network.

7. Specify the ID and name of the VLAN. Network Director maps this VLAN ID to the VXLAN overlay network.

The VLAN name is auto-populated, which you can edit, if needed.

8. Specify the IP address and subnet mask of the default layer 3 gateway network address. The default gateway enables the traffic to be routed between different VXLANs networks within the VRFs of the tenant.

9. Click **Auto Select Ports** to list the ports or LAGs that are connected to the devices. You can select these ports or LAG and assign it to the overlay network.

10. Click the **Ports Selection** tab to select the ports that you want to be part of the overlay network.

The Ports Selection tab displays the leaf devices of Layer 3 Fabrics that are part of the overlay fabric in the data center. When you select one or more leaf devices in the left pane, the ports associated with the device are displayed in a table in the right pane.

11. To preview a port selection, select a port and click **Preview**. You can edit the port selection listed in the preview by clicking **Edit**.

12. Select a leaf device or devices and click **Select by Range** to specify the port range that you want to include in the overlay network. You can select the port type as **Normal Ports** or **Channelized Ports**.

The normal ports are physical ports such as et, ge, and xe. Channelized ports are et ports that are 40-Gigabit, 100-Gigabit and can be converted to xe ports. For example, et-0/0/0 (40-Gigabit) can be converted to xe-0/0/0:0, xe-0/0/0:1, xe-0/0/0:2, xe-0/0/0:3. Similarly a 10-Gigabit and 100-Gigabit can be converted to 25-Gigabit.

NOTE: You can also select multiple devices from the left pane by the Shift key pressed and selecting devices. All the ports for selected devices are listed.

13. Click the **LAG Selection (Multi-homed)** tab to select the link aggregation groups (LAGs) that you want to be part of the overlay network. You can select one or more LAGs from the existing LAGs that are created in Network Director or create and add new LAGs.

NOTE: The lag you created will have a temporary system generated name with the prefix LAG; for example, LAG0 or LAG1. When you save or deploy a tenant, the available aggregate Ethernet interface name for example, ae0 or ae1 replaces this LAG name. The LAG name must be common for all the leaf devices across a host.

14. To select from a list of existing LAGs, click **Select Existing LAG**. The LAG Selection (multi-homed) window opens. Select the LAGs that you want to include in the overlay network and click **OK**. Network Director adds the LAGs that you selected to the LAG Selection (multi-homed) table.

15. Do the following to create a new LAG:

- a. Click **Add** in the LAG Selection (multi-homed) table. The Create LAG window opens.
- b. Select the Layer 3 Fabric on which you want to create the LAG.
- c. Click **Add** to specify the device and port details for the LAG. Network Director adds a row to the Device and Port details table.
- d. In the Device column, click and select the leaf device from which you want to select a port for the LAG.
- e. In the Port column, click to select the ports that you want to add to the LAG.
- f. Repeat steps c through e to add ports from another leaf device.

The Port column lists only those ports that are not part of any LAG or port selection tab.

- g. To remove a device and port entry from the LAG, select the row from the Device and Port details table and click **Remove**.
 - h. Click **OK** to save the details and create the LAG. Network Director adds the LAG that you created, to the LAG Selection (multi-homed) table.
16. To delete a LAG from the overlay network, select the LAG in the LAG Selection (multi-homed) table and click **Delete**.
 17. Click **OK** to save the details and add the overlay network for the tenant. Network Director lists the overlay network that you added in the Overlay Networks table in the Create Tenant page.
 18. Repeat steps 4 through 17 to add more overlay networks.
 19. Click **Preview & Deploy** to review the changes and deploy the tenant details to the overlay fabric.

Modifying a Tenant

To modify a tenant:

1. Select a tenant in the Manage Tenants page and click **Edit**.
The Edit Tenant page appears.
2. You can modify the tenant description, overlay network and its details, and the VRF details. You can also add a new overlay network by clicking **Add**.
3. To modify the overlay network details, select the overlay network you want to modify and click **Edit**.
The Edit Overlay Network for Tenant page opens.
4. You can modify VNI, VLAN ID, L3 Gateway Network Address, VLAN Name, Port Selection, and LAG Selection (multi-homed) details and click **OK**. For details, see [“Creating a Tenant” on page 838](#).

NOTE: You cannot change a VLAN name after an overlay network is deployed.

5. To edit the VRF instance details and the loopback network address, you can either:
 - Click the VRF instance name in the table.
 - Select an overlay network and click **Edit VRF Details**.

The Edit VRF Instance for tenant page appears.

6. Modify the description and the loopback network address and click **OK**.

NOTE: You can create a VRF instance while you create the overlay network (this VRF instance has the tenant scope). However, you cannot delete it directly as it might be used by other overlay networks. Therefore, to enable deleting make sure you do not map this VRF instance to the VLAN while creating the tenant. You can however, edit the details of a VRF instance. If you want to delete a VRF instance, you must first delete the deployed VLANs, to which the VRF instance is associated. After you delete these VLANs, the VRF instance is removed after you save and deploy the configuration.

Viewing Tenant Configuration Details

After you successfully deploy a tenant, you can view the configuration details of the tenant from the View Tenant Configuration page. Click the **Deployed** link corresponding to the tenant name under the Deployment State column in the Manage Tenant Fabrics landing page to open this page.

The View Configuration page displays the details as described in [Table 188](#).

Table 188: View Tenant Configuration Details

Field	Description
Tenant Name	Name of the overlay fabric for which the configuration details are being viewed.
Description	Description of the overlay fabric.
Overlay Fabric Name	Device list that is mapped to the selected overlay fabric.
View VNI	Click View VNI to view the VNI details of the overlay network that are mapped to the tenant. You can view the following details: <ul style="list-style-type: none"> • Ports selection, that are part of the overlay network. • Lag selection (multi-homed), that are part of the overlay network

Table 188: View Tenant Configuration Details (continued)

Field	Description
View VRF	<p>Click View VRF, to view the VRF instance configuration details that are mapped to the tenant. You can view the following details:</p> <ul style="list-style-type: none"> • Devices—Devices that are mapped to the tenant. • Element—Name of the VRF instance mapped to the tenant. • Configurations—Configuration details of the VRF instance for the selected device. • Status—Deployment status of the tenant. • Message—The error message if any. <p>NOTE: To go back to the Overlay Network table click < Back to Overlay Networks link.</p>
VNI	VXLAN Network identifier of the tenant.
VLAN ID	VLAN ID of the tenant.
VLAN Name	VLAN name of the tenant.
L3 Gateway Network Address	IP address and subnet mask of the default L3 gateway.
Ports	Port that is part of the overlay network.
LAGs	Lag that is part of the overlay network
VRF Instance	VRF instance that is assigned to the VXLAN overlay network.
Status	Deployment status of the tenant.

RELATED DOCUMENTATION

| [VXLAN—EVPN Overlay Overview](#) | 828

Configuring VRRP Profiles

IN THIS CHAPTER


- [Understanding VRRP Profiles | 845](#)
- [Creating and Managing VRRP Profiles | 846](#)

Understanding VRRP Profiles

Virtual Router Redundancy Protocol (VRRP) enables hosts on a LAN to make use of redundant routing devices on that LAN without requiring more than the static configuration of a single default route on the hosts. The routing device on which VRRP is enabled share the IP address corresponding to the default route configured on the hosts. At any time, one of the routing devices is the master (active) and the others are backups. If the master fails, one of the backup routers becomes the new master, providing a virtual default routing platform and enabling traffic on the LAN to be routed without relying on a single routing device. Using VRRP, a backup routing device can take over a failed master router within a few seconds and without any interaction with the hosts.

Routing devices on which VRRP is enabled dynamically elect the master and backup devices. You can also configure the assignment of the master and the backup routers by specifying the priorities from 1 through 255 for master election, with 255 being the highest priority. VRRP functions by the default master sending advertisements to the backup devices at regular intervals. The default interval is 1 second, but you can set this interval. If a backup device does not receive an advertisement for the set period, the backup device with the next highest priority takes over as master and begins forwarding packets. To minimize network traffic, VRRP is designed in such a way that only the device that is acting as the master sends out VRRP advertisements at any given point in time. The backup devices do not send any advertisement until and unless they take over as the master.

The following figure illustrates a basic VRRP topology. In this example, routers A, B, and C are running VRRP and together they function as a virtual router. The IP address of this virtual router is 10.10.0.1 (the same address as the physical interface of Router A).



Because the virtual router uses the IP address of the physical interface of router A, router A is the master router, while routers B and C function as backup VRRP routers. Clients 1 through 3 are configured with the default gateway IP address of 10.10.0.1. As the master router, router A forwards packets sent to its IP address. If the master virtual router fails, the backup router configured with the higher priority becomes the master virtual router and provides uninterrupted service for the LAN hosts. When router A recovers, it becomes the master virtual router again.

RELATED DOCUMENTATION

[Creating and Managing VRRP Profiles | 846](#)

[Creating and Managing Port Profiles | 414](#)

[Creating and Managing VLAN Profiles | 502](#)

Creating and Managing VRRP Profiles

IN THIS SECTION

- [Managing VRRP Profiles | 847](#)
- [Creating VRRP Profiles | 848](#)
- [Specifying VRRP Settings for an EX Switching, Campus Switching ELS, or Data Center Switching ELS or non-ELS | 848](#)

VRRP profiles enable grouping of VRRP parameters and applying them to one or more interfaces. You can configure the attributes for this profile by using the VRRP option under Profiles. You can also choose this profile as an in-line profile in a Port profile and a VLAN profile.

- VRRP on Port profile—Select VRRP in Port profile if you want to configure VRRP on a physical interface. The VRRP settings in Port profile are displayed only when you select the Service Type as Custom and Family Type as Routing. The VRRP attributes such as group ID and priority are applied to the device during the profile assignment.
- VRRP on VLAN profile—Select the VRRP in VLAN profile if you want to configure VRRP on an integrated routing and bridging (IRB) interface. The VRRP attributes such as group ID and priority are applied to the device during the profile assignment.

This topic describes:

Managing VRRP Profiles

From the Manage VRRP Profiles page, you can:

- Create a new profile by clicking **Add**.
- Modify an existing profile by selecting it and clicking **Edit**.
- View information about a profile by selecting the group and clicking **Details** or by clicking the profile name.
- Clone a profile by selecting a profile and clicking **Clone**.
- Delete profiles by selecting the profiles and clicking **Delete**.

TIP: You cannot delete profiles that are in use—that is, assigned to objects or used by other profiles. To see the current assignments for a profile, click the profile name

The following table describes the information provided about VRRP profiles on the Manage VRRP Profiles page. This page lists all VRRP profiles defined for your network, regardless of your current selected scope in Network view.

Table 189: Managing Profiles


Field	Description
Profile Name	Name given to the profile when the profile was created.

Table 189: Managing Profiles (continued)

Field	Description
Description	<p>Description of the profile that was entered when the profile was created. If the profile was created by using the CLI and then discovered by Network Director, the description is: <i>Profile created as part of device discovery.</i></p> <p>NOTE: To display the entire description, you might need to resize the Description column by clicking the column border in the heading and dragging it.</p>
Family Type	The device family on which the profile was created: EX Series Switches, Campus Switching ELS, or Data Center Switching.

Creating VRRP Profiles

To create VRRP profiles for EX Series switches, Campus Switching ELS, Data Center Switching, or Data Center Switching ELS:

1. Click  in the Network Director banner.
2. Under Views, select one of the following views: **Logical View**, **Location View**, **Device View**, or **Custom Group View**.
3. Under **Tasks**, expand **Wired** and click **VRRP**.
The Manage VRRP Profiles page opens.
4. Click **Add**.
The Device Family Chooser appears.
5. Select **Switching (EX)**, **Campus Switching ELS**, **Data Center Switching ELS**, or **Data Center Switching Non-ELS**.
The Create VRRP Profile page appears for the selected family with the appropriate fields for configuring that family.

Specifying VRRP Settings for an EX Switching, Campus Switching ELS, or Data Center Switching ELS or non-ELS

Use the Create VRRP Profile page to define a common set of VRRP attributes, which you can then apply to a group of interfaces. These directions address creating a VRRP profile for EX Series switches.

Table 190: VRRP Profile Settings

Field	Action
Profile Name	Type the name of the profile. You can use up to 64 characters in the profile name of profiles created for wired devices. Profile name must not contain special characters or spaces. Note that profiles that are automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character in the profile name.
Description	Type a description for the profile.
VRRP Configuration Settings	
Family	Select the IPv4 or IPv6 address family.
VRRP Group Identifier [0 - 255]	Select the VRRP group identifier, which identifies the virtual routing device where the packet is routed to. Each VRRP group is identified by a unique virtual identifier. MAC addresses ranging from 00:00:5e:00:01:00 through 00:00:5e:00:01:ff are reserved for VRRP groups, as defined in RFC 2338. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address, which is 0 through 255.
Advertise Interval ([1-255] secs for IPv4/ [100-40000] msec for IPv6):	Configure the interval in milliseconds between VRRP IPv4 and IPv6 advertisement packets.
Fast Interval in msec [10-40950]	Configure the interval, in milliseconds, between VRRP advertisement packets. All devices in the VRRP group must use the same advertisement interval. Range: 100 through 999 milliseconds Default: 1 second
Authentication Type(for IPv4)	<ul style="list-style-type: none"> • simple—Use a simple password. The password is included in the transmitted packet. • md5—Use the MD5 algorithm to create an encoded checksum of the packet.
Authentication Key(for IPv4)	Configure a VRRP IPv4 authentication key or password. You also must specify a VRRP authentication scheme by including the authentication-type statement. All devices in the VRRP group must use the same authentication scheme and password. For simple authentication, the password can contain 1 through 8 characters. For MD5 authentication, it can contain 1 through 16 characters. If you include spaces, enclose all characters in quotation marks (" ").

Table 190: VRRP Profile Settings (continued)

Field	Action
Preempt	Determine whether or not a backup device can preempt a master device: When no-preempt is configured, the backup device cannot preempt the master device even if the backup device has a higher priority.
Hold Tim in secs [0 - 3600]	Set the hold time before a higher-priority backup device preempts the master device. Range: 0 through 3600 seconds Default: 0 seconds
Accept Data	Determine whether or not an interface accepts packets destined for the virtual IP address This feature helps to debug connectivity issues by making devices respond to ping packets on virtual IP.
Virtual Link Local Address (IPv6)	Configure a virtual link local address for the VRRP IPv6 groups. You must explicitly define a virtual link local address for each group. The virtual link local address must be in the same subnet as the physical interface address.
Virtual IP Addresses (IPv4)	
IP Addresses	The addresses of the virtual routers in a VRRP IPv4 group. You can configure up to eight addresses. Do not include a prefix length. If the address is the same as the interface's physical address, the interface becomes the master virtual device for the group.
Virtual IP Addresses (IPv6)	
IP Addresses	The addresses of the virtual routers in a VRRP IPv6 group. You can configure up to eight addresses. Do not include a prefix length. If the address is the same as the interface's physical address, the interface becomes the master virtual device for the group.

RELATED DOCUMENTATION

[Understanding VRRP Profiles | 845](#)

[Creating and Managing Port Profiles | 414](#)

[Creating and Managing VLAN Profiles | 502](#)

Configuring Wireless Access Points and Radios

IN THIS CHAPTER

- Understanding Access Point Bias for Controllers | 852
- Understanding Wireless Radio Channels | 856
- Understanding WMM Power Save and WLAN Client Battery Life | 859
- Understanding Adaptive Channel Planner | 861
- Understanding Auto Tune Power Policy for Wireless Radios | 866
- Understanding Wireless Scanning | 869
- Understanding Distributed Access Point Behavior on a Layer 3 Network | 873
- Understanding How To Add Access Points to a Wireless Network By Using Network Director | 878
- Understanding Radio Profiles | 879
- Understanding Auto AP Profiles | 883
- Understanding WLAN Service Profiles | 885
- Understanding Wireless Mesh | 894
- Understanding Wireless Encryption and Ciphers | 899
- Understanding PSK Authentication | 903
- Understanding Web Portals | 905
- Understanding Local Switching on Access Points | 907
- Understanding Wireless Bridging | 912
- Understanding Wireless Interference | 914
- Understanding Rogue Access Points | 917
- Understanding Rogue Clients | 923
- Understanding an SSID Masquerade | 926
- Understanding Ad-Hoc Networks | 927
- Understanding LLDP and LLDP-MED | 930
- Importing RingMaster Data to Network Director | 931
- Creating and Managing a Radio Profile | 932
- Assigning a Radio Profile to Radios | 952
- Specifying Custom Radio Profile Setup Settings | 957
- WLAN Setup | 970

- [Configuring Wireless Mesh and Bridging | 976](#)
- [Creating and Managing Wireless Auto AP Profiles | 980](#)
- [Assigning an Auto AP Profile to Controllers | 991](#)
- [Understanding Bonjour | 995](#)
- [Creating and Managing mDNS Profiles | 997](#)
- [Assigning an mDNS Profile to Devices | 1001](#)
- [Creating and Managing an mDNS VLAN List | 1002](#)
- [Creating and Managing Local Switching Profiles | 1005](#)
- [Assigning a Local Switching VLAN Profile to Existing Access Points | 1011](#)
- [Assigning a Local Switching Profile During Access Point Configuration | 1013](#)
- [Creating and Managing Remote Site Profiles | 1014](#)
- [Assigning Remote Site Profiles to Access Points | 1024](#)
- [Creating and Managing RF Detection Profiles | 1026](#)
- [Assigning RF Detection Profiles to Controllers | 1033](#)
- [Configuring Link Layer Discovery Protocol \(LLDP\) on an Access Point | 1035](#)

Understanding Access Point Bias for Controllers

IN THIS SECTION

- [How Do I Determine the Bias Settings I should Use? | 853](#)
- [Example of a Layer 3 Network With Multiple Controllers | 854](#)
- [A Third Option for Access Point Bias: Sticky | 854](#)
- [How Do I Set the Controller Bias for an Access Point? | 855](#)
- [How Can I Determine the Bias of an Access Point by Looking at a Controller? | 855](#)
- [What About Controller Clusters? | 855](#)

In wireless networks, distributed access points can have a bias for different controllers, which means they can have a preference for some controllers. Bias is used only when access points are distributed over the network, not when access points are directly plugged into a configured port on a controller. Bias does not effect the initial controller discovery process performed by the access point. Instead, for distributed access points only, bias is referenced by the controllers and the first controller an access point finds relays the bias settings to an access point.

This topic describes:

How Do I Determine the Bias Settings I should Use?

Refer to [Table 191](#) to determine bias settings for access points.

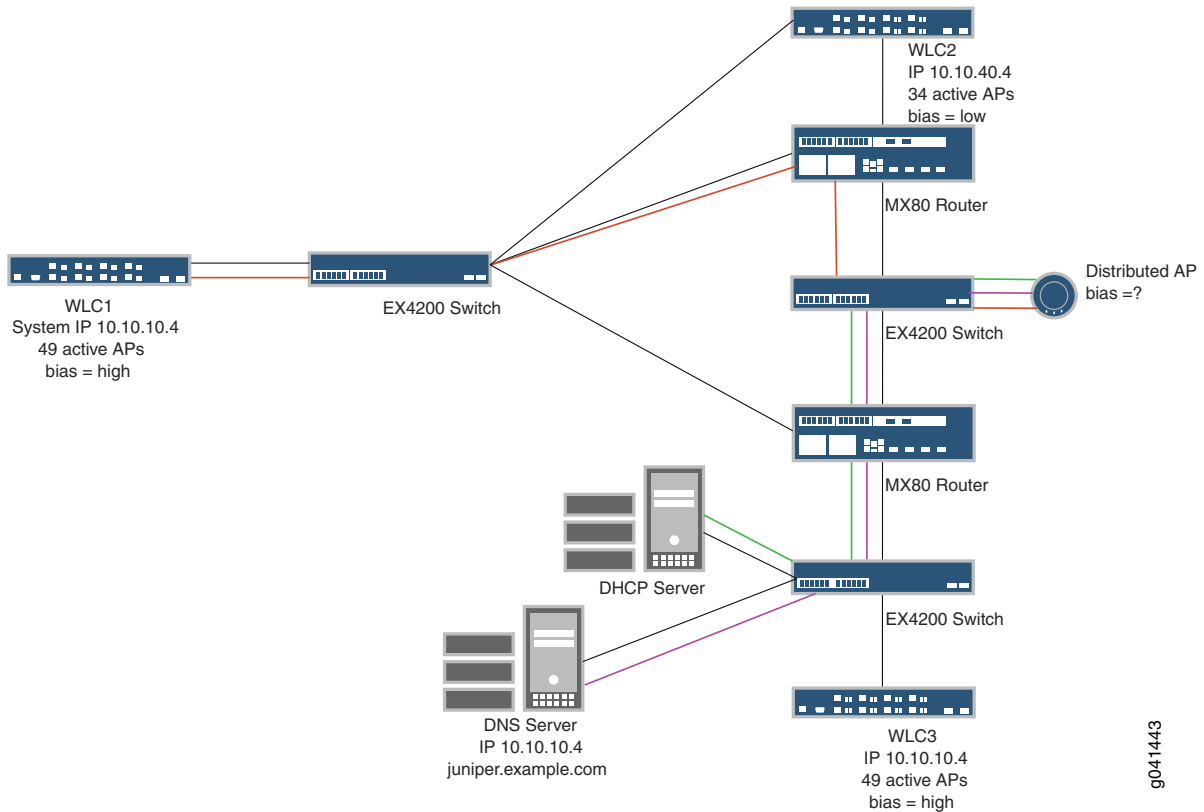
Table 191: Configuring Bias on Controllers and Access Points

Situation	What You Should Do	Result
Access point is directly wired and connected to a controller. Controller ports are configured as direct attach ports.	Bias settings are ignored in this case. You do not need to configure the access point bias. If you do configure it, that configuration is ignored.	The access point associates directly to the connected controller. An access point always attempts to boot on access point port 1 first, and if a controller is directly attached on access point port 1, the access point boots from there regardless of the bias settings.
There is one controller on the network.	Configure access points as usual. Bias will have no effect since there is only one controller.	Access points will associate to the controller, using one of these methods to discover the controller: static configuration, DHCP option 43, DNS, or subnet directed broadcast.
There are three controllers on the network, set up as shown in Figure 34 .	In this situation, bias will make a difference. The access point has a high bias for WLC1 and WLC3, and a low bias for WLC2.	<p>Since the controllers are in another subnet, the access points locate a controller and receive the bias settings and the capacity of each controller.</p> <p>Access points are then directed to associate to a high bias controller (WLC1 or WLC3 in this case) if one is available .</p> <p>When the bias for two controllers is the same, as it is in this illustration, the controller with the greatest capacity to add more active access points is selected.</p>

Example of a Layer 3 Network With Multiple Controllers

This Layer 3 network would make use of bias settings since there are several controllers that access points can associate to.

Figure 34: Controllers With High Bias and Low Bias for an Access Point



A Third Option for Access Point Bias: Sticky

Access points actually have three bias options for controllers, high, low, and sticky. The third option, sticky, means that an access point stays on whatever controller it is associated to, regardless of whether a high bias controller is available or not. Even if a sticky access point is using a low bias controller when a high bias controller becomes available, the sticky access point stays with the low bias controller. The sticky option is useful because it avoids unscheduled access point resets that would normally occur as the access point transitions from one controller to another.

An access point with a sticky bias stays on the controller until an administrator intervenes by resetting the access point to have it discover and connect to a high bias controller.

How Do I Set the Controller Bias for an Access Point?

When you add an access point, you must first select a controller—this controller is the one you are setting the bias for. If you want the access point to be able to use different controllers, you must set the bias for each controller. You would add the access point to each controller following the directions [“Adding and Managing an Individual Access Point” on page 1156](#).

How Can I Determine the Bias of an Access Point by Looking at a Controller?

If an access point is associated with a controller, it appears on the list you see when you click **Wireless Network > View Inventory** (under Device Management) in Network Director. If the access point is using that controller, the Connection field on the controller will indicate Up. If the controller’s Connection field indicates Up/Redundant, it means that cluster/network resiliency is enabled. In that case, each access point has a primary link to one controller and a backup/redundant link to another controller. The redundant controller connection is also listed as Up/Redundant.

What About Controller Clusters?

Clusters are a subset of a mobility domain. Clustering creates a logical group of controllers (and their associated access points) that share network and user information for failover support. Every access point in a cluster has a primary access point manager—this can be any controller in the cluster. For more information, see [“Creating a Mobility Domain for Wireless LAN Controllers” on page 1053](#).

The concept of bias is superseded in cluster configuration scenarios by automatic redundancy and/or optional access point affinity configuration. Access point affinity is configured per controller in a cluster, so bias has no bearing or effect.

RELATED DOCUMENTATION

[Adding and Managing an Individual Access Point | 1156](#)

[Creating a Mobility Domain for Wireless LAN Controllers | 1053](#)

[Understanding Distributed Access Point Behavior on a Layer 3 Network | 873](#)

[Network Director Documentation home page](#)

Understanding Wireless Radio Channels

IN THIS SECTION

- [What WLAN Channels Are Available? | 856](#)
- [What Channels Are not Available? | 857](#)
- [How Do I Know Which Channels I Should Use? | 857](#)
- [How Do I Avoid Co-Channel Interference? | 858](#)
- [DFS Channels | 858](#)
- [802.11n Channels can be Wider and Work on Both Bands | 858](#)
- [How Are Channel Numbers Assigned? | 859](#)
- [How Do I Know What Channel an Access Point Is Using? | 859](#)

Each wireless radio operates on a configured radio frequency (RF) channel identified by numbers. A radio assigned to a particular channel both transmits and receives all traffic on that channel.

Depending upon the network configuration, some channels might have less interference than others. Choosing the right channel lets you optimize performance.

This topic describes:

What WLAN Channels Are Available?

There are 14 channels designated for wireless networks in the 2.4-GHz frequency band and 42 channels in the 5-GHz frequency band.

The 14 channels in the 2.4-GHz band are spaced 5 MHz apart. The protocol requires 25 MHz of channel separation, which means that it is possible for adjacent channels to overlap and then interfere with each other. For this reason, only channels 1, 6, 11 are typically used in the US to avoid interference. In the rest of the world, the four channels 1, 5, 9, 13 are typically recommended. The 2.4-GHz frequency band is heavily used because most devices can operate on that band.

The 5-GHz band is actually four frequency bands: 5.1-GHz, 5.3-GHz, 5.4-GHz, and 5.8-GHz. The 5-GHz band has a total of 24 channels with 20 MHz bandwidth available. Unlike the 2.4-GHz band, the channels are non-overlapping, therefore all channels have the potential to be used in a single wireless system. Formerly, only 802.11a devices used this band, but now this band is used for the newest 802.11ac technology.

What Channels Are not Available?

Because each country has different regulatory requirements, the country code determines which channels you can configure on the radios. When you specify the country of operation for an access point, the radios are restricted to using the valid channels for that country.

The FCC (United States) requires that devices operating the 5-GHz band must employ dynamic frequency selection (DFS) and transmit power control (TPC) capabilities. This is to avoid interference with weather-radar and military applications. Additional channels in the 5-GHz band are restricted to avoid interference with Terminal Doppler Weather Radar (TDWR) systems. This eliminates the use of channels 120, 124, and 128. Channels 116 and 132 can be used if they are separated by more than 30 MHz (center-to-center) from a TDWR located within 35 km of the device.

How Do I Know Which Channels I Should Use?

For best performance, choose a channel at least 5 channels apart from your neighbors' networks. Determine this by completing a site survey—a site survey includes a test for RF interference.

Try to use non-overlapping channels (1, 6, 11 typically in the US), or minimize overlap of signals by using channels as far apart as possible from other networks in range.

You also need to know what channel your clients are capable of using, so that you can provide connection for that channel. For example, you can use Microsoft Windows 7 Device Manager to find the channel number at which your Microsoft Windows client can operate by following these steps:

1. Click **Start > Control Panel**.
2. Click **Device Manager > Network Adapters**.
3. Right-click the link describing Wi-Fi and select **Properties** from the menu.
4. Click the **Advanced** tab.

Two columns are displayed: Property and Value. The list under Property tells you the wireless capabilities of the computer. The current operating properties are highlighted—this tells you the current mode of operation.

Once you know the networking transmission standards used by clients (802.11n, 802.11g, 802.11b, or 802.11a) you can determine which channels the device can use—see [Table 192](#) for details.

Table 192: Channels a Device Can Use

Device Capability	Band and Channel Width used:	Typical Channel Use
802.11b	2.4-GHz band, 20 MHz channel width	1, 6, 11 (US) or 1, 5, 9, 13

Table 192: Channels a Device Can Use (continued)

Device Capability	Band and Channel Width used:	Typical Channel Use
802.11g	2.4-GHz band, 20 MHz channel width	1, 6, 11 (US) or 1, 5, 9, 13
802.11n	2.4-GHz band, 20 MHz channel width 5-GHz band, 40 MHz channel width	1, 6, 11 (US) or 1, 5, 9, 13 (36,1) (40,-1) (44,1) (48,-1) (52, 1) (56,-1) (60,1) (64,-1) (100,1) (104,-1) (108,1) (112,-1) (116,1) (120,-1) (124,1) (128,-1) (132,1) (136,1) (149,1) (153,-1) (157,1) (161,-1)
802.11a	5-GHz band, 40 MHz channel width	3, 11

How Do I Avoid Co-Channel Interference?

Prevent interference and signal overlapping by doing an initial site survey of your wireless spectrum before deploying a wireless network. Once you discover all nearby signals, you can choose an optimum wireless channel that will provide the best performance with the least signal interference.

To automatically avoid co-channel interference, enable automatic channel tuning in a Radio profile—all radios using that Radio profile will then use algorithms to switch to optimum channels. See [“Creating and Managing a Radio Profile” on page 932](#) to enable automatic channel tuning.

DFS Channels

In countries where Dynamic Frequency Selection (DFS) is required, Juniper Networks devices perform the appropriate check for radar on the 802.11a band. If radar is detected on a channel, the access point radio stops using the channel for the amount of time specified in the country’s regulations. Log messages are generated when this occurs. To enable DFS, see [“Creating and Managing a Radio Profile” on page 932](#).

802.11n Channels can be Wider and Work on Both Bands

802.11n and 802.11ac devices work on the 5-GHz radio band as well as the more-populated 2.4-GHz radio band. On the 5-GHz band, 802.11n and 802.11ac channels can be either 80 MHz, 40 MHz or 20 MHz wide. This is one reason that 802.11n and 802.11ac devices are faster than 802.11a/b/g/n devices.

NOTE: 40 MHz channels work only with 802.11n and 802.11ac on the 5-GHz band—40 MHz channels cannot not be configured on a 2.4-GHz radio.

802.11n and 802.11ac radios configured for the 5-GHz band have a primary channel and a secondary channel. The primary channel is listed using the channel number, and the secondary channel adds another 20 MHz to make the channel 40 MHz. Therefore, the secondary channel is either the channel above (1) or the channel below (-1) the primary channel. This notation keeps you from configuring non-contiguous channels. For example, if the primary channel is 36 and the secondary channel is 40, the combination would be (36,1). If the primary channel is 44 and the secondary channel is 40, the notation would be (44,-1).

The following channels, listed as (primary channel, secondary channel) are supported for 802.11n at 5-GHz with 40 MHz bandwidth: (36,1) (40,-1) (44,1) (48,-1) (52, 1) (56,-1) (60,1) (64,-1) (100,1) (104,-1) (108,1) (112,-1) (116,1) (120,-1) (124,1) (128,-1) (132,1) (136,1) (149,1) (153,-1) (157,1) (161,-1)

How Are Channel Numbers Assigned?

Juniper Networks access point radios use channel auto-tuning by default. You can change the channel tuning interval, channel tuning range, or the channel tuning holddown settings. You can also turn off auto-tuning. For more information, see [“Understanding Adaptive Channel Planner” on page 861](#) and to turn off channel auto-tuning, see [“Creating and Managing a Radio Profile” on page 932](#).

How Do I Know What Channel an Access Point Is Using?

If the access point is already installed and operating, use the Equipment tab of the monitor mode in Network Director to view the access point channel numbers.

RELATED DOCUMENTATION

[Creating and Managing a Radio Profile | 932](#)

[Adding and Managing an Individual Access Point | 1156](#)

[Understanding Adaptive Channel Planner | 861](#)

[Network Director Documentation home page](#)

Understanding WMM Power Save and WLAN Client Battery Life

IN THIS SECTION

- [How Does WMM Power Save Extend Battery Life? | 860](#)
- [Where is WMM Defined? | 860](#)
- [How is WMM Power Save Implemented on Juniper Networks WLANs? | 861](#)

- [WMM Power Save is Disabled by Default | 861](#)
- [Why Should I Enable WMM Power Save in a Radio Profile? | 861](#)

Wi-Fi mobile devices need as much battery power as possible because the power demands of voice, audio, and video applications are ever-increasing. To maintain laptop battery life in these situations, WMM Power Save was certified by the Wi-Fi Alliance. WMM Power Save is an addition to WMM, the technology that enables Quality of Service (QoS) functionality in Wi-Fi networks by prioritizing traffic from different applications. WMM is a required feature for 802.11n capable devices—almost all modern Wi-Fi devices support it.

WMM Power Save is optimized for latency-sensitive applications such as voice, audio, or video, but benefits any Wi-Fi device. With WMM Power Save, the same amount of data can be transmitted in fewer frames in a shorter time, while enabling the Wi-Fi device to preserve power in a low-power, dozing state in between transmissions.

This topic describes:

How Does WMM Power Save Extend Battery Life?

Power save uses mechanisms from 802.11e and legacy 802.11 to save power for battery powered equipment and fine-tune power consumption. Products targeted for power-critical applications such as mobile phones, smart phones, and other portable power devices typically use WMM Power Save.

The underlying concept of WMM Power Save is that the client triggers the release of buffered data from the access point by sending an uplink data frame. Upon receipt of that data trigger frame, the access point releases previously buffered data stored in each of its four queues. Queues can be configured to be either trigger enabled (receipt of a data frame corresponding to each queue acts as a trigger), or delivery enabled (data stored in all queues is released upon receipt of a frame).

WMM operates by dividing traffic into four access categories: background, best effort, video, and voice. QoS policy determines the different handling of each access category. The result is that different packets are handled differently.

Where is WMM Defined?

WMM was a precursor to the 802.11e standard. Before the 802.11e standard could be ratified, some organizations agreed on and published a draft standard called WMM. Once the 802.11e standard was finalized, WMM became an enhancement to the 802.11e standard, and was referred to as 802.11e quality-of-service (QoS) enhancements. Both the original WMM standard and 802.11e are now deprecated, but the industry continues to use that terminology.

How is WMM Power Save Implemented on Juniper Networks WLANs?

To take advantage of WMM Power Save functionality, both the Wi-Fi client and the access point must be Wi-Fi CERTIFIED for WMM Power Save. In addition, the applications used also need to support WMM Power Save to inform the client of the requirements of the traffic they generate.

WMM is supported by Radio profiles, along with the corresponding QoS policies that describe access classes. When a wireless client using WMM Power Save associates with an access point using a Radio profile that includes WMM, the client selects the access classes (voice, video, best effort, background) in WMM Power Save.

WMM Power Save is Disabled by Default

WMM Power Save is disabled by default on access points, even though it saves client battery life, because clients using Power Save must send a separate PSpoll to retrieve each unicast packet buffered by the access point radio. This can sometimes slow performance, depending on the network configuration. Also, your applications need to support WMM Power Save to inform the client of the requirements of the traffic they generate.

Why Should I Enable WMM Power Save in a Radio Profile?

WMM Power Save preserves client battery life, especially for applications such as voice and video. Wi-Fi CERTIFIED for WMM Power Save devices can operate in any Wi-Fi network and coexist with 802.11 legacy power save mechanisms.

RELATED DOCUMENTATION

[Creating and Managing a Radio Profile | 932](#)

[Network Director Documentation home page](#)

Understanding Adaptive Channel Planner

IN THIS SECTION

- [Why Use Adaptive Channel Planner? | 862](#)
- [When Should I Use Adaptive Channel Planner? | 862](#)
- [How Does Adaptive Channel Planner Work? | 863](#)
- [How Do I Configure Adaptive Channel Planner? | 864](#)

- [What Are Adaptive Channel Planner Results? | 864](#)
- [When Is Adaptive Channel Planner Most Beneficial? | 865](#)
- [What Happens When Severe Interference Is Detected? | 866](#)

Adaptive Channel Planner automatically makes channel-tuning decisions for access point radios on the basis of the RF data gathered by access points. Channel auto-tuning is configured in a Radio profile in Network Director and operates on all access point radios by default. This topic explains channel automatic tuning.

This topic describes:

Why Use Adaptive Channel Planner?

The state of a wireless network is dynamic—to continually tune each access point radio for optimum performance would be time consuming. Auto-tuning performs the following tuning based on feedback from access point scanning:

- Chooses a random starting channel for newly deployed access points
- Automatically tunes channels of access points to minimize co-channel interference
- Mitigates spectral interference

You can specify which channels auto-tuning can choose from.

In addition, auto-tuning maintains tuning configurations across access point and controller reboots. For more information about scanning, see [“Understanding Wireless Scanning” on page 869](#).

When Should I Use Adaptive Channel Planner?

Channel auto-tuning is beneficial in these situations:

- When you are deploying a new network—no channel configuration exists yet.
- When you add or move access points.
- When the performance of an existing configuration needs improvement.
- When the network is experiencing interference.
- When you need to tune around channels used by radar when there is either weather or military radar operating in your area. Weather radar frequently operates at large commercial airports.
- When interference makes a channel unusable. For example, a continuous wave transmitter, like some video surveillance cameras, will make a channel unusable.

Adaptive Channel Planner Improves Performance

Auto-tuning optimizes performance—you do not need to make repeated channel changes in a reactive mode. Channel tuning for performance operates best over a regular time frame, such as a day. We recommend that you schedule auto-tuning to collect data over a long period of time, then make changes at a time when only a few users are affected.

The duration of the sample period must match the retuning period so that the algorithm does not generate a reactive response to short term-transients. For example, the whole network should not retune each time a door opens or closes. Instead, it is better to choose and maintain tuning that takes into account how much of the time the door is open and closed (average samples over a long period). Changing channels can be disruptive to active clients—therefore, it is best not to do it frequently and at times when service disruptions can be tolerated.

Adaptive Channel Planner Resolves Interference

If a channel becomes unusable, the wireless auto-tuning algorithm immediately changes the radio to a working channel but the changes do not persist. The algorithm tries to revert to the configured channel at intervals. This emergency tuning is similar to DFS channel handling.

Adaptive Channel Planner Is Used by Dynamic Frequency Selection to Comply with Country Regulations

If an access point detects radar on a channel it is currently using, regulations require that it cease using the channel immediately. Auto-tune enables the radio to be switched to a usable channel rather than taking the radio out of service.

Dynamic Frequency Selection (DFS) requirements apply only to radios operating in the 5-GHz band (802.11a radios). Auto-tuning enables the system to retune channels in the 5-GHz band when radar is detected and regulations require the system to cease using a channel.

When a controller learns that an access point 802.11a radio has detected radar on a channel, Auto-tune immediately switches the radio to another channel. The event is time stamped along with the state of the channel. For the next 30 minutes, the controller blocklists the radar channel—after this, the channel is automatically added back to the eligible channel list.

NOTE: If radar is detected on a radio with its Auto-tune channel disabled, the radio goes out of service as required by the Federal Communications Commission (FCC) and European Telecommunications Standards Institute (ETSI).

DFS is configured in a Radio profile—see [“Creating and Managing a Radio Profile”](#) on page 932.

How Does Adaptive Channel Planner Work?

The controller evaluates the radios' scanning results for possible channel changes every 3600 seconds (1 hour). An algorithm running as a distributed algorithm on controllers and access points uses the following parameters to determine whether to change the channel on a radio:

- Amount of noise on the channel
- Packet retransmission count—see [“Monitoring the Percentage of RF Packet Retransmissions” on page 1338](#) for more information.
- Utilization calculated on the number of multicast packets per second that a radio can send on a channel while continuously sending fixed-size frames over a period of time.
- Phy error count, which is the number of frames received by the radio with physical layer errors. A high number of Phy errors can indicate the presence of a non-802.11 device using the same RF spectrum.
- Received cyclic redundancy check (CRC) error count. A high number of CRC errors can indicate a hidden node or co-channel interference.

The thresholds for these parameters are not configurable. RF auto-tuning also can change a radio channel when the channel tuning interval expires, if a channel with less disturbance is detected. Disturbance is based on the number of neighbors the radio has and the RSSI of each neighbor. A radio also can change channels before the channel expires to respond to interference or radar.

How Do I Configure Adaptive Channel Planner?

Adaptive Channel Planner is an Advanced Setup option of Radio profiles—see [“Creating and Managing a Radio Profile” on page 932](#). In addition to disabling and enabling Tune Channel, you can change these settings:

- Tune Transmit Power
- Ignore Clients
- Tune Channel Range (11a)
- Channel Tuning Interval
- Channel Tuning Holddown
- TX Power Backoff Timer
- Power Tuning Interval
- Power Ramp Interval

For descriptions of these settings and directions for creating Radio profiles, including automatic channel tuning, see [“Creating and Managing a Radio Profile” on page 932](#).

What Are Adaptive Channel Planner Results?

Because auto-tuning automatically optimizes channels used by an access point, the following results occur:

- Co-channel interference is minimized. Co-channel interference occurs when multiple neighboring radios use the same frequency (channel). Excessive co-channel interference occurs if a poor channel plan is used. Auto-tune channel optimizes the channel plan such that nearby access points avoid using the same channel.
- Spectral interference from non-802.11 devices, such as instruments, microwaves, cordless phones, and surveillance cameras, is mitigated.
- The interference domain is defined. An interference domain is a set of radios in a mobile domain that can interfere with one another. Interference domains are temporary and non-configurable.
- Tuning is maintained across access point and controller reboots.
- Newly deployed access points are assigned a random starting channel. This also happens at the first invocation. The purpose is to quickly choose a reasonable channel rather than waiting for the tuning algorithm to do its next update. It turns out that just choosing random channels for all radios usually results in a reasonably good tuning. The old method of setting the channel to a single default channel guaranteed that the system started at the worst possible tuning—all radios on the same channel.

NOTE: Keep in mind Bluetooth devices and frequency hopping cordless phones spread their communications over the entire 2.4-GHz frequency band, so it is not possible to avoid them by changing channels. However, they are designed to coexist with other transmitters, so their interference is minimal.

When Is Adaptive Channel Planner Most Beneficial?

The most beneficial time to turn on auto-configuration is when you first deploy your wireless network or after you make significant changes to your network or facilities. Turn on auto-configuration of channels when you:

- First deploy your wireless network or after you make significant changes to your network or facilities.
- Deploy new access points that have not had channels configured.
- Add more access points to an existing configuration or move existing access points, for example during a remodel.
- Discover that a neighboring organization deployed or changed a wireless network in close proximity to some of your access points.
- Have reason to believe the current tuning is sub-optimal, for example when users experience poor service in certain locations.

During regular network operation, there is usually no need to retune channels frequently. A working network should be in data collection mode most of the time. It is better to configure channel Auto-tune to collect data over a long period of time and then make changes at a time when only a few users are

affected. If you do not have a time when disruptions can be tolerated, you might want to run it only occasionally, on demand.

What Happens When Severe Interference Is Detected?

Severe interference usually occurs only when something changes in the interference domain. For example, a neighboring company could reconfigure their wireless channels, or a new source of interference, such as a continuously transmitting video surveillance camera, is deployed.

If RF scanning detects severe interference, auto-tuning immediately chooses a new temporary channel assignment to mitigate the problem. Auto-tuning becomes active, even if you have configured it to wait. If the interference persists, auto-tune takes it into account the next time it retunes channels. In this case, auto-tuning uses the tuning algorithm to change to a usable channel if one is available. Otherwise, the channel will revert back to the previously tuned channel. The difference with auto-tuning done in a crisis is that the emergency changes are not saved, and Auto-tune will periodically try to revert the network to the saved channels. Only changes made by you are saved.

RELATED DOCUMENTATION

[Creating and Managing RF Detection Profiles | 1026](#)

[Creating and Managing a Radio Profile | 932](#)

[Understanding Auto Tune Power Policy for Wireless Radios | 866](#)

[Configuring a Controller | 1037](#)

[Network Director Documentation home page](#)

Understanding Auto Tune Power Policy for Wireless Radios

IN THIS SECTION

- [How Does Wireless Transmit Power Work? | 867](#)
- [How Does Auto Tune Power Policy Work? | 867](#)
- [When is Auto Tune Power Policy Most Helpful? | 868](#)
- [How Do I Turn Off an Auto Power Policy? | 868](#)
- [What Changes Can I Make to an Auto Tune Power Policy? | 868](#)

The amount of power an access point uses affects the coverage area of the wireless network. The higher the power level on access points, the larger the coverage area of a wireless network. Usually, you want your wireless network to cover all areas, but with minimal overlap between access points that share the same channel—this minimizes co-channel interference. Configuring each access point radio's power manually can be time consuming and tedious for a large installation. For one thing, in order to ensure complete coverage while minimizing co-channel interference, you need to consider nearby access points sharing the same channel as well as understand signal propagation issues. For example, the walls and windows in your facility affect signals. Instead of manually configuring power, you can configure access points to automatically tune their radios' power based on RF data they collect about neighboring access points. With automatic power tuning, which is configured in the Network Director Radio profile under the Power & RF Settings tab, access points adjust their radios' power levels automatically, based on the power levels of all neighboring access points.

Radios get most configurations, including power tuning, from the associated Radio profile. For directions to create a Network Director Radio profile, see [“Creating and Managing a Radio Profile” on page 932](#).

NOTE: Automatic channel tuning is also available, and automatic power tuning and channel tuning can be used together. For more information about automatic channel tuning, see [“Understanding Adaptive Channel Planner” on page 861](#).

This topic describes:

How Does Wireless Transmit Power Work?

Transmit power, like other sound pressure, is measured in decibels. Because the measurement is logarithmic instead of linear, an increase of 6 dB will double the range of coverage of a radio. Valid power values depend on the country of operation. The default transmit power on all access point radio types is either the highest setting allowed for the country of operation or highest setting supported on the hardware, whichever is lower.

How Does Auto Tune Power Policy Work?

Power tuning computation is performed on the access point itself without any help from the controller. Access points listen for nearby access points on the same channel and then adjust their power to provide good coverage while avoiding co-channel interference.

The power tuning algorithm automatically adjusts to changes when needed, for example when one access point is removed.

Auto-tuned power settings are not persistent—auto-tuning must be enabled for the changes to occur. If you turn off auto-tuning, the access point will go back to the configured power setting.

When is Auto Tune Power Policy Most Helpful?

When an access point radio is first enabled, automatic power tuning can assign initial power settings compatible with surrounding access points.

How Do I Turn Off an Auto Power Policy?

You can turn off automatic power tuning by editing Radio Profiles—this turns off auto power for all access points using the Radio Profile. See [“Creating and Managing a Radio Profile” on page 932](#). You can also turn off automatic power tuning for a single radio when you add or edit an individual access point—for directions, see [“Adding and Managing an Individual Access Point” on page 1156](#).

What Changes Can I Make to an Auto Tune Power Policy?

When power tuning is enabled in a Radio profile, you can change the wireless transmit power backoff timer, power tuning interval and the power ramp interval for that Radio profile.

- *Transmit Power Backoff Timer* changes the interval at which radios reduce power after temporarily increasing the power to maintain the minimum data rate for an associated client. At the end of each power-backoff interval, radios that temporarily increased their power reduce it by 1 dBm every 10 seconds. The power backoff continues in 1 dBm increments after each interval until the power returns to expected setting.
- *Power Tuning Interval* is the number of seconds between reevaluations of power. Power changes can only take place after an evaluation or when an anomaly occurs. You can change the wait interval between evaluations from the default 600 seconds.
- *Power Ramp Interval* is the rate at which power is increased or decreased on radios in a Radio profile until the optimum power level calculated by RF auto-tuning is reached. You can change the 1 dBm increment to increase and decrease in larger or smaller steps.

RELATED DOCUMENTATION

[Creating and Managing a Radio Profile | 932](#)

[Understanding Adaptive Channel Planner | 861](#)

[Network Director Documentation home page](#)

Understanding Wireless Scanning

IN THIS SECTION

- [What Is the Difference Between Passive and Active Scanning? | 869](#)
- [What Channels Are Scanned? | 870](#)
- [How Does Scanning Work? | 870](#)
- [What Additional Information Is Learned by an Active Scan? | 871](#)
- [What Happens to Scanned Information? | 872](#)
- [CTS-to-self During Scanning | 872](#)
- [What Is Spectral RF Scanning? | 872](#)

All wireless access point radios continually scan for other RF transmitters. While 802.11b/g/n radios scan in the 2.4-GHz to 2.4835-GHz spectrum, 802.11a radios (and sometimes 802.11n radios) scan in the 5.15-GHz to 5.85-GHz spectrum. There are two scanning methods, passive scanning and active scanning. By default, radios perform both types of scans on all channels allowed by the country of operation, which is the regulatory domain set during initial access point deployment. While both types of scanning are on by default, active scanning is performed only on channels on which local government regulations allow it to transmit. Channels that are not authorized for unlicensed use and channels that require radar detection with dynamic frequency selection (DFS) are excluded from active scanning.

A radio in sentry mode is a dedicated scanner (no data transmission) providing better RF detection because the radio spends more time scanning each channel.

This topic describes:

What Is the Difference Between Passive and Active Scanning?

During passive scans, the radio listens for beacons and probe responses. If you use only passive mode, the radio scans once per second, and audits packets on the wireless network. Passive scans are always enabled and cannot be disabled because this capability is also used to connect clients to access points.

Active scans are enabled by default but can be disabled in a Radio profile. During active scans, the radio sends probe-any requests (probe requests with a null SSID name) to solicit probe responses from other devices. In other words, access points actively look for other devices, in addition to listening for them.

What Channels Are Scanned?

RF scanning can be performed on a variety of different sets of channel ranges or frequencies. The scan can be configured in the Radio profile to scan either operating channels, regulatory channels, or all channels.

NOTE: An access point will never transmit on channels that are not authorized for transmission.

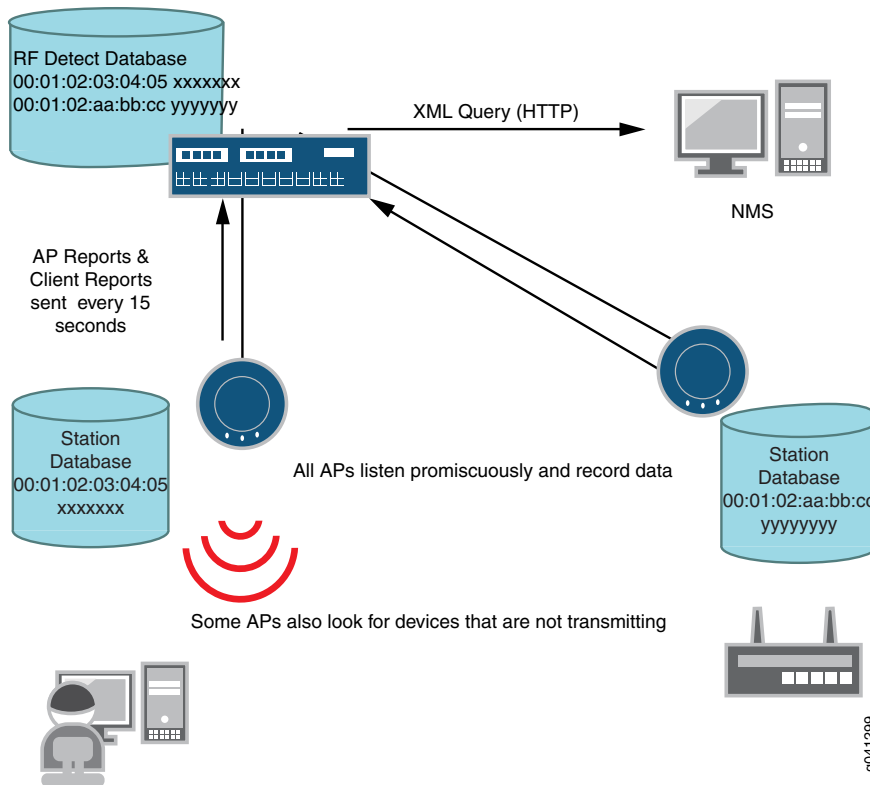
In a Radio profile, you can change the channels a radio actively scans. These are the three options for active scanning :

- **Operating:** Only the current channel is scanned and audited.
- **Regulatory:** Only regulatory channels are scanned and audited. If the radio is configured for 802.11b/g/n, the most commonly used channels, such as 1, 6, or 11, are scanned and audited more frequently.
- **All:** All channels are scanned and audited.

How Does Scanning Work?

To scan outside of the operating range, the access point must change channels. These off-channel scans are performed once per second, and a different channel in the range is scanned each second until it cycles through all in-scope channels. The access point will go off channel for about 30ms (3% of the time). Scans are scheduled to avoid interfering with beacon transmission. Radio transmit queues are drained prior to channel change. Then the probes are sent once channel change is completed. Note that the scan frequency is reduced if voice, video traffic, or heavy load is detected. Also, the CTS-to-self feature can be configured to silence clients on the operating channel while access point goes off channel. See [Figure 35](#) for more details.

Figure 35: Scanned Data Is Stored in the Station Databases and Sent to the Controller



How Does Active Scanning Work?

The active-scan algorithm is sensitive to high-priority (voice or video) traffic or heavy data traffic. Active-scan scans for 30 milliseconds once every second, unless either of the following conditions is true:

- High-priority traffic (voice or video) is present at 64 Kbps or higher. In this case, active-scan scans for 30 milliseconds every 60 seconds.
- Heavy data traffic is present at 4 Mbps or higher. In this case, active-scan scans for 30 milliseconds every 5 seconds.

What Additional Information Is Learned by an Active Scan?

Active scanning is more thorough and provides more information than passive scanning. If you select active mode, the radio actively sends probes on other channels and then audits the packets on the wireless network. The probe response received by an active scan normally contains the BSSID and WLAN SSID of the access point answering the probe. This is how the Permitted SSID list gets a list of all SSIDs on the air, even if a device is not currently sending signals.

What Happens to Scanned Information?

Scanned information is stored and used by the:

- Network security system to identify rogue access points (see [“Understanding Fault Mode in Network Director” on page 1448](#)).
- Network Director reports (see [“Alarm Summary Report” on page 1503](#), [“Audit Trail Report” on page 1505](#), and [“Alarm History Report” on page 1500](#)).
- Automatic access point power feature (see [“Understanding Auto Tune Power Policy for Wireless Radios” on page 866](#)).
- Automatic access point channel feature (see [“Understanding Adaptive Channel Planner” on page 861](#)).
- RF Neighborhood monitoring (see [“Monitoring the RF Neighborhood” on page 1340](#)).
- Spectral RF scanning feature (see [“Monitoring the RF Spectrum of a Radio” on page 1343](#)).
- Spectral RF scanning feature (see [“Monitoring the RF Spectrum of a Radio” on page 1343](#)).

CTS-to-self During Scanning

The clear to scan CTS-to-self feature can be configured in a Radio profile to silence clients on the operating channel while an access point goes off channel to scan. This option is also part of the Radio profile.

What Is Spectral RF Scanning?

The electromagnetic spectrum includes all possible frequencies of electromagnetic radiation. Wireless communication uses the low frequencies used for radio communication, but other objects can affect these frequencies. Spectral analysis reports objects with any electromagnetic properties.

Spectral analysis starts with the RF detect function. The primary function of RF detect is to detect and classify 802.11 devices, but for spectral analysis, that function is extended to recognize non-802.11 sources of interference. Because the primary function of RF detect is to locate 802.11 devices, interfaces are wireless-LAN-centric; for example, frequencies are displayed in terms of 802.11 channels.

In Network Director, two spectrograms are provided under monitoring, one for a radio's channels and one for the results of a spectrum sweep. The second spectrogram also includes graphing of the duty cycle of the radio.

Channel auto-tuning (see [“Understanding Adaptive Channel Planner” on page 861](#)) defines a transmission ID for access point radios. With spectral analysis, that definition is extended to all interference sources.

With spectral monitoring enabled, an enabled access point radio will drop the current clients and scan for any device with a radio signal. For directions to implement spectral monitoring, see [“Monitoring the RF Spectrum of a Radio” on page 1343](#).

RELATED DOCUMENTATION

[Monitoring the RF Spectrum of a Radio | 1343](#)

[Alarm Summary Report | 1503](#)

[Audit Trail Report | 1505](#)

[Alarm History Report | 1500](#)

[Understanding Auto Tune Power Policy for Wireless Radios | 866](#)

[Understanding Adaptive Channel Planner | 861](#)

[Network Director Documentation home page](#)

Understanding Distributed Access Point Behavior on a Layer 3 Network

IN THIS SECTION

- [What Is a Layer 3 Network? | 874](#)
- [How Does an Access Point Find a Controller on a Layer 3 Network? | 875](#)

Access Points can be connected to a wireless network two different ways, either directly to a controller or through a switch. When an access point is connected directly to a controller, operation is straightforward—the access point communicates with the controller by default.

An access point that is connected to a network instead of directly to a controller is referred to as a distributed access point—the operation of a distributed access point is more complex because the access point has to complete a number of steps to find, then connect to, a working controller. This topic describes how a distributed access point boots on a Layer 3 network.

This topic describes:

What Is a Layer 3 Network?

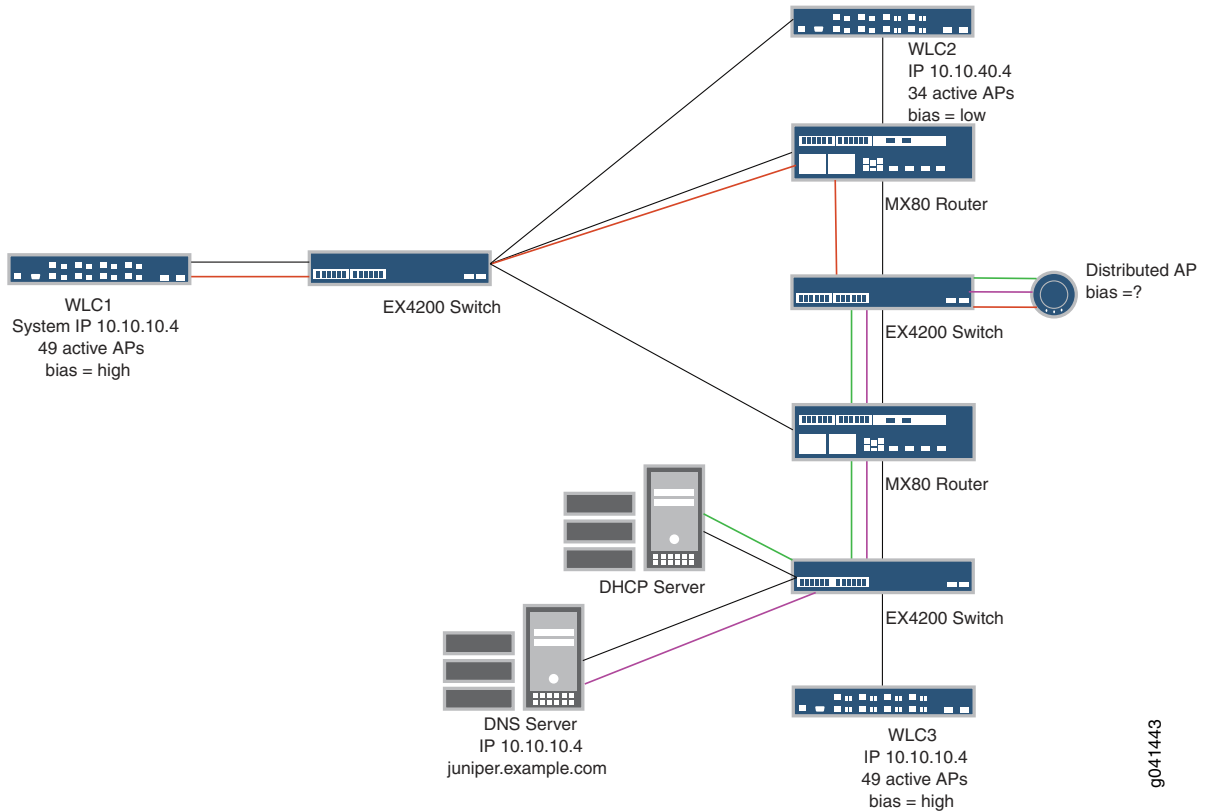
Layer 2 networks forward all broadcast traffic so that any broadcast traffic is transmitted. A large Layer 2 network, forwarding all of this data, can suffer from congestion and decreased network efficiency. In contrast, Layer 3 devices restrict broadcast traffic such as ARP and DHCP broadcasts to the local network. This reduces overall traffic levels by allowing administrators to segment networks into subnetworks and restrict broadcasts to that subnetwork. This means there is a limit to the size of a Layer 2 network, but a Layer 3 network, properly configured, can grow to any size.

This explanation uses the following hardware components to form a typical Layer 3 network:

- One DHCP server
- One DNS server
- Two routers
- Two switches
- Three controllers
- One access point set to Distributed Connection and High Bias for controller1 and controller3. The access point is set to low bias for controller2. For more information about bias, see ["Understanding Access Point Bias for Controllers"](#) on page 852.

[Figure 36](#) illustrates this network topology. As you can see, the access point is in a different subnet than the controllers, so the access point must negotiate its way to one of the controllers. In this case, the access point must also select one of three available controllers using the bias settings on the access point and controllers.

Figure 36: Access Point Booting Over a Layer 3 Network



How Does an Access Point Find a Controller on a Layer 3 Network?

In a Layer 3 network, a distributed access point behaves as described in [Table 193](#). Interactions are color coded as indicated in the table. Interaction begins with the first action in the table (access point sends a discover message to the DHCP server) and ends with a list of SSIDs on the access point.

Table 193: Access Point Connecting to a Controller in a Layer 3 Network

Access Point Action	Response	Color of Route in Figure
<p>Access point sends a discover message from port 1 on the access point to the DHCP server.</p>	<p>DHCP server replies with a DHCP offer message containing:</p> <ul style="list-style-type: none"> • IP address for the access point • Default router IP address for the access point IP subnet • DNS server address • Domain name <p>NOTE: DHCP can optionally provide Option 43 information to the access point, which specifies controller information directly without need for the DNS step.</p>	green
<p>Access point sends a DHCP request message to the DHCP server.</p>	<p>DHCP server sends an ACK to the access point.</p>	green
<p>Access point sends a broadcast find controller message to the IP subnet broadcast address.</p> <p>When the access point is unable to locate a controller on the subnet connected to it, the access point sends a DNS request for wlan-switch.</p>	<p>The DNS server sends the system IP address of the controller mapped to wlan-switch. DNS returns controller1 in this example.</p>	purple
<p>Access point sends a unicast find controller message to controller1.</p>	<p>Controller1 receives the find controller message and compares the bias settings on each controller. More than one controller has a high bias for this access point, therefore controller1 selects the controller with the greatest capacity to add new active access point connections. In this example, controller1 has more capacity. Controller1 sends its own IP address to the access point in the find controller reply message.</p>	orange
<p>Access point contacts controller1 and determines whether to use a locally stored operational image or download an image from the controller.</p>	<p>Controller1 sends an access point image if requested.</p>	orange

Table 193: Access Point Connecting to a Controller in a Layer 3 Network (*continued*)

Access Point Action	Response	Color of Route in Figure
Once the operational image is loaded, the access point requests configuration information from controller1.	Controller1 sends SSID information to the access point.	orange

RELATED DOCUMENTATION

Mobility System Software Quick Start Guide, *Preparing the Network for Distributed APs*

Mobility System Software Quick Start Guide, *Using the Quickstart Command.*

[Adding and Managing an Individual Access Point | 1156](#)

[Understanding Access Point Bias for Controllers | 852](#)

[Network Director Documentation home page](#)

Understanding How To Add Access Points to a Wireless Network By Using Network Director

There are two ways to add an access point to your wireless network. Both methods have advantages and disadvantages as shown in [Table 194](#).

Table 194: Adding Access Points to a Wireless Network With Network Director

Method	What You Do	Pros and Cons
The controller discovers access points and provides the list to Network Director. You can convert any or all of these temporary configurations to persistent configurations.	Enable automatic access point discovery on a controller by following the directions “Creating and Managing Wireless Auto AP Profiles” on page 980 . Convert the automatically discovered access points to persistent configurations by following the directions “Converting Automatically Discovered Access Points to Manually Configured Access Points” on page 1248 . (For an explanation of Auto AP profiles, see “Understanding Auto AP Profiles” on page 883 .)	All access points in range are automatically discovered and listed first as access points on the controller and then as devices in Network Director. This method is very efficient. The drawback to this method is that non-company access points can be discovered and added if they are in range. You can overcome this drawback by converting the access points that belong to your network.
You add an access point one at a time.	See “Adding and Managing an Individual Access Point” on page 1156 for directions for this method of adding access points.	Adding one access point at a time can be tedious if you need to add a large number of access points. However, adding one access point at a time from Network Director is more explicit because that method inputs more information and does more configuration. Controllers give preference to specifically configured, persistent access points over those access points using an Auto AP profile. If a configured access point is discovered by a fully utilized controller, the controller disconnects an access point using an Auto AP profile and accepts connection from the configured access point instead.

RELATED DOCUMENTATION

[Adding and Managing an Individual Access Point | 1156](#)

[Creating and Managing Wireless Auto AP Profiles | 980](#)

[Converting Automatically Discovered Access Points to Manually Configured Access Points | 1248.](#)

[Understanding Auto AP Profiles | 883](#)

[Network Director Documentation home page](#)

Understanding Radio Profiles

IN THIS SECTION

- [RF Scanning | 879](#)
- [Spectral Scanning | 880](#)
- [Dynamic Frequency Selection \(DFS\) Channels | 880](#)
- [RFID Asset Tracking | 880](#)
- [WMM Power Save | 881](#)
- [Countermeasures | 881](#)
- [Limiting Client Power | 882](#)
- [802.11n Channel Width | 882](#)
- [Automatic Channel Tuning | 882](#)
- [Automatic Power Tuning | 882](#)
- [IEEE 802.11 | 882](#)
- [Long and Short Preamble Length | 882](#)

Radio profiles contain configuration information for radios. A single Radio profile can be applied to many radios that share a common configuration. This topic describes Radio profile configuration concepts. To create, assign, or delete a Radio profile, see [“Creating and Managing a Radio Profile” on page 932.](#)

This topic describes:

RF Scanning

All wireless radios continually scan for other RF transmitters. While 802.11b/g radios scan in the 2.4-GHz to 2.4835-GHz spectrum, 802.11a radios scan in the 5.15-GHz to 5.85-GHz spectrum. There are two scanning methods, passive scanning and active scanning. By default, radios perform both types of scans on all channels allowed by the country of operation. (The country of operation is the regulatory domain set during initial access point deployment.)

While both types of scanning are on by default, active scanning is done only on channels on which local government regulations allow it to transmit. Channels that are not authorized for unlicensed use and channels that require radar detection with dynamic frequency selection (DFS) are excluded from active scanning.

For more information, about scanning see [“Understanding Wireless Scanning” on page 869](#).

Spectral Scanning

The electromagnetic spectrum includes all possible frequencies of electromagnetic radiation. Wireless communication uses unlicensed radio bands where different types of devices (microwave ovens, cordless phones, video surveillance cameras for example) might emit radio signals that interfere with Wi-Fi service. Spectral analysis detects and reports on radio transmissions in the frequency bands used by Wi-Fi equipment. For more information, about scanning see [“Understanding Wireless Scanning” on page 869](#) and [“Monitoring the RF Spectrum of a Radio” on page 1343](#).

Dynamic Frequency Selection (DFS) Channels

The regulatory bodies of most countries now allow many 5-GHz (802.11a) channels that are assigned for use by radar systems to be used for by 802.11 wireless networking, provided there is no radar operating in the vicinity of the wireless network deployment. To use these channels, the wireless access point must implement radar detection and immediately vacate any channel on which radar is active. This requirement is referred to as Dynamic Frequency Selection (DFS). To accomplish this, a DFS mechanism was created to have unlicensed devices detect the presence of a radar system on the device channel and, if the level of the radar is above a certain threshold, vacate that channel and select an alternate channel.

When you deploy an access point, you must configure the correct country of operation to ensure compliance with local government regulations. Failing to do so might cause harmful interference to other systems. In certain countries, such as the United States, access point models are country-specific and will not operate if you configure a different country of operation. Since the corresponding DFS requirements for each available country of operation are programmed into the access point, you do not have to specifically configure the DFS channels. All you need to do is enable DFS in a Radio profile and the appropriate channel avoidance will be implemented.

RFID Asset Tracking

RFID stands for Radio-Frequency Identification. RFID tracking devices consist of transponders, capable of carrying as much as 2,000 bytes of data, with an attached antenna. These devices provide a unique identifier for the object they are attached to.

There are various kinds of RFID, such as the card readers on the doors at companies—the Juniper Networks wireless system does not support those kinds of RFIDs. Juniper Networks access points are capable of reading 802.11 RFID tags that *chirp* short 802.11 messages periodically, for example AeroScout tags. Because the RFID tag must be scanned by an access point radio to retrieve the identifying information,

you must configure RFID scanning in a Radio profile. Then the information can be scanned by access points in that Radio profile and fed to a transceiver to interpret the information.

RFID tags replace traditional barcode solutions for inventory and tracking, and provide the following advantages:

- You can track items without climbing ladders, crawling under furniture, or having to be in the direct line of sight of tracked items.
- Most things can be tracked, including people.
- The time it takes to conduct a complete inventory is significantly reduced, while the accuracy of the data is greatly increased.
- The time it takes to locate an asset, for example an expensive medical device, is greatly reduced.

WMM Power Save

The receiver portion of an 802.11 radio consumes considerable battery power if it is left on all the time. WMM Power Save enables mobile devices to minimize the length of time a receiver is fully powered on, thereby extending extend battery life. With WMM Power Save, the same amount of data can be transmitted in fewer frames in a shorter time, while allowing the Wi-Fi device to preserve power in a low-power, dozing state in between transmissions. To take advantage of WMM Power Save functionality, both the Wi-Fi client and access point must use WMM Power Save. In addition, the applications used also need to support WMM Power Save to inform the client of the requirements of the traffic they generate. WMM is a required feature for 802.11n capable devices—almost all modern Wi-Fi devices support it.

For more information about WMM Power Save, see [“Understanding WMM Power Save and WLAN Client Battery Life” on page 859](#).

Countermeasures

Countermeasures are actions used by controllers and access points to thwart rogue devices attempting to use your network. You can also apply countermeasures to suspect devices. Countermeasures are enabled in a Radio profile and take place automatically when an access point discovers a device.

NOTE: Countermeasures cannot be configured with Network Director Release 1.0. Use the MSS CLI to configure countermeasures.

Limiting Client Power

In Radio profiles, you can limit the maximum power level allowed on clients that associate to your access points. Then, when a client associates with an access point, the access point transmits the maximum power level setting allowed in that country to the client. Clients automatically configure their power to match the access point requirements.

802.11n Channel Width

802.11n channels can be either 40 MHz or 20 MHz wide if you use the 5-GHz radio band. See [“Understanding Wireless Radio Channels” on page 856](#) for more information and for instructions to indicate a 40 MHz channel width.

Automatic Channel Tuning

Channel auto-tuning automatically makes channel tuning decisions for access point radios based on the RF data gathered by access points. Channel auto-tuning is configured in a Radio profile and constantly operates on all access point radios by default. For more information about automatic channel tuning, see [“Understanding Adaptive Channel Planner” on page 861](#).

Automatic Power Tuning

The amount of power an access point uses affects the coverage area of the access point. The higher the power level, the larger the coverage area of an access point. Usually, you want your access points to cover all areas with minimal overlap. With automatic power tuning, which is controlled by the Radio profile, access points adjust the power levels of radios automatically, based on the power levels of all neighboring access points. For more information about automatic power tuning, see [“Understanding Auto Tune Power Policy for Wireless Radios” on page 866](#).

IEEE 802.11

IEEE 802.11 is a set of standards for implementing wireless local area network (WLAN) computer communication in the 2.4-GHz and 5-GHz frequency bands. For more information about IEEE 802.11, see [“Understanding the IEEE 802.11 Standard for Wireless Networks” on page 1076](#).

Long and Short Preamble Length

The preamble is part of the IEEE 802.11b physical layer specification. Specifically, the preamble is a data header section that contains information the access point and clients need when both sending and receiving packets.

In general, use the short preamble type in high network traffic areas, and use the long preamble to provide more reliable communication in noisy (high interference) networks.

Note that all 802.11b devices must support the long preamble format, but can also optionally support the short preamble. Because the short preamble is default, if an 802.11b device has trouble communicating with other 802.11b devices, try using the long preamble.

RELATED DOCUMENTATION

[Understanding the Network Director User Interface | 85](#)

[Creating and Managing a Radio Profile | 932](#)

[Understanding Wireless Scanning | 869](#)

[Understanding the IEEE 802.11 Standard for Wireless Networks | 1076](#)

[Understanding Call Admission Control | 1068](#)

[Understanding Auto Tune Power Policy for Wireless Radios | 866](#)

[Understanding WMM Power Save and WLAN Client Battery Life | 859](#)

[Understanding Wireless Radio Channels | 856](#)

[Network Director Documentation home page](#)

Understanding Auto AP Profiles

IN THIS SECTION

- [How Are Specifically Configured Access Points Different from Access Points Configured with Auto AP? | 884](#)
- [How Should I Use Auto AP Profiles? | 884](#)
- [How Does Auto AP Work? | 885](#)

If you set up an Auto AP profile for a wireless controller, discovered distributed access points that are not specifically configured on that controller use a temporary configuration from the Auto AP profile. When you set up an Auto AP profile for these distributed access points, you provide only the basic information that an access point needs to operate. Auto AP access points do not include all the information that individually configured access points do. In addition, these Auto AP access point configurations are not persistent on the controller—they go away when the access points are not present.

NOTE: A controller can have only one Auto AP profile.

This topic describes:

How Are Specifically Configured Access Points Different from Access Points Configured with Auto AP?

Controllers give preference to specifically configured, persistent access points over those access points using an Auto AP profile. If a configured access point is discovered by a fully utilized controller, the controller disconnects an access point using an Auto AP profile and accepts connection from the configured access point instead.

In this case, the disconnected access point begins the boot process again to find another controller with an Auto AP profile. When the access point is disconnected, the access point clients experience a service disruption, and attempt to associate with another available access point to reconnect to the SSID. If another access point is not available to a client, the client can still reconnect after the disconnected access point locates a new controller and finishes the boot and configuration process.

How Should I Use Auto AP Profiles?

The Auto AP feature is useful when you need to add multiple access points all at once. Typical use of the Auto AP profile would be to:

- Configure an Auto AP profile, assign the profile to a controller, and then deploy the controller configuration. To configure an Auto AP profile, see [“Creating and Managing Wireless Auto AP Profiles” on page 980](#). To assign an Auto AP profile to a controller, see [“Assigning an Auto AP Profile to Controllers” on page 991](#). To deploy the controller configuration, see [“Deploying Configuration to Devices” on page 1180](#).
- Convert access points found by Auto AP on that controller to persistently configured access points following the directions in [“Converting Automatically Discovered Access Points to Manually Configured Access Points” on page 1248](#).
- Turn off the Auto AP profile by editing the profile—see [“Creating and Managing Wireless Auto AP Profiles” on page 980](#).

How Does Auto AP Work?

When an Auto AP profile is assigned to a controller, the controller assigns a valid number and a name to unassigned, distributed access points. The controller also configures the access point settings and radio parameter settings supplied in the Auto AP profile.

NOTE: The Auto AP profile does not configure SSIDs, encryption parameters, or any other parameters managed by the WLAN Service profile. You still need to configure a WLAN Service profile for each SSID.

RELATED DOCUMENTATION

[Creating and Managing Wireless Auto AP Profiles | 980](#)

[Assigning an Auto AP Profile to Controllers | 991](#)

[Deploying Configuration to Devices | 1180](#)

[Converting Automatically Discovered Access Points to Manually Configured Access Points | 1248](#)

[Understanding How To Add Access Points to a Wireless Network By Using Network Director | 878](#)

[Network Director Documentation home page](#)

Understanding WLAN Service Profiles

IN THIS SECTION

- [SSID | 886](#)
- [Mapping WLAN Service Profiles to Additional Profiles | 887](#)
- [SSID Encryption | 887](#)
- [Associated Authentication Profile | 888](#)
- [Associated Authorization Profile | 888](#)
- [VLAN Use | 888](#)
- [Bandwidth Limit for Client Sessions | 888](#)
- [Load Balancing Between Access Points | 888](#)
- [Using Proxy ARP | 889](#)
- [Restricting DHCP | 889](#)

- Client Types | 889
- Call Admission Control Settings for Voice | 890
- Retry Count | 890
- Client Timeouts | 890
- 802.11n Settings | 890
- Maximum Bandwidth Used by a WLAN Service Profile's SSID | 891
- Maximum Transmission Unit Parameter | 891
- Client Probing of Idle Clients | 892
- Enable Pre-Shared Key (PSK) for WPA or WPA2 | 892
- Create a Pre-Shared Key (PSK) Phrase for WPA or WPA2 | 892
- Create a Pre-Shared Key (PSK) Raw Phrase for WPA or WPA2 | 892
- Enforce Data Rates | 892
- Retry Count for Sending Frames | 892
- WPA Encryption Type Used | 893
- Shared Key Authentication Values | 893
- Radio Transmit Rates Used | 893
- WMM Power Save | 893

A wireless LAN (WLAN) Service profile is a set of configurations, including a unique SSID, that provides clients part of a wireless connection to the wireless network. You must have at least one WLAN Service profile with an SSID on your wireless network for operation. Note that there are no default profiles provided by MSS—you must configure each WLAN Service profile. This topic describes the parameters configured in a WLAN Service profile.

There are many parameters, either optional or mandatory, that are associated with every SSID. The WLAN Service profile provides some of these parameters but not all of them.

This topic describes the parameters configured in a WLAN Service profile:

SSID

The SSID name is the most important configuration in a WLAN Service profile. Only the SSID name and WLAN name are actually required to create a WLAN Service profile—the other parameters have default values.

Wireless networks are identified by unique network names such as Juniper Networks_meetings or employee_patio. The unique name is known as a service set identifier, or SSID and this electronic identifier

serves as a password for certain online communications. Most controllers are set to broadcast their SSID using WPA or WPA2 encryption. If an SSID is not broadcast, you must manually configure one or more SSIDs on clients so that the clients automatically find and connects to those SSIDs when it is in range of the controller. On a PC running Microsoft Windows OS, this setting is typically found in the Network Control Panel.

Beaconing the SSID Name

By default, SSIDs are beaconed, which means that an SSID advertises its name on the air. You might want to disable beaconing for security reasons, although doing so can make it more difficult for clients to access the WLAN. When you disable beaconing for an SSID, the radio still sends beacon frames, but the SSID name in the frames is blank. For a non-beaconed SSID, radios respond only to directed 802.11 probe requests that match the non-beaconed SSID string.

NOTE: Disabling beaconing is not particularly effective as a security measure because any sniffer can easily detect the name of the SSID.

Mapping WLAN Service Profiles to Additional Profiles

WLAN Service Profiles rely on other profiles, such as Authentication profiles and Authorization profiles, to provide additional parameters for the SSID. You link these other profiles as part of the WLAN configuration. (For more information, see [“Understanding Network Director SSID Configuration Using Profiles” on page 1064.](#)) The completed WLAN Service profile is then mapped to a Radio profile during configuration of the Radio profile. It is the Radio profile that is actually deployed to controllers, pulling all of the other mapped profiles along with it.

SSID Encryption

Encryption protects information within a wireless session by reading that information in a data stream and altering it to make it unreadable to users outside the network.

Encryption (ssid-type) is either on (Crypto) or off (Clear). If encryption is on, additional configuration indicates the type of encryption (WPA and/or WPA2), any added ciphers such as TKIP (tkip-mc-time) or CCMP. For more information, see [“Understanding Wireless Encryption and Ciphers” on page 899.](#)

Both WPA and WPA2 are enabled by default in Network Director. Clients that use only WPA associate by using WPA and all clients capable of using WPA2 associate by using WPA2.

NOTE: The original wireless encryption method, WEP, is not supported in Network Director.

Authentication Used for Encryption Methods

The standard for wireless LAN authentication is the IEEE 802.1X standard, which is based on the IETF's Extensible Authentication Protocol (EAP). EAP and 802.1X together provide an authentication framework.

The second way to authenticate encrypted traffic is pre-selected key (PSK) authentication. If you use PSK, you must also provide the key or password. For more information, see [“Understanding Wireless Encryption and Ciphers” on page 899](#).

Associated Authentication Profile

An Authentication profile must be associated with each WLAN Service profile—you do this by selecting an Authentication profile while configuring the WLAN. Authentication profiles are described in [“Understanding Authentication Profiles” on page 381](#).

Associated Authorization Profile

An Authorization profile is also mapped to WLAN Service Profiles—you do this by selecting an Authorization profile while configuring the WLAN. Authorization profiles are described in [“Understanding Wireless Authorization Profiles” on page 395](#).

VLAN Use

Clients usually switch VLANs when they switch controllers, usually as a result of roaming, but you can make initially assigned VLANs persist over different controllers. If an 802.1X user is not assigned to a VLAN by AAA, and subsequently roams to a controller where the VLAN he was in does not exist, a tunnel is set up so that he stays in that VLAN. This does not work for Web portal clients, however.

Bandwidth Limit for Client Sessions

You can limit bandwidth for clients to prevent one client from hogging bandwidth.

Load Balancing Between Access Points

RF load-balancing is the ability to reduce network congestion over an area by distributing client sessions across the access point radios with overlapping coverage in the area. Load balancing automatically occurs on the mobility domain to ensure maximum failover capability. For more information, see [“Understanding Load Balancing for Wireless Radios” on page 1070](#).

Using Proxy ARP

Proxy address resolution protocol (ARP) is a technique by which a device on a network answers requests for a different device. Because the proxy knows the location of the traffic's destination, it offers its own IP address then sends the traffic on to the true destination.

Usually, wireless clients receive an IP address from a router. If you want to, you can have the controller respond to wireless clients' search for a destination, and then forward the traffic to the router.

Restricting DHCP

The dynamic host configuration protocol (DHCP) is used to configure network devices with IP addresses from a DHCP server.

You can configure a controller to capture but not forward any wireless client traffic except DHCP traffic during authentication and authorization. This is referred to as restricting DHCP and enables a controller to authenticate and authorize new clients more quickly.

Client Types

You decide which client types to support on a WLAN. If you put all clients on one WLAN, they will be reduced to the speed of the slowest client—we do not recommend doing this.

Possible client types are:

- 802.11n clients use newer technology that can produce throughput link rates above 54 MBps, if you select only this client type and enforce the data rate (speed). Typical clients include laptops, PCs, and streaming video.

NOTE: The Wi-Fi Alliance requires that high-throughput (802.11n) transmissions use WPA2 and CCMP.

- 802.11g clients can have throughput link rates as fast as 54 MBps, with a more average rate of 19 MBps. Typical clients include older laptops and PCs.
- 802.11a clients can have throughput link rates as fast as 54 MBps, with a more average rate of 19 MBps.
- 802.11b clients can have throughput link rates as fast as 10 MBps.

Call Admission Control Settings for Voice

Call admission control (CAC) regulates the addition of new real-time media sessions on access point radios, guaranteeing a higher quality of service to a fixed number of clients by limiting either the number of concurrent sessions or the number of concurrent phone calls.

There are two CAC methods, WMM and SVP. WMM is enabled by default and is used by all newer 802.11n devices. SVP includes all the configurations of WMM, but adds Spectralink phones at the top of the priority list. SVP is a legacy technology—Spectralink's new phones use WMM.

As the name indicates, call admission control applies to real-time media traffic as opposed to data traffic. Call admission control mechanisms and quality-of-service settings work together to protect voice traffic from the negative effects of other voice traffic and keep excess voice traffic off the network. For more information, see [“Understanding Call Admission Control” on page 1068](#).

Retry Count

The retry count is the number of times a channel resends a frame without getting a response. You can configure 1 - 15 attempts, with each attempt taking more time and affecting throughput. We recommend configuring five attempts.

You can also specify either a long retry count or a short retry count. The difference between the two methods is the length of the pause between attempts. In general, a short retransmission works best in heavy traffic and is used most often. You might want to try a long transmission if the network is experiencing a lot of interference.

Client Timeouts

Timeouts are used to disconnect clients under certain circumstances. You can change the timeframes for these timeouts:

- User Idle—180 seconds
- Handshake attempt (logon)—20 milliseconds
- Web portal session—0 seconds (which means no timeout)

802.11n Settings

802.11n is the most recent wireless technology that utilizes different mechanisms than previous versions of 802.11. Therefore, there are settings that apply only to 802.11n traffic.

Guard Intervals

A guard interval is the interval is observed before the next bit of traffic is transmitted. This guard interval ensures that bit transmissions do not interfere with one another. As long as the echoes fall within this

interval, they do not affect the receiver's ability to safely decode the actual data, because data is interpreted only outside the guard interval—it eliminates intersymbol interference. In normal 802.11 operation, the guard interval is 800 ns. In 802.11n operation, short guard intervals of 400 ns are supported. Shorter guard intervals between symbols increases throughput. Legacy devices might require long guard intervals. By reducing this interval (called *short guard interval*), data bits are transmitted in shorter intervals and provide for increased throughput.

Frame Aggregation

You can enable frame aggregation for certain frame types in 802.11n. Multiple packets of application data can be aggregated into a single packet called an aggregated MAC protocol data Unit (A-MPDU). This improves performance because the number of packets is reduced.

After transmission of every frame, an idle time called Interframe Spacing (IFS) is observed before transmitting the subsequent frame. When frames are aggregated, fewer IFS intervals are used, which in turn reduces the time for data transmission. In addition, when clients operating in 802.11n send acknowledgement for block of aggregated packets instead of individual packets, overhead involved in frame acknowledgements and increasing overall throughput is reduced.

MAC Service Data Unit (MSDU) Length

With 802.11n, you can change the maximum length for a MAC service data unit (MSDU) to reduce the overhead associated with each transmission. An MSDU is the service data unit received from the logical link control (LLC) sub-layer which lies above the medium access control (MAC) sub-layer in a protocol stack. When 802.11n is an enabled client type for this WLAN, you can configure the maximum aggregated MSDU packet length. This enables joining multiple packets together into a single transmission unit, which reduces the overhead associated with each transmission. MSDU default length is 4K.

MPDU Length

MPDU can have a maximum length for frame aggregation.

Maximum Bandwidth Used by a WLAN Service Profile's SSID

The speed of a computer network is most commonly stated in bandwidth units of Megabits per second (Mbps) or Gigabits (Gbps). This standard measure of communication capacity (data rate) is advertised by all computer networking equipment. When you indicate a maximum bandwidth for a WLAN, you are indicating the highest data rate you support with a given WLAN Service profile. Devices in each category have a maximum bandwidth, so the maximum bandwidth also determines which devices are supported. For example, the 802.11g standard for wireless networking supports a maximum bandwidth of 54 Mbps, including overhead. 802.11n supports a maximum bandwidth of 600 Mbps.

Maximum Transmission Unit Parameter

The maximum transmission unit (MTU) of the communications protocol of a layer is the size (in bytes) of the largest protocol data unit that the layer can forward. MTU parameters usually appear in association with a communications interface such as a NIC card or serial port.

Client Probing of Idle Clients

Idle client probing sends periodic keepalives from a radio to non-transmitting clients. By default, a radio sends idle-client probes every 10 seconds to each client with a session on the radio to verify that the client is still active. The probes are unicast null-data frames. Normally, an active client sends an ACK in reply to an idle-client probe. If a client does not send any data or respond to idle-client probes before the user idle timeout expires, the client session is disassociated.

Enable Pre-Shared Key (PSK) for WPA or WPA2

WPA-PSK is an authentication mechanism in which users provide credentials to verify whether or not to allow them access to a network. This requires that a single password be entered into each WLAN node (access points, wireless routers, client adapters, bridges). When the passwords match, a client is granted access to a WLAN.

Create a Pre-Shared Key (PSK) Phrase for WPA or WPA2

A pre-shared key (PSK) phrase is the hexadecimal secret phrase used for authenticating WPA or WPA2 clients. Note that either WPA or WPA2 security must be enabled for this to have any effect.

Create a Pre-Shared Key (PSK) Raw Phrase for WPA or WPA2

A pre-shared key (PSK) raw phrase is the raw hexadecimal secret phrase used for authenticating WPA or WPA2 clients. Note that either WPA or WPA2 security must be enabled for this to have any effect.

Enforce Data Rates

By default, a client can associate with and transmit data to an access point by using a slower data rate than the mandatory or standard rate, although the access point does not necessarily transmit data back to the client at the slower rate.

When you enforce data rates, a connecting client must transmit at one of the mandatory or standard rates to associate with the access point. Clients transmitting at slower rates cannot associate with the access point.

Retry Count for Sending Frames

Retry-count settings indicate how many times that the network sends either a long unicast frame or short unicast frame without receiving an acknowledgement. If you set retry count to zero, frames are sent once with no retries.

NOTE: The fragmentation threshold uses the short-retry-count for frames shorter than 2346 bytes and uses the long-retry-count for frames that are 2346 bytes or longer.

WPA Encryption Type Used

Enable and choose WPA encryption—either WPA or WPA2. Either or both TKIP and CCMP cipher algorithms for WPA and WPA2 can be added.

NOTE: The Wi-Fi Alliance requires that high-throughput (802.11n) transmissions use WPA2 and CCMP.

Shared Key Authentication Values

Shared key authentication is a process by which a client gains access to a WLAN by using an encryption key. The key, obtained in advance by the client, must match a key stored at the access point. To begin the connection process, the client sends a request for authentication to the access point. The access point responds by generating a sequence of characters called a challenge text for the computer. The computer encrypts the challenge text with the key and transmits the message back to the access point. The access point decrypts the message and compares the result with the original challenge text. If there are no discrepancies, the access point sends an authentication code to the connecting computer. Finally, the computer accepts the authentication code and becomes part of the network for the duration of the session or for as long as it remains within range of the original access point. If the decrypted message does not precisely agree with the original text, the access point does not allow the computer to become part of the network.

Radio Transmit Rates Used

Radio transmit rates supported by access point radios have defaults, but you can change the transmit rates for the radios. Each type of radio (802.11a, 802.11b, 802.11g, and 802.11n) providing service to an SSID has a set of rates the radio is enabled to use for sending beacons, multicast frames, and unicast data. The rate you set also specifies the rates clients must support to associate with a radio.

WMM Power Save

WMM Power Save is disabled by default, even though it saves client battery life, because clients that use power save must send a separate PSpoll to retrieve each unicast packet buffered by the access point radio.

This increases bandwidth and affects performance. For more information, see [“Understanding WMM Power Save and WLAN Client Battery Life”](#) on page 859.

RELATED DOCUMENTATION

[Creating and Managing a WLAN Service Profile | 1090](#)

[Understanding WMM Power Save and WLAN Client Battery Life | 859](#)

[Understanding Wireless Authorization Profiles | 395](#)

[Understanding Wireless Encryption and Ciphers | 899](#)

[Understanding the Network Director User Interface | 85](#)

[Understanding Network Director SSID Configuration Using Profiles | 1064](#)

[Network Director Documentation home page](#)

Understanding Wireless Mesh

IN THIS SECTION

- [Example Mesh Topology With One Access Point Wired | 895](#)
- [Why Use Mesh? | 896](#)
- [How Does Mesh Work? | 896](#)
- [Planning a Mesh Portal | 897](#)
- [How Do I Set Up and Configure Mesh? | 898](#)
- [How Do I Configure a Mesh Access Point? | 898](#)
- [Security Between a Mesh AP and the Mesh Portal AP | 898](#)

WLAN mesh allows an access point to communicate with the network using a radio link to route network traffic toward its destination. This is an alternative to using wired links for each access point. Mesh networking is useful when you need to provide wireless coverage to an area where wired network connection is not practical. Instead, the remote access point uses a wireless link to another access point to provide access to the rest of the network. See [Figure 37](#) where the mesh has one access point wired to a switch and two access points that are not hard-wired.

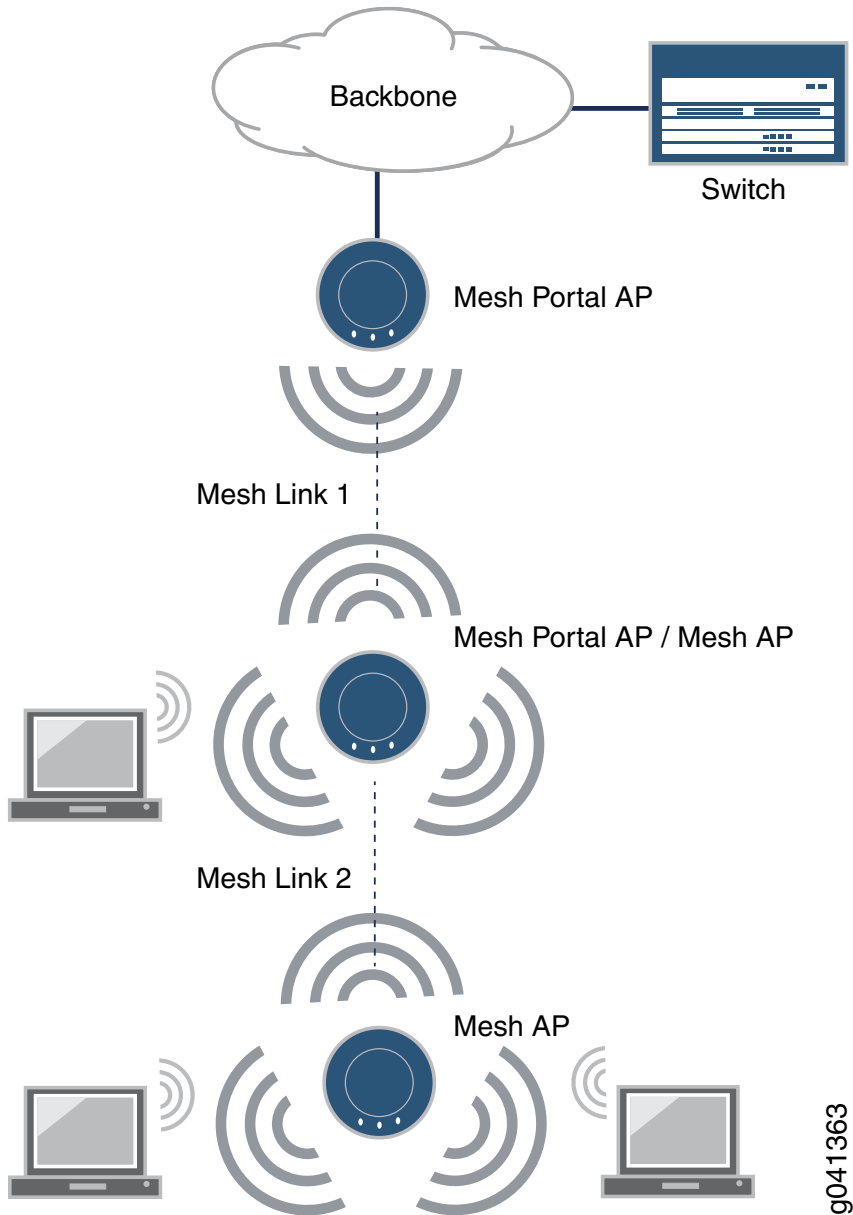
NOTE: WLAN mesh is supported only for access points with dual radios.

This topic describes:

Example Mesh Topology With One Access Point Wired

Many mesh options and topologies are available. In this example, the mesh has one wired access point and all other access point connections are wireless as shown in [Figure 37](#).

Figure 37: Example Mesh Topology



g041363

Why Use Mesh?

WLAN mesh is useful when running an Ethernet cable to a location is either inconvenient, expensive, or impossible.

How Does Mesh Work?

As shown in [Figure 37](#), there are three components in a mesh deployment:

- *Mesh Portal*—any access point providing a wireless signal to another access point. There will always be one mesh portal connected to a wired network. Additional access points with two radios can perform the optional role of both mesh access point and mesh portal access point.

NOTE: An access point must have two radios if it performs the dual role of both mesh access point and mesh portal access point. The wired access point gets power from the switch—additional portal access points need their own power. (The access points can be using PoE without a network connection.)

- *Mesh AP*—a wireless access point without a wired connection (untethered). Clients associated with a mesh access point have the same connectivity to the network they would have when associated to a wired access point.

NOTE: An access point must have two radios if it performs the dual role of both mesh access point and mesh portal access point.

- *Mesh Link*—a Layer 2 transparent bridge with a mesh portal and a mesh access point as endpoints.

Once mesh is configured, the mesh portal access points beacon a mesh services SSID on the radio used for the mesh link, which is also commonly called the backhaul. When a mesh access point is booted, it finds access points beaconing a mesh SSID, selects the mesh portal access point with the greatest signal strength (RSSI value), and then establishes a secure connection to the mesh portal SSID. Once this connection is established, the access point can offer services configured by the controller. If, after 60 seconds, no link is established, the remote mesh access point reboots. If the remote access point fails to connect to the mesh access point it has chosen, it tries another, and so on until it has tried all mesh access points. If all attempts to connect fail, the remote access point tries each access point again and repeats until it connects or reboots after 60 seconds.

Planning a Mesh Portal

The following recommendations provide the most stable mesh services on a wireless network:

- Dedicate one radio to client services and one radio to mesh services. We recommend that you dedicate the 802.11a radio (radio 2) to mesh services and the 802.11g radio (radio 1) to client services.
- Dedicate the entire mesh portal access point to mesh services if you anticipate a need for the additional bandwidth.

- Limit the physical length of the mesh link to 3/8ths of a mile (1.09 km) or less if you have configured MSS 6.0.4 or earlier. Later versions of MSS can support distances up to 1 mile (1.6 km) for Mesh Links, depending on the RF characteristics of the access points and their antennae, location and configuration.
- For the best performance, minimize the number of mesh access points that connect to a mesh portal and avoid multi-hop mesh deployments. Never exceed a mesh width of 10 mesh portal access points or a mesh depth of four mesh access points per mesh portal access point.

How Do I Set Up and Configure Mesh?

Mesh requires a dedicated mesh SSID (named in a WLAN Service profile), a dedicated Radio profile, and access points configured for untethered mesh operation. You configure access points for mesh while they are connected to a controller, then untether them and place them into the mesh location.

If you offer client service from the mesh portal, you must use dual-radio access points so that one radio can be used for mesh link communications (using the SSID reserved for this purpose) while the other radio is used for client associations.

How Do I Configure a Mesh Access Point?

Before a Mesh access point can be installed in a location untethered from the network, it must be connected to a controller and preconfigured for mesh services, including the mesh services SSID, and the pre-shared key for establishing the connection between the mesh access point and the mesh portal. For directions, see [“Configuring Wireless Mesh and Bridging” on page 976](#).

NOTE: When using external antennas in conjunction with mesh configurations, enable mesh mode before configuring the external antenna. After adding and configuring the external antenna, reboot the access point.

Mesh access points must be configured for local switching—see [“Understanding Local Switching on Access Points” on page 907](#).

Security Between a Mesh AP and the Mesh Portal AP

Security is configured as part of the mesh WLAN Service profile, where the required pre-shared key (PSK) authentication and PSK key are set up in the mesh SSID. The PSK key must match the one you configure on the mesh access points.

RELATED DOCUMENTATION

[Configuring Wireless Mesh and Bridging | 976](#)

[Understanding Wireless Bridging | 912](#)

[Understanding Local Switching on Access Points | 907](#)

[Network Director Documentation home page](#)

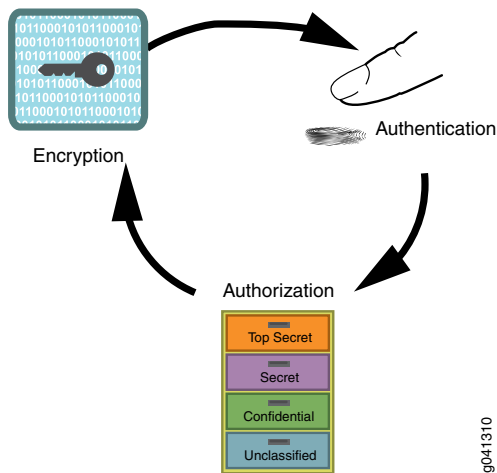
Understanding Wireless Encryption and Ciphers

IN THIS SECTION

- [Wired Equivalent Privacy \(WEP\) was the Original Wireless Encryption | 900](#)
- [WPA Encryption Replaced WEP | 900](#)
- [WPA2 Is the Strongest Encryption Available | 901](#)
- [Security Ciphers for WPA and WPA2 | 901](#)
- [Which Encryption Method Should I Use? | 902](#)

Wireless network security relies on a combination of encryption, authentication, and authorization to provide maximum protection for a WLAN. Encryption is focused on protecting the information within a session, reading information in a data stream and altering it to make it unreadable to users outside the network. This topic discusses encryption.

Figure 38: Network Security Is Provided by Encryption, Authentication, and Authorization



Juniper Networks access points support all three standard types of wireless access point-client encryption: the legacy encryption Wired Equivalent Privacy (WEP), Wi-Fi Protected Access (WPA), and WPA2 (also called RSN). Encryption type is configured in WLAN Service profiles under the Security Settings tab. For information about applying encryption, see [“Creating and Managing a WLAN Service Profile” on page 1090](#).

This topic describes:

Wired Equivalent Privacy (WEP) was the Original Wireless Encryption

WEP was the original security algorithm for IEEE 802.11 wireless networks, introduced as part of the original 802.11 standard.

WPA Encryption Replaced WEP

WPA addressed the vulnerabilities of WEP, the original, less secure 40 or 104-bit encryption scheme in the IEEE 802.11 standard. WPA also provides user authentication—WEP lacks any means of authentication.

WPA replaced WEP with a stronger encryption technology called Temporal Key Integrity Protocol (TKIP) with Message Integrity Check (MIC). It also provides a scheme of mutual authentication using either IEEE 802.1X/Extensible Authentication Protocol (EAP) authentication or pre-shared key (PSK) technology.

NOTE: You can simultaneously apply both WPA and WPA2 to an SSID. Clients use WPA2 if they have the capability—otherwise the client uses WPA. WPA2 is recommended unless you need to provide access to for legacy devices. All 802.11n devices support WPA2.

WPA2 Is the Strongest Encryption Available

WPA2 is the certified version of the full IEEE 802.11i specification. Like WPA, WPA2 supports either IEEE 802.1X/EAP authentication or PSK technology. It also includes a new advanced encryption mechanism using the Counter-Mode/CBC-MAC Protocol (CCMP) called the Advanced Encryption Standard (AES).

WPA was based on the 802.11i draft, while WPA2 is based on 802.11i final standard. Where WPA encryption was specifically designed to work with some wireless hardware that supported WEP, WPA2 offers stronger security but is not supported by earlier hardware designed for WEP.

NOTE: The Wi-Fi Alliance requires that high-throughput (802.11n) transmissions use WPA2 and CCMP. You can simultaneously apply both WPA and WPA2 to an SSID. Clients use WPA2 if they have the capability—otherwise the client uses WPA.

Security Ciphers for WPA and WPA2

Standard security ciphers are part of both WPA and WPA2 encryption. You choose whether you want to apply either the newer CCMP, or TKIP (an upgrade of original WEP programming), or both for each WLAN Service profile. Both cipher suites dynamically generate unique session keys for each session and periodically change the keys to reduce the likelihood of a network intruder intercepting enough frames to decode a key. The two available ciphers are:

- *Counter Mode with Cipher Block Chaining Message Authentication Code Protocol (CCMP)*—CCMP provides Advanced Encryption Standard (AES) data encryption for WPA and WPA2. To provide message integrity, CCMP also uses the Cipher Block Chaining Message Authentication Code (CBC-MAC).

A radio using WPA/WPA2 with CCMP encrypts traffic for only WPA CCMP clients but not for TKIP clients. The radio disassociates from TKIP clients unless you selected both CCMP and TKIP.

NOTE: The Wi-Fi Alliance requires that high-throughput (802.11n) transmissions use WPA2 and CCMP.

- *Temporal Key Integrity Protocol (TKIP)*—TKIP uses the RC4 encryption algorithm, a 128-bit encryption key, a 48-bit initialization vector (IV), and a message integrity code (MIC). A radio using WPA/WPA2 with TKIP encrypts traffic for only WPA TKIP clients but not for CCMP clients. The radio disassociates from CCMP clients unless you selected both CCMP and TKIP.

TKIP is most useful for upgrading security on devices originally using WEP — it does not address all of the security issues facing WLANs and may not be reliable or efficient enough for sensitive corporate and government data transmission. The 802.11i standard specifies the Advanced Encryption Standard (AES) in addition to TKIP. AES is an additional cipher stream that adds a higher level of security and is approved for government use.

NOTE: TKIP is not permitted for 802.11n-based transmissions. It is only supported for legacy (802.11b, 802.11g and 802.11a) transmissions, which are limited to a maximum of 54 Mbps.

Which Encryption Method Should I Use?

WPA2 is the most secure encryption method available for wireless networks—we recommend using WPA2 with the CCMP cipher whenever possible. WPA2 with CCMP is the only option permitted for high throughput 802.11n transmissions. Eventually, WPA encryption with TKIP will be obsolete as you replace older devices that use only TKIP.

If you need to accommodate legacy devices with an SSID, enable WPA encryption with the TKIP cipher. Keep in mind that this has an effect on performance. The additional AES cipher takes more computing power to run than simple TKIP does, therefore older, smaller devices may not support it.

NOTE: You can create different WLAN Service profiles (SSIDs) for different levels of encryption. This maximizes the use of WPA2 security.

Security always affects performance, so it is really up to you how much bandwidth and processing time you want to devote to it. With newer devices, this is much less of an issue because new devices have plenty of resources for the highest level of security, WPA2 with CCMP.

RELATED DOCUMENTATION

[Creating and Managing a WLAN Service Profile | 1090](#)

[Network Director Documentation home page](#)

Understanding PSK Authentication

IN THIS SECTION

- [What Is PSK? | 903](#)
- [How Does PSK Work? | 903](#)
- [When Would I Use PSK Authentication? | 904](#)
- [Why Would I not Use PSK Authentication? | 904](#)
- [How Is WPA Encryption Different from WPA-PSK Encryption? | 905](#)

Pre-Shared Key (PSK) is a client authentication method that uses a string of 64 hexadecimal digits, or as a passphrase of 8 to 63 printable ASCII characters, to generate unique encryption keys for each wireless client. PSK is one of two available authentication methods used for WPA and WPA2 encryption on Juniper Networks wireless networks. PSK is not the default authentication method when creating a WLAN Service profile because the other choice, 802.1X authentication, is the standard and is stronger.

NOTE: 802.1X and PSK authentication types can be applied simultaneously—clients will use the most secure option that they are capable of using. For more information about 802.1X authentication, see [“Understanding the IEEE 802.11 Standard for Wireless Networks”](#) on page 1076.

This topic describes:

What Is PSK?

There are two WPA forms of encryption available with Network Director: Wi-Fi Protected Access (WPA) and the newer WPA2. Pre-shared key (PSK), a shared secret method, can be added to either encryption method:

- WPA/WPA2 Enterprise (requires a RADIUS server) and provides coverage for large entities.
- WPA/WPA2 Personal (also known as WPA-PSK) is appropriate for use in most residential and small business settings.

How Does PSK Work?

With PSK, you configure each WLAN node (access points, wireless routers, client adapters, bridges) not with an encryption key, but rather with a string of 64 hexadecimal digits, or as a passphrase of 8 to 63

printable ASCII characters. Using a technology called TKIP (Temporal Key Integrity Protocol), that passphrase, along with the network SSID, is used to generate unique encryption keys for each wireless client. Those encryption keys are constantly changed. When clients connect, the PSK authentication users provide the password to verify whether to allow them access to a network. As long as the passwords match, a client is granted access to the WLAN.

NOTE: You have the option to encrypt the PSK plain-English passphrase.

When Would I Use PSK Authentication?

PSK was designed for home and small office networks that do not require the complexity of an 802.1X authentication server. Some reasons to use PSK authentication are:

- PSK is simple to implement, as opposed to 802.1X authentication, which requires a RADIUS server.
- Your legacy clients might not support 802.1X or the latest WPA2 standard. You can use both WPA/WPA2 and PSK simultaneously to accommodate all clients.

Why Would I not Use PSK Authentication?

Even if you have a small company, there are drawbacks to using PSK authentication. For example:

- If an administrator leaves the company, you should reset the PSK key. This can become tiresome and be skipped.
- If one user is compromised, then all users can be hacked.
- PSK cannot perform machine authentication the way that IEEE 802.1X authentication can.
- Keys tend to become old because they are not dynamically created for users upon login, nor are the keys rotated frequently. You must remember to change the keys and create keys long enough to be a challenge to hackers. PSK is subject to brute force key space search attacks and to dictionary attacks.
- Because WPA2-Personal uses a more advanced encryption type, additional processing power is required to keep the network functioning at full speed. Wireless networks that use legacy hardware for access points and routers can suffer speed reductions when WPA2-Personal is used instead of WPA, especially when several users are connected or a large amount of data is moving through the network. Because WPA2-Personal is a newer standard, firmware upgrades can also be required for some hardware that previously used WPA exclusively.

How Is WPA Encryption Different from WPA-PSK Encryption?

The primary difference between WPA and WPA2-Personal are the encryption ciphers used to secure the network. WPA can use only the encryption cipher Temporal Key Integrity Protocol (TKIP). WPA2-Personal can use TKIP, but because TKIP security keys are less secure, the WPA2 protocol usually uses the Advanced Encryption Standard. AES uses a much more advanced encryption algorithm that cannot be defeated by the tools that overcome TKIP security, making it a much more secure encryption method.

RELATED DOCUMENTATION

[Understanding the IEEE 802.11 Standard for Wireless Networks | 1076](#)

[Understanding WLAN Service Profiles | 885](#)

[Creating and Managing a WLAN Service Profile | 1090](#)

[Network Director Documentation home page](#)

Understanding Web Portals

IN THIS SECTION

- [Why Use a Web Portal on Your Wireless Network? | 906](#)
- [How Does MSS Support Web Portals? | 906](#)
- [How Does Web Portal WebAAA Work? | 906](#)
- [How Are Web Portals Created in Network Manager? | 907](#)

WebAAA provides a simple and universal way to authenticate any user or device by using a Web browser. A common application of WebAAA is to control access for guests on your network. When a user requests access to an SSID or attempts to access a Web page before logging onto the network, MSS displays a login page in the user's browser.

Aggregated Web portal information is typically presented on a grid layout customized for a particular audience.

This topic describes:

Why Use a Web Portal on Your Wireless Network?

Employees have login names and passwords for daily access, but sometimes temporary access is needed for visitors, meeting rooms, or event access. A Captive Portal enables temporary users to enter your network after they complete some steps on your Web page.

How Does MSS Support Web Portals?

WebAAA provides a simple and universal way to authenticate any user or device by using a Web browser. A common application of WebAAA is to control access for guests on your network. When a user requests access to an SSID or attempts to access a Web page before logging onto the network, MSS displays a login page to the user's browser. After the user enters a username and password, MSS validates the user information about the local database or RADIUS servers and grants or denies access based on whether the user information is found. MSS redirects an authenticated user back to the requested Web page, or to a page specified by the administrator. WebAAA, like other types of authentication, is based on an SSID or on a wired authentication port. You can use WebAAA on both encrypted and unencrypted SSIDs. If you use WebAAA on an encrypted SSID, you can use static WEP or WPA with PSK as the encryption type. MSS provides a default login page but you can alternately add custom login pages and configure MSS to display these pages instead.

How Does Web Portal WebAAA Work?

For a wireless users, the connection process begins when the network interface card (NIC) associates with an SSID. MSS starts a portal session for the user and assigns the user to the VLAN set associated with the SSID's Service Profile. A Web browser is opened and sends a DNS request for the IP address of the home page or a URL requested by the user.

After the user enters a username and password, MSS validates the user information about the local database or RADIUS servers and grants or denies access based on whether the user information is found. MSS redirects an authenticated user back to the requested Web page, or to a page specified by the administrator.

MSS does the following:

- First, MSS intercepts the DNS request, and uses MSS DNS proxy to obtain the URL IP address from the network DNS server.
- Then MSS sends the IP address to the user's browser.
- MSS then serves a login page to the WebAAA user.

The user enters a username and password in the WebAAA login page.

MSS authenticates the user by checking RADIUS or the MX local database for the username and password. If the user information is present, MSS authorizes the user based on the authorization attributes set for the user.

TIP: MSS ignores the VLAN-Name or Tunnel-Private-Group-ID attribute associated with the user, and leaves the user in the VLAN associated with the SSID Service profile (if wireless) or with the web-portal-wired user (if the user is on a wired authentication port).

After authentication and authorization are complete, MSS changes the user session from a portal session with the name web-portal-ssid or web-portal-wired to a WebAAA session with the user name. The session remains connected, but now the session is identity-based instead of a portal session.

MSS redirects the browser to the URL initially requested by the user or, if the URL VSA is configured for the user, redirects the user to the URL specified by the VSA. The Web page for the URL that the user is redirected appears in the browser window.

How Are Web Portals Created in Network Manager?

See [“Creating and Managing Device Common Settings” on page 291](#).

RELATED DOCUMENTATION

[Creating and Managing Device Common Settings | 291](#)

[Network Director Documentation home page](#)

Understanding Local Switching on Access Points

IN THIS SECTION

- [Why Use Local Switching? | 908](#)
- [How Does Local Switching Work? | 908](#)
- [When to Use Local Switching? | 909](#)
- [When not to Use Local Switching? | 910](#)
- [VLAN Profiles and Local Switching | 911](#)
- [How Do I Configure Local Switching? | 911](#)
- [Does a Web Portal Work with Local Switching Configured? | 911](#)
- [Does QoS Policy Enforcement Work with Local Switching Configured? | 911](#)
- [What Happens if the Controller or WAN Link Goes Down? | 911](#)
- [Is Local Switching Included in any Other Wireless Features? | 912](#)

Local packet switching makes packets switch directly from access points to the wired network, instead of passing through a controller. When an access point is configured to perform local switching, the controller is removed from the forwarding path for client data traffic and the client VLAN is directly accessible through the wired interface on the access point. Packets can be switched directly to and from this interface.

Instead of using the WLAN controller for deep packet inspection, packet forwarding, and in some cases encryption, local switching offloads these functions to access points on a per service set identifier (SSID) or per application basis. This approach is more efficient and more reliable, because it leverages the processing capacity that each new access point adds to the network, and reduces the likelihood of congestion and high latency than can cripple real-time applications. Only control traffic gets tunneled back to the controller and the data traffic stays local on the switch.

Why Use Local Switching?

By allowing local switching at the access point, WLAN controllers are relieved of the packet forwarding overhead, and have more of their processing capacity available for dealing with control-plane traffic and security policy administration. The result is fewer controllers are needed to support all the access points in the network. Bypassing the controllers for traffic forwarding, also results in more efficient traffic flows across the network core, because it eliminates backhauling traffic across the network to be switched by the WLAN controllers in the Data Center.

How Does Local Switching Work?

When local switching is enabled on an access point, control traffic is managed by a controller and data traffic is handled by the local switches using CAPWAP.

Local switching uses two different data tunneling methods, depending on whether a VLAN is configured. If there is a VLAN for the local switching session and it is set to overlay mode, switching is done the same way it is done with local switching disabled—using an overlay tunnel. When no VLAN is configured on the access points, local switching uses a mobility domain tunnel to transport packets.

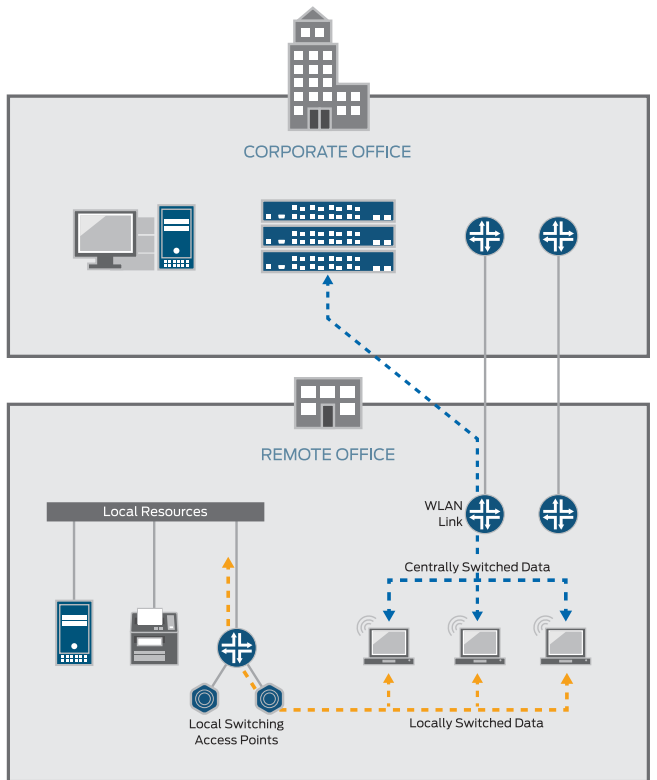
The following is a summary of the two tunnels:

- **Overlay tunnel**—A data tunnel that exists between a controller and access points. When a client session is not locally switched, the overlay tunnel is used for client data traffic. Sessions that use this type of tunnel include:
 - Sessions with local switching disabled.
 - Sessions with local switching enabled but with the VLAN for the session set to overlay mode in the VLAN profile on the access point.
- **Mobility Domain tunnel**—Data tunnel between two controllers, two access points, or a controller and an access point. It is used when an access point or controller with local switching enabled does not have the client VLAN configured on it and the traffic is tunneled from another controller or access point.

When to Use Local Switching?


When access points are deployed at remote locations, they are often configured to utilize a local switching configuration, so that user traffic stays at the remote site and the access points send the controller only various status updates and client authentication frames. Because the access points and controller exchange very little traffic in this state, they can often work quite well over a low bandwidth connection. Local switching solutions for remote offices eliminate the additional expense of distributed controllers at each and every remote site as shown in [Figure 39](#).

Figure 39: Remote Office Using Wireless Local Switching



Applications that require fast switching, such as voice, benefit from local switching—see [Figure 40](#).

Figure 40: Voice Application Using Wireless Remote Switching



When not to Use Local Switching?

Local switching is extremely beneficial for simplifying management, but there can be issues when it is used in large environments with hundreds of access points. With too many access points requiring trunk and VLAN configuration on every access point port, in addition to the time required to manage them all on the downlinks back to the core, it can be more cost effective to deploy additional controllers. Additional issues include:

- Instability due to extremely large broadcast domains spanning every access layer switch.
- The large number of MAC addresses that have to be learned on every switch.
- Administrative overhead needed to create a dot1q trunk for every access point deployed.

The following restrictions apply to access point local switching:

- Restricting Layer 2 forwarding for a VLAN is not supported if the VLAN is configured for local switching.
- The DHCP restrict feature is not supported for locally switched clients.
- When the **set ap apnum port portnum type** command is used to specify a port on a directly attached access point, the access point cannot be configured for local switching. However, a directly connected access point with an unspecified port can perform local switching.
- IGMP snooping is not supported with local switching.

VLAN Profiles and Local Switching

A VLAN profile consists of a list of VLANs and tags. When a VLAN profile is applied to an access point, traffic for the specified VLANs is locally switched by the access point instead of by the controller.

The tag on the VLAN for access points that are using local switching must be unique and not match any tag on any VLAN configured on the controller. This is because the wireless operating system uses both the name and the VLAN ID for various tasks such as moving client sessions, authentication requests, roaming, and data traffic.

How Do I Configure Local Switching?

A VLAN profile consists of a list of VLANs and tags. When a VLAN profile is applied to an access point, traffic for the specified VLANs is locally switched by the access point instead of by the controller.

From Network Director, enable local switching for access points by following the directions in [“Creating and Managing Local Switching Profiles” on page 1005](#), [“Assigning a Local Switching VLAN Profile to Existing Access Points” on page 1011](#), [“Assigning a Local Switching Profile During Access Point Configuration” on page 1013](#), [“Creating and Managing Remote Site Profiles” on page 1014](#), and [“Assigning a Local Switching Profile During Access Point Configuration” on page 1013](#).

Does a Web Portal Work with Local Switching Configured?

Yes, a Web portal works with local switching configured. When the VLAN for the Web Portal service exists locally on the controller, the controller must have an IP interface configured on this VLAN (either statically or by enabling the DHCP-client feature on the controller for this VLAN). However, with the local switching feature, it is possible that the VLAN for the WLAN client exists only at the uplink port of the access point with which the client is associated.

Does QoS Policy Enforcement Work with Local Switching Configured?

Yes, policy enforcement works with local switching configured. An ingress map determines how DSCP values are classified into CoS values. An egress map determines how CoS values are marked into DSCP. The controller and associated access points share the same set of maps, which means the QoS policy enforcement will work with or without local switching enabled.

What Happens if the Controller or WAN Link Goes Down?

Locally switched access points are managed either through the WAN or through the Internet by controllers at headquarters, and will maintain local session persistence indefinitely, if the WAN link goes down. If the connection to the controller is lost, wireless services continue uninterrupted—connected clients maintain wireless connection to the access point, new clients can connect and authenticate locally, and the Wireless Intrusion Detection System (WIDS) continues to work.

Is Local Switching Included in any Other Wireless Features?

Local Switching Profiles can also be included in Remote Site Profiles (see [“Creating and Managing Remote Site Profiles” on page 1014](#)), when adding an access point to Network Director (see [“Adding and Managing an Individual Access Point” on page 1156](#)), and in WLAN Voice Profiles.

RELATED DOCUMENTATION

[Creating and Managing VLAN Profiles | 502](#)

[Creating and Managing Local Switching Profiles | 1005](#)

[Assigning a Local Switching VLAN Profile to Existing Access Points | 1011](#)

[Creating and Managing Wireless Auto AP Profiles | 980](#)

[Creating and Managing Remote Site Profiles | 1014](#)

[Network Director Documentation home page](#)


Understanding Wireless Bridging

IN THIS SECTION

- [Why Use Wireless Bridging? | 913](#)
- [How Does Wireless Bridging Work? | 914](#)
- [How is Wireless Bridging Configured? | 914](#)

Wireless bridging is a common technique to accommodate multiple buildings in one company or other entity. Legitimate wireless bridging between two buildings is accomplished with an extension of wireless mesh, a service that enables access points to connect to each other using their radios. You must use mesh to enable bridging. (For more information about wireless mesh, see [“Understanding Wireless Mesh”](#) on page 894.) The most typical application of wireless bridging is to provide network connectivity between two buildings using an outdoor wireless link, as shown in [Figure 41](#).

Figure 41: Wireless Bridging Between Two Buildings



A wireless bridge uses two access points as the bridge endpoints in a transparent Layer 2 bridge. As long as there is a clear RF line of sight between the access points, and the distance between the access points is within the capability of the access points, the bridge can connect another building or remote area.

This topic describes:

Why Use Wireless Bridging?

Wireless bridging is frequently used with outdoor access points to connect two buildings. In [Figure 41](#), the network is located in Building 1. Instead of running cable between the two buildings, which is more expensive, two outdoor access points have been configured for wireless mesh. The access point outside of Building 2 is connected to the network in Building 1—that is why the access point on both buildings is a dual-radio network portal. For a description of mesh portals and mesh access points, see [“Understanding Wireless Mesh”](#) on page 894.

How Does Wireless Bridging Work?

Wireless bridging is a special case wireless mesh network. A wireless bridge is established between a wired mesh portal access point and an associated mesh access point as shown in [Figure 41](#). A wireless bridge can also consist of two mesh portals. In this case, both outdoor access points must be dual-radio access points, with one of their radios configured with a mesh WLAN Service profile.

When a bridge is operating, a mesh access point serving as a bridge endpoint picks up packets from the wired port and transfers them to the other bridge endpoint. A simple source-destination learning mechanism is used to avoid forwarding unnecessary packets across the bridge.

How is Wireless Bridging Configured?

You need to do the following to create a wireless bridge:

- Configure a mesh SSID (WLAN Service profile)—be sure to enable Bridging when you enable Mesh on the Basic Settings tab.
- Connect access points to a controller and configure them to work as the points of a wireless bridge.
- Disconnect the access points from the controller and place them in the bridge configuration.

Understanding Wireless Interference

IN THIS SECTION

- [What Causes Wireless Radio Frequency Interference? | 915](#)
- [Effects of Interference Seen by Clients | 915](#)
- [You Can Monitor RF Interference with Network Director | 915](#)
- [What Is RF Jamming? | 916](#)

Wi-Fi interference is a common and troublesome issue. The lack of wires that makes WLAN so attractive is also the feature that makes other consumer devices capable of causing Wi-Fi interference. Your WLAN network might be working fine one day and sluggish the next day, without you having made any network changes, all due to interference.

This topic describes:

What Causes Wireless Radio Frequency Interference?

Because the air is shared by all transmitters, transmissions by any device at the same frequency as an access point's radio can cause interference. Because 802.11 wireless networks operate in unlicensed bands used by many technologies, such as microwave ovens, video surveillance cameras, cordless phones, they are subject to interference. In addition, wireless access points sharing the same channel might interfere with each other. The effect of interference is highly dependent on the strength of the transmission and the distance from the interferer. Access points closest to and on the same channel as an interferer will be affected more than those that are further away.

These are some common causes of wireless interference:

- Leaving the channel number on each radio set to the default value can result in high interference among the radios because too many radios are sharing the bandwidth on one channel.
- Hidden nodes in a wireless network referring to nodes that are out of range of other nodes or a collection of nodes. A hidden node can generate a high number of cyclic redundancy check (CRC) code errors.
- Co-channel interference or adjacent channel interference can result from setting radios to bands that have overlapping channels. The channels might not all be in use by your network—neighboring company signals can also cause interference.
- Some non-network devices, such as microwave ovens, car alarms, cordless phones, or wireless video cameras can interfere with wireless channels. Most often, these devices are using the 2.4-GHz frequency.
- Bad electrical connections can cause broad RF spectrum emissions.
- RF jamming is a deliberate attempt to disrupt the network with a powerful signal.

Effects of Interference Seen by Clients

Your network clients might notice the results of interference before you do. They might complain of network slowdown, but not of data loss. This slowdown might not be immediately obvious with low capacity data transmission because, if interference is intermittent, packets eventually get through. Therefore, there is no packet loss, just retransmissions that take time. Another possibility is that some devices, such as microwaves, reduce throughput without blocking it entirely. Complaints will increase when more users log in, increasing data capacity until data loss occurs, or when Voice Over IP calls are placed. VoIP requires significant bandwidth because resending voice is not an option—the result is dropped or jittery voice transmission.

You Can Monitor RF Interference with Network Director

Network Director includes a Monitor Mode that displays the compiled RF data gathered by scanning the mobility domain of your network. There are two monitors that indicate network interference for access points and radios:

- **AP Interference Sources**—On this chart, interference sources for each radio or each access point (either one can be selected) are sorted into categories such as Microwave Oven, Phone FHSS, and Continuous Wave. The occurrences of each source are tracked and added—the sum appears on a bar chart. The categories with the tallest bars are those causing the most interference on this radio or access point, but the sum of all interference is what really matters. For more information, see [“Monitoring RF Interference Sources on Wireless Devices” on page 1329](#).
- **Radio Interference Sources**—The pie chart for radio interference is also sorted into categories such as Microwave Oven, Phone FHSS, and Continuous Wave. In addition, a list of details are provided about each source of interference—time last seen, transmitter ID, listener MAC address of the access point that found the interference, channel, RSSI, duty cycle, and percentage of compliance to the common information model (CIM). For more information, see [“Monitoring RF Interference Sources on One Radio” on page 1325](#).
- **RF Interference For Radios on One Access Point**—The bar chart that displays each radio on an access point lets you compare the interference experienced on two radios when the access point has dual radios. With single radio access points, you see the same data as you do for the radio interference pie chart, but in bar chart format. For more information, see [“Monitoring RF Interference Sources For Radios on One Access Point” on page 1328](#).

What Is RF Jamming?

RF jamming is a DoS attack. The goal of RF jamming is to take down an entire WLAN by overwhelming the radio environment with high-power noise. A symptom of an RF jamming attack is excessive interference. If an access point radio detects excessive interference on a channel, and channel auto-tuning is enabled, the radio changes to a different channel. The radio continues to scan on an active data channel and on other channels and reports the results to the controller.

Jamming occurs at the physical layer of the network, saturating the channel or band with noise and making it difficult or impossible for a receiving radio to detect a real transmission. Think of jamming as trying to hear someone talking as a siren goes off. The increased noise floor results in a poor signal-to-noise ratio (SNR), usually detected by the clients as poor signal quality. Jamming can also be detected by an access point, which then triggers dynamic temporary tuning of the channel if automatic channeling is enabled. However, selecting a different channel does not always stop jamming. An experienced attacker will often use all available channels in the attack.

RELATED DOCUMENTATION

[Monitoring RF Interference Sources on Wireless Devices | 1329](#)

[Monitoring RF Interference Sources on One Radio | 1325](#)

[Monitoring RF Interference Sources For Radios on One Access Point | 1328](#)

[Troubleshooting Excessive Wireless Interference | 1332](#)

Understanding Rogue Access Points

IN THIS SECTION

- [What is a Rogue Access Point? | 917](#)
- [How Are Rogue Access Points and Rogue Clients Identified By Controllers? | 918](#)
- [How are Rogue access points and Rogue Clients Classified as Rogue? | 918](#)
- [What Harm Can a Rogue Access Point Do? | 920](#)
- [Section | ?](#)
- [What Can I do To Prevent Rogue Access Points? | 921](#)
- [How Do I Prevent a Benign Access Point From Being Classified as a Rogue? | 923](#)

One of the most common wireless security threats is the rogue access point—it is used in many attacks, both DoS and data theft. Many other rogue access points, however, are deployed by employees wanting unfettered wireless access—these access points are called soft access points. Other rogues are located in neighboring companies using your network for free access. Typically low-cost and consumer-grade, these access points often do not broadcast their presence over the wire and can only be detected over-the-air. Because they are typically installed in their default mode, authentication and encryption are not enabled, thereby creating a security hazard. Because wireless LAN signals can traverse building walls, an open access point connected to the corporate network the perfect target for war driving. Any client that connects to a rogue access point must be considered a rogue client because it is bypassing the authorized security procedures put in place by the IT department.

This topic includes the following:

What is a Rogue Access Point?

A rogue access point is a device not sanctioned by an administrator, but is operating on the network anyway. This could be an access point set up by either an employee or by an intruder. The access point could also belong to a nearby company.

These are some reasons to suspect that an access point is a rogue:

- The SSID of the access point is neither your network SSID nor listed in the permitted SSID list. (See .) The access point may not be broadcasting an SSID at all. Check the SSID of an access point using any of these methods:
 - From the MSS CLI,
 - From Network Director, .
 - From MSS,
- The access point is masquerading one of your SSIDs. Access points masquerading your SSID are rogue by default—you can, however, change that policy. See [“Understanding an SSID Masquerade” on page 926](#).
- The access point is an ad-hoc access point, formed directly between two client devices. See [“Understanding Ad-Hoc Networks” on page 927](#). Ad-hoc access points are rogues by default—you can, however, change that policy—see [Configuring Your Ad-Hoc AP Policy](#).
- Network management features of the access point, such as SNMP, HTTP, and Telnet have been disabled.
- The access point's MAC address does not appear in ARP tables.
- The access point is operating as a bridge—see [“Understanding Wireless Bridging” on page 912](#) .
- The access point is listed in the rogue list, where it has been added by an administrator.

How Are Rogue Access Points and Rogue Clients Identified By Controllers?

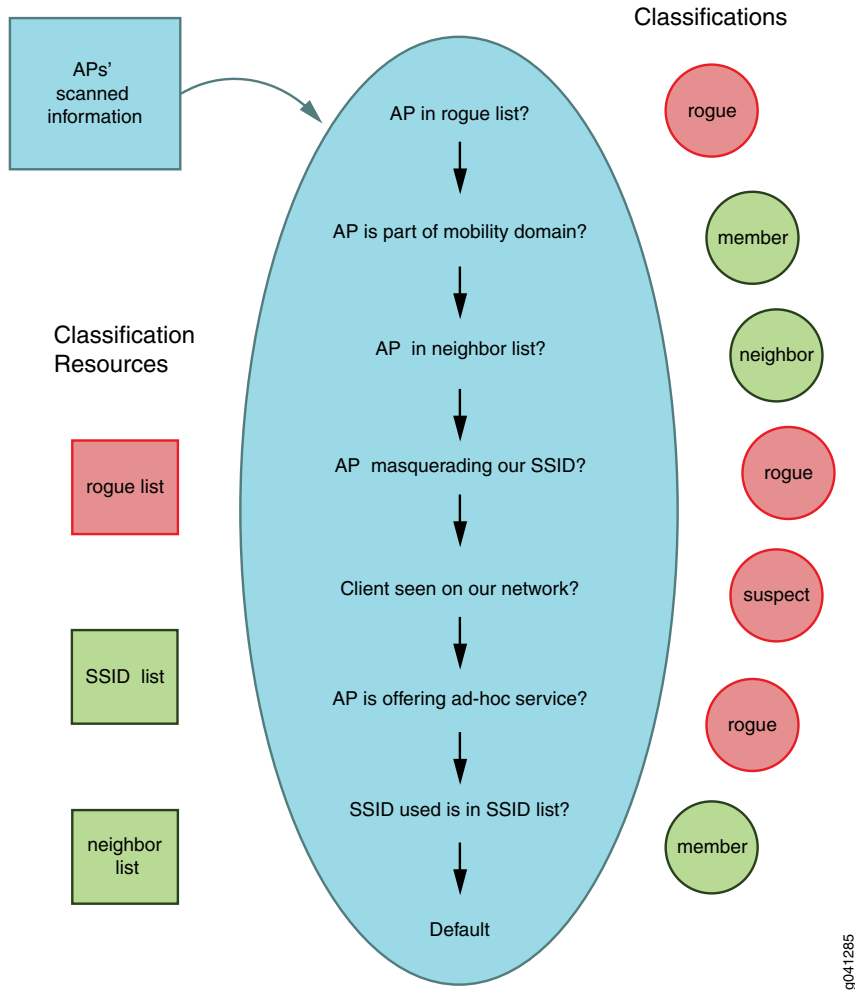
Wireless radios automatically scan the RF spectrum for other access points transmitting in the same spectrum. The RF scans discover third-party transmitters in addition to other Juniper radios. Controllers consider all non-Juniper transmitters to be suspects (potential rogues) by default. If the device is a Juniper device, but the MAC address is not in the appropriate database, a series of rules determine whether that device is a rogue. Once an access point is declared a rogue, it is reported by MSS:

- From the MSS CLI,
- From Network Director,
- From MSS,

How are Rogue access points and Rogue Clients Classified as Rogue?

Controllers use a set of rules, illustrated in [Figure 42](#), in order to classify unknown access points as either members, neighbors, suspects, or rogues.

Figure 42: How Scanned Information is Used to Classify Access Points



The definition of each classification--member, neighbor, suspect, or rogue—is listed in [Table 195](#).

Table 195: Classifications Define a Rogue

RF Classification	Description
Member	Access point is in this mobility domain. Access point fingerprint (also referred to as signature) is used to securely identify member access points.
Neighbor	Trusted device (good neighbor) is listed in the permitted third-party SSID list. Usually, this access point is part of a neighboring wireless network or mobility domain.
Suspect	Not enough information to classify this access point as neighbor or rogue. You may decide to add it to the rogue list, SSID list or neighbor list.

Table 195: Classifications Define a Rogue (continued)

RF Classification	Description
Rogue	Rogue device (bad guy) on the air. For example, unauthorized access point on an enterprise network.

You Can Change Some Rogue Classification Rules

Classification rules are either built-in or selected by you from a set of pre-defined rules. Built-in rules are constant and cannot be changed. User rules are the rules that let you configure certain classification behaviors.

Notice that the first classification rule eliminates access points in the rogue list and cannot be altered. Two configurable rules default to rogue classification and you can set a third to classify the default condition as rogue.

User Rule	#	Rules for RF Classification	Classification
N	1	If access point in rogue list	rogue
N	2	If access point is part of mobility domain	member (never a rogue)
N	3	If access point in neighbor list	neighbor (never a rogue)
Y	4	If access point is Masquerading our SSID	rogue (default)
Y	5	Client or Client DST MAC seen in network	rogue (default)
Y	6	If access point is acting as an Ad-hoc device	skip-test (default) can be set to rogue
N	7	If SSID is in SSID list	neighbor (never a rogue)
Y	8	Default Classification	suspect (default) can be set to rogue

What Harm Can a Rogue Access Point Do?

Rogue access points and their clients undermine the security of an enterprise network by potentially allowing unchallenged access to the network by any wireless user or client in the physical vicinity. Rogue access points can also interfere with the operation of your enterprise network. Rogue access points can do the following damage:

- Allow a hacker to conduct a man-in-the-middle attack. The attacker makes independent connections with the victims and relays messages between them, making them believe that they are talking directly to each other over a private connection, when in fact the entire conversation is controlled by the attacker.
- Flood the network with useless data, creating a denial of service.
- Send fake SSIDs advertising attractive features such as free Internet connectivity. Once a user connects, the fake SSID is added to the client's wireless configuration and the client begins to broadcast the fake SSID, thereby infecting other clients. .
- Provide a conduit for the theft of company information..

What Can I do To Prevent Rogue Access Points?

There are a number of actions you can take that make it more difficult for a rogue to penetrate your network. See [Table 196](#) for details.

Table 196: Preventing Rogue Access Points

Do This:	Result of Action:
Establish strict rules and make sure they are well published.	<p>Only authorized IT staff can connect networking equipment.</p> <p>All devices that connect to the network, including wireless access points, conform to company security policies.</p> <p>NOTE: Some colleges even expel students who are caught with rogue access points or ad-hoc networks.</p>
Change the rogue classification rules.	By default, unknown devices are classified as suspects. When you change this default to rogue, the controller automatically classifies any third-party access point or client as a rogue, and you can optionally isolate the access point by dropping all packets to and from the device.
Eliminate benign access points from the rogue list so that real rogues stand out.	When you add safe networks' SSIDs and/or vendor names to the list of SSIDs allowed on the network, these access points cannot be classified as rogues. .
Add known intruders to the rogue list.	Third-party access points are isolated when they are added to the rogue list. All packets are dropped to and from access points.

Table 196: Preventing Rogue Access Points (continued)

Do This:	Result of Action:
Use strong security	The IEEE 802.11i security standard uses IEEE 802.1X for mutual authentication between the network and the client. This means that clients that try to access network resources must be authenticated by the network. In a similar vein, the client verifies the authenticity of the network infrastructure it is attaching to before beginning data transmission. With 802.1X, the credentials used for authentication, such as login passwords, are never transmitted without encryption over the wireless medium. In addition, 802.1X provides dynamic per-user, per-session encryption keys, removing the administrative burden and security issues associated with static encryption keys. Security is configured in WLAN profiles.
Use active access point scanning in addition to passive scanning.	Active scans send probes with a null SSID name to look for rogue access points and clients. Active scan is enabled by default on radio-profiles. We recommend that you do not change this setting.
Make sure the Juniper wireless signature is enabled and change it on a regular basis.	A wireless signature is a set of bits in a management frame sent by an access point as an identifier. If someone attempts to spoof management packets from a Juniper access point, Network Director can detect the spoof attempt.
Make sure logging is enabled and check log messages and traps for suspicious activity.	By default, a controller generates a log message when a rogue is detected or disappears. For details, see “Collecting Logs for Troubleshooting” on page 107 .
Immediately investigate ad-hoc access points and either add security to them or eliminate them.	An ad hoc network is one that is formed directly between two client devices. Ad hoc networks pose a threat to the enterprise because the security checks imposed by the infrastructure are bypassed. One of the dangers is an employee who brings in a wireless-enabled laptop, plugs it into a wired port at work, and leaves the wireless interface enabled. In this scenario, a hacker in a neighboring area could connect directly to the client, creating a security threat. The hacker at this point could look for information on the employee’s client device, and potentially gain access to the corporate network through the simultaneous wireless and wired interfaces. This situation may place the enterprise in violation of regulatory policies for its industry. The security hole provided by ad-hoc access points is not the ad-hoc network itself but the bridge it provides into other networks.
Immediately investigate wireless bridge frames and eliminate the source.	An attacker often sets up a laptop with two wireless adaptors—one card is used by the rogue access point and the other is used to forward requests through a wireless bridge to the legitimate access point. .

Table 196: Preventing Rogue Access Points (*continued*)

Do This:	Result of Action:
Use managed switches on your network and use their port-based security to allow only certain MAC addresses or disable unused ports.	An access point randomly plugged into ports on this switch will not work.
Consider using static IP addresses instead of having them assigned by a DHCP server.	When you use static IP addresses, an intruder who installs a rogue access point needs to manually assign an IP address to the access point before it can gain access to the network.
Enable automatic countermeasures to immediately react to rogues or suspect rogues.	Countermeasures can attack or isolate rogue and/or suspect transmitters using various methods of attack.

How Do I Prevent a Benign Access Point From Being Classified as a Rogue?

access points belonging to your mobility domain are never classified as rogues.

Presence of third-party access points on a permitted SSID list or OUI list does not guarantee that the device will not be classified as a rogue for other reasons. The only sure way to be sure a non-mobility domain device is not classified as a rogue is to add the device or vendor to the neighbor list.

Neighbors are devices known to be part of a neighboring network and non-threatening. Vendors can also be added to the neighbor list, so that all of the devices from that vendor become neighbors.

RELATED DOCUMENTATION

[Understanding Rogue Clients](#) | 923

Understanding Rogue Clients

IN THIS SECTION

- [What Defines a Rogue Client?](#) | 924
- [How Are Rogue Clients Detected?](#) | 925
- [What Can I do To Prevent Rogue Client Damage?](#) | 925

- [How Do I Prevent a Benign Client From Being Classified as a Rogue? | 925](#)
- [How Do I Make Sure An SSID Won't Be Classified as Rogue? | 925](#)
- [How Can I Make Sure a Device is Classified as Rogue? | 925](#)

A rogue client is a client that does not belong to your company, but is operating on the network anyway. Rogue clients might be trying to steal information, could be trying to disrupt normal wireless service by launching attacks, or might be simply trying to use your wireless access.

This topic includes the following:

What Defines a Rogue Client?

A client is automatically classified as a rogue if it is listed in the rogue list, where it has been added by an administrator - see [“Creating and Managing RF Detection Profiles” on page 1026](#). In addition, the MSS OS automatically classifies a client as a rogue for the following reasons:

- Ad-hoc clients such as laptops, PDA's, and printers attempting to connect directly, without using an access point, are rogues by default. You can change this by [“Creating and Managing RF Detection Profiles” on page 1026](#).
- Any client that connects to a rogue access point is considered a rogue client because it is bypassing the authorized security procedures put in place by the IT department. .
- The client has violated a network policy.

Additional indications that a client might be a rogue include:

- The client sends multiple frames with prolonged duration. The duration value in the frame indicates the duration in milliseconds for which the channel is reserved, so long duration values can disrupt legitimate users. .
- The client is not connected to an access point—this is called an unassociated client—but is sending packets anyway. The packets are probably forged packets that the client is injecting into the wireless network.
- The client is probing for 'any' SSID. When access points are not configured properly, they can allow clients to connect with 'any' SSID. A client tries to connect using 'any' SSID it would most probably be a rogue client.
- The client is repeatedly sending association requests to an access point. This could be indication of a flood attack. The threshold for triggering a flood message is 100 frames of the same type from the same MAC address, within a one-second period. If MSS detects more than 100 of the same type of wireless

frame within one second, it generates a log message. The message indicates the frame type, the MAC address of the sender, the listener (MP and radio), channel number, and RSSI. To see the log,

How Are Rogue Clients Detected?

Access points have the ability to identify clients, including rogue clients. This information is passed to the controller and to Network Director. In Network Director, rogue clients generate a fault. From the CLI, use the command **show rfdetect clients** to see all clients. From RingMaster, click **Alarms > Query >** select options including the word rogue > **OK**.

What Can I do To Prevent Rogue Client Damage?

The sooner you detect a rogue client or suspicious activity, the easier it is to stop the intrusion.

NOTE: A client belonging to your network mobility domain cannot be classified as a rogue, but you can add them to the blacklist—see [“Creating and Managing RF Detection Profiles” on page 1026](#).

How Do I Prevent a Benign Client From Being Classified as a Rogue?

Clients belonging to your mobility domain are never classified as rogues.

Presence of other clients on a permitted SSID list or OUI list does not guarantee that the device will not be classified as a rogue for other reasons. The only sure way to be sure a non-mobility domain client is not classified as a rogue is to add the device or vendor to the neighbor list.

Neighbors are devices known to be part of a neighboring network and non-threatening. Vendors can also be added to the neighbor list. For directions, see [“Creating and Managing RF Detection Profiles” on page 1026](#).

How Do I Make Sure An SSID Won't Be Classified as Rogue?

MSS maintains a permitted SSID list, which is a list of SSIDs allowed in the mobility domain. If a detected SSID is not on the list, MSS generates a message.

How Can I Make Sure a Device is Classified as Rogue?

If you know the MAC address of a client, you can add it to the rogue list—for directions, see [“Creating and Managing RF Detection Profiles” on page 1026](#).

RELATED DOCUMENTATION

| [Creating and Managing RF Detection Profiles](#) | 1026

Understanding an SSID Masquerade

IN THIS SECTION

- [Fake SSID Attacks](#) | 926
- [Fake BSSID Attacks](#) | 926
- [Detecting Fake SSID Attacks and Fake BSSID Attacks](#) | 927

SSID Masquerade is a situation where a hacker pretends to be part of your network by using an access point that sends an SSID that looks like one of your legitimate SSIDs.

TIP: For an explanation of the terms SSID and BSSID, see [“Understanding the Network Terms SSID, BSSID, and ESSID”](#) on page 1061.

Fake SSID Attacks

With fake SSID attacks, an access point joins the network pretending to be a legitimate access point. Even though access point beacons and probe responses must carry information about the WLAN, including the BSSID, access point MAC addresses (the basis of the BSSID) are easily reconfigured—therefore, any 802.11 device can transmit packets that appear to originate from another access point or MAC address. This fake access point then becomes a conduit for stealing sensitive company information.

Fake BSSID Attacks

With fake BSSID attacks, Wi-Fi users can be tricked into associating with a phony access point. Also called Evil Twin, AP Phishing, Wi-Fi Phishing, Hotspotter, or Honeypot AP, these attacks use fake access points with faked login pages to capture credentials and credit card numbers, launch man-in-the-middle attacks, or infect wireless hosts.

Detecting Fake SSID Attacks and Fake BSSID Attacks

All SSIDs belonging to a mobility domain are part of a SIFA cluster that can be compared with SSIDs seen on the air. If an access point does not belong to your mobility domain and is seen using your SSID(s), the controller logs that instance, and a log and a trap are generated. If you want to override this alarm, you can add the device to the ignore list and the device will no longer be considered a threat.

MSS detects two kinds of these access point attacks by monitoring the neighbor table: fake access point BSSID attacks and fake access point SSID attacks. The usual purpose of a fake access point attack is to penetrate the network unobserved and obtain information, but they can also sometimes be used to infect the network.

Access points detect any packet type that is using one of the access point's own BSSIDs. When an access point detects that one or more of its BSSIDs are being spoofed, the access point creates a record indicating that this particular BSSID is being spoofed. There will be one record per spoofed BSSID.

With fake access point BSSID attacks, Wi-Fi users are tricked into associating with a phony access point. Also called Evil Twin, AP Phishing, Wi-Fi Phishing, Hotspotter, or Honeypot AP, these attacks use phony access points with faked login pages to capture credentials and credit card numbers, launch man-in-the-middle attacks, or infect wireless hosts.

With fake SSID attacks, an access point joins the network pretending to be a legitimate access point. Even though access point beacons and probe responses must carry information about the WLAN, including the BSSID, access point MAC addresses (the basis of the BSSID) are easily reconfigured so any 802.11 device can transmit packets that appear to originate from another access point or MAC address. This fake access point then becomes a conduit for stealing sensitive company information.

For directions to configure a wireless masquerade policy with Network Director, see [“Creating and Managing RF Detection Profiles” on page 1026](#).

RELATED DOCUMENTATION

| [Creating and Managing RF Detection Profiles | 1026](#)

Understanding Ad-Hoc Networks

IN THIS SECTION

- [Why Are Ad-Hoc Networks a Security Risk? | 928](#)
- [How Do I Detect an Ad-Hoc Network? | 928](#)

- [Are All Ad-Hoc Networks Malicious? | 928](#)
- [How Do I Know Whether an Ad-Hoc Network Is Malicious? | 929](#)

Ad hoc is Latin and means *for this purpose*. An ad hoc network is one that is formed directly between two client devices for a specific reason. An ad-hoc network might not be an intentionally malicious attack on the network, but it poses a threat to the enterprise because the security checks imposed by the infrastructure are bypassed, and it steals bandwidth from your infrastructure users.

Why Are Ad-Hoc Networks a Security Risk?

There are two common attacks that can be launched from ad-hoc clients on your network. Fake SSIDs can be sent from ad-hoc networks advertising attractive SSIDs such as free Internet connectivity. Once a user connects, the fake SSID is added to the client's wireless configuration and the client begins to broadcast the fake SSID, thereby infecting other clients. Also, ad-hoc clients are capable of forwarding data by flooding in addition to the classic routing technique—this forwarding can be used for flood attacks.

How Do I Detect an Ad-Hoc Network?

MSS detects and reports ad-hoc networks. In Network Director, the fault *Adhoc User Detected* appears in the RF Detect category of faults (see [“Alarms by Category Monitor” on page 1466](#), [“Alarm Detail Monitor” on page 1459](#), [“Understanding the Fault Mode Tasks Pane” on page 1452](#), and [“Alarm Summary Report” on page 1503](#).)

In Network Director, configure an ad-hoc network policy from [“Creating and Managing RF Detection Profiles” on page 1026](#).

Are All Ad-Hoc Networks Malicious?

Most ad-hoc networks are not created with malicious intentions. Laptops, PDA's, and printers with wireless enabled are simply attempting to connect to each other without using an access point—this is also called peer-to-peer networking. The security hole provided by ad-hoc networking is not the ad-hoc network itself, but the bridge it provides into other networks. One of the common scenarios is an employee who brings in a wireless-enabled laptop, plugs it into a wired port at work, and leaves the wireless interface enabled. In this scenario, a hacker in a neighboring area could connect directly to the client, creating a security threat. The hacker at this point could look for information on the employee's client device, and potentially gain access to the corporate network through the simultaneous wireless and wired interfaces. This situation might place the enterprise in violation of regulatory policies for its industry.

How Do I Know Whether an Ad-Hoc Network Is Malicious?

Does the access point have characteristics of a benign device or characteristics of a threatening device? Check the characteristics in [Table 197](#) for information.

Table 197: Characteristics of Benign Rogues and Threatening Rogues

Benign rogues tend to be:	Threatening rogues tend to be:
off of the network	on the network
secured	not secured
using a foreign SSID	using your SSID. Access points masquerading your SSID are rogue by default, but this is configurable. See “Creating and Managing RF Detection Profiles” on page 1026.
using weak RSSI	using strong RSSI
without clients	attracting clients
associating only with untrusted stations	actively associated with your stations—this indicates a man-in-the-middle attack
consuming some bandwidth	consuming a lot of bandwidth—this indicates a DoS flood or port scan

RELATED DOCUMENTATION

[Creating and Managing RF Detection Profiles | 1026](#)

[Alarms by Category Monitor | 1466](#)

[Alarm Detail Monitor | 1459](#)

[Understanding the Fault Mode Tasks Pane | 1452](#)

[Alarm Summary Report | 1503](#)

[Creating and Managing RF Detection Profiles | 1026](#)

Understanding LLDP and LLDP-MED

Link Layer Discovery Protocol (LLDP) is a vendor-neutral link layer protocol used by network devices to advertise the attributes of a device. These attributes contain type, length, and value descriptions and are referred to as TLVs. LLDP supported devices use TLVs to receive and send information to neighboring devices. Details such as configuration information, device capabilities, and device identity can also be advertised by this protocol. MSS supports the following TLVs:

- Port Description
- System Name
- System Description
- System Capabilities
- Management address

You can configure the frequency of LLDP updates, the amount of time to hold information before discarding it, and initialization delay time. You can also select the LLDP TLVs to be sent and received.

Link Layer Discovery Protocol-Media Endpoint Discovery (LLDP-MED) is an extension to LLDP that operates between endpoint devices such as IP phones to provide support for voice over IP (VoIP) applications. LLDP-MED endpoints determine the capabilities of a connected device and whether those capabilities are enabled.

TIP: LLDP and LLDP-MED cannot operate simultaneously on a network. By default, access points send only LLDP packets until LLDP-MED packets are received from an endpoint device. The access point then sends out LLDP-MED packets until it receives LLDP packets.

LLDP can be configured on both a controller and on an access point. See [“Configuring a Controller” on page 1037](#) (Configuring LLDP on a Controller) and [“Configuring Link Layer Discovery Protocol \(LLDP\) on an Access Point” on page 1035](#).

RELATED DOCUMENTATION

[Configuring Link Layer Discovery Protocol \(LLDP\) on an Access Point | 1035](#)

[Configuring a Controller | 1037](#)

Importing RingMaster Data to Network Director

You can use the Import RingMaster Data task in the System Settings to import data from the network plan that is active in RingMaster. A network plan is a repository for all hardware, site, and configuration details for a single or multi-site wireless network. It can contain wireless network data that spans multiple sites or campuses, which include one or more buildings and associated floors, and possibly outdoor areas.

Network Director uses this data to discover devices, build profiles, configure the Location view, build the Topology view (including a floor plan, if one is defined in the network plan), and to assign devices to various entities in Location and Topology views.

Before you import RingMaster data to Network Director:

- Ensure that you have the correct network plan activated in RingMaster.
- Ensure that the RingMaster server is up and is accessible through the network.
- If your RingMaster server uses any port other than Port 443, then you must open that port from the Junos Space Network Management Platform server.

To import RingMaster data to Network Director:

1. From the Tasks menu in the Logical View, Location View, Device View or the Custom Group View, click **Device Discovery > Import RingMaster Data**.
2. Specify the RingMaster server details and user credentials of the user who has administrative access to the RingMaster server, in the appropriate fields.
3. Click **Import**. Network Director uses the details that you specified in this page to connect to RingMaster and start importing data.

After you click Import, Network Director performs the following tasks:

- a. Imports the list of controllers and access points that are managed by the RingMaster server that you specified.
- b. Initiates discovery of these controllers and access points.
- c. Reads the configuration of the discovered devices.
- d. Creates the necessary profiles based on the device configuration.
- e. Assigns profiles to devices.
- f. Performs location setup based on the location data that is available in the network plan.

- g. Assigns devices to locations.
- h. Imports floor plan, if defined, to the topology view.

You can use the Job Management page to track the progress of each of these tasks.

RELATED DOCUMENTATION

[Discovering Devices in a Physical Network | 204](#)

[Troubleshooting Device Discovery Error Messages | 217](#)

[Network Director Documentation home page](#)

Creating and Managing a Radio Profile

IN THIS SECTION

- [Managing Radio Profiles | 933](#)
- [Radio Profiles Need an Associated WLAN Profile | 934](#)
- [Creating a Radio Profile | 935](#)
- [Specifying Quick Setup Radio Profile Settings | 935](#)
- [Specifying Custom Radio Profile Setup Settings | 939](#)
- [What To Do Next | 951](#)

Radio profiles control aspects of radio behavior for radios in a WLAN. A radio can have only one profile assigned at a time. Use the Network Director Radio profiles to both create a Radio profile and link it to existing WLAN Service profiles (SSIDs). For more information about the aspects and parameters of a radio file, see [“Understanding Radio Profiles” on page 879](#).

Radio profiles can be used by multiple WLAN Service profiles, so you really only need to create one Radio profile for wireless operation. However, there are a number of reasons to create multiple Radio profiles:

- For best network performance, use separate Radio profiles for transmissions using long and short guard intervals. The guard interval is the space between symbols (characters) being transmitted—it eliminates intersymbol interference. In normal 802.11 operation, the guard interval is 800 ns. In 802.11n operation, short guard intervals of 400 ns are supported. Shorter guard interval between symbols increases throughput. Guard intervals are configured in WLAN Service profiles.
- For best network performance, use separate Radio profiles for voice and data traffic because network issues such as packet loss, latency, and jitter affect VoIP much more than they affect data. With separate profiles, you can optimize settings for data in one profile and VoIP in another profile that includes call admission control and user idle timeout. These parameters are configured in Radio profiles.
- When all radios are using one channel, that channel becomes crowded. Channel use is determined in Radio profiles.
- 802.11n devices cannot perform at optimum speed when a radio also services slower 802.11b/g devices. Separate Radio profiles should be used for 802.11n and 802.11b/g devices.

This topic describes:

Managing Radio Profiles

From the Manage Radio Profiles page, you can:

- Create a new Radio profile by clicking **Add**. For directions, see [“Creating a Radio Profile” on page 935](#).
- Modify an existing profile by selecting it and clicking **Edit**.
- Assign the Radio profile to radios by selecting a Radio profile, and clicking **Assign**. For directions, see [“Assigning a Radio Profile to Radios” on page 952](#).
- Change Radio profile assignments by selecting a Radio profile, and clicking **Edit Assignment**.
- View information about a profile by selecting a Radio profile and clicking **Details**.
- Delete profiles by selecting a Radio profile and clicking **Delete**.

TIP: You cannot delete profiles that are in use—that is, assigned to objects or used by other profiles. To see the current assignments for a profile, select the profile and click **Details**.

- Clone a profile by selecting a profile and clicking **Clone**.

[Table 198](#) describes the information provided about Radio profiles on the Manage Radio Profiles page. This page lists all Radio profiles defined for your network, regardless of your current selected scope in the network view.

Table 198: Radio Profile Fields

Field	Description
Profile Name	Name given to the profile when the profile was created.
Device Family	Wireless controllers (WLC)
Description	Up to 256 alphanumeric characters, provided when the profile was created, including spaces and special characters.
Spectral Scan	Indicates whether spectral scanning is Enabled or Disabled. For more information, about scanning see “Understanding Wireless Scanning” on page 869 .
WLAN Service Profiles	Lists associated WLAN Service profiles. For more information, about WLAN Service profiles, see “Understanding WLAN Service Profiles” on page 885 .
Assignment State	A Radio profile can be Deployed on a controller , Unassigned , or Pending Deployment . For deployment directions, see “Deploying Configuration to Devices” on page 1180 .
Creation Time	Date and time when the profile was created.
Last Updated Time	Date and time when the profile was last modified.
User Name	The username of the person who created or modified the profile.

TIP: All columns might not be currently displayed. To show or hide fields in the table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

Radio Profiles Need an Associated WLAN Profile

You need at least one WLAN Service profile to create a Radio profile—they are mapped during configuration of the Radio profile. You can either use an existing Radio profile (follow the directions [“Creating and Managing a WLAN Service Profile” on page 1090](#)) or you can create a WLAN profile during the creation of the Radio profile.

Creating a Radio Profile

Radio profiles provide a set of radio parameters that can be applied to multiple radios. At minimum, you must:

- Specify a unique Radio profile name.
- Link at least one WLAN Service profile to the Radio profile.

To create a Radio profile:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.

3. In the Tasks pane, expand **Wireless**, expand **Profiles**, and then click **Radio**.

The Manage Radio Profiles page opens, displaying the list of all currently configured Radio profiles.

4. Click **Add**.

The Create Wireless (WLC) Radio Profile page opens. It consists of two tabs, **Quick Setup** and **Custom Setup**—**Quick Setup** is displayed.

5. Either complete the Quick Setup as described in both the online help and in [“Specifying Quick Setup Radio Profile Settings” on page 935](#), or complete the Custom Setup as described in both the online help and in [“Specifying Custom Radio Profile Setup Settings” on page 957](#).

6. Click **Done**.

The Radio profile appears in the list on the screen. You can now assign this profile to radios by following the directions in [“Assigning a Radio Profile to Radios” on page 952](#).

Specifying Quick Setup Radio Profile Settings

For a Radio profile minimum setup, specify the settings listed in [Table 199](#).

Table 199: Quick Radio Profile Setup

Field	Action
Name and Description	<p>The values for these attributes might have been transferred from the Basic Settings tab.</p> <p>If not, type a unique name that identifies the profile.</p> <p>Use up to 32 characters for wireless profile names. Profile names must not contain special characters or spaces. Note that profiles automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character.</p>
Task: Select Existing WLAN Service profiles	Click Select under WLAN Service Profiles and then select any or all listed WLAN Service profiles to associate with this Radio profile. Radios will broadcast the SSIDs named in the WLAN Service profile(s).

Table 199: Quick Radio Profile Setup (continued)

Field	Action
Task: Create and Add a New WLAN Service profile	

Table 199: Quick Radio Profile Setup (continued)

Field	Action
	<p>1. Click Create under WLAN Service Profiles.</p> <p>The Create WLAN Service Profile for Radio Profile window opens.</p> <p>2. Provide the following settings for a WLAN Profile:</p> <ol style="list-style-type: none"> Profile Name—Type a unique name that identifies the profile. Use up to 32 characters for wireless profile names. Profile names must not contain special characters or spaces. Note that profiles automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character. Description Service Profile Type—Select 802.1X, Voice, Web Portal, Open Access or Custom. SSID—Provide an SSID to be broadcast to clients. <p>3. Indicate security settings for the WLAN Profile by selecting either RSN, WPA, or Static WEP. Provide this additional information:</p> <ol style="list-style-type: none"> RSN—Select either AES (CCMP) or TKIP. WPA—Select either AES (CCMP) or TKIP. Select either 802.1X Authentication or PSK Authentication or both. Static WEP—Provide one or more WEP keys. <p>4. Indicate authentication settings for the WLAN Service profile. Enable either Configure Authentication Settings, or Select Existing Authentication.</p> <p>If you enabled Configure Authentication Settings, select either Create RADIUS Server or Select RADIUS Server.</p> <p>If you selected Create RADIUS Server, provide this information:</p> <ol style="list-style-type: none"> RADIUS Server Address Secret <p>If you enabled Select Existing Authentication, click Select, indicate one of the listed profiles, and then click OK. The Authentication Profile name appears in the Authentication Profile field in the Create WLAN Service Profile for Radio Profile</p>

Table 199: Quick Radio Profile Setup (continued)

Field	Action
	<p>window.</p> <p>5. Indicate authorization settings for the WLAN Service profile. Select either Configure Authorization Settings, or Select Existing Authorization.</p> <p>If you enabled Configure Authorization Settings, provide either a VLAN Name or a VLAN Pool Name.</p> <p>If you enabled Select Existing Authorization, click Select, indicate one of the listed profiles, and then click OK. The Authorization Profile name appears in the Authorization Profile field in the Create WLAN Service Profile for Radio Profile window.</p> <p>6. Click OK.</p> <p>The selected WLAN Service profile appears on the list of WLAN Service Profiles on the Basic Setup tab.</p> <p>7. Click Done or, to reconfigure default advanced settings, click Advanced Setup.</p>

Next, assign this Radio profile to radios. For directions, see [“Assigning a Radio Profile to Radios” on page 952](#).

NOTE: Assigned settings from any profile, including this one, have lower priority than settings made directly to a controller or an access point. For more information, see [“Adding and Managing an Individual Access Point” on page 1156](#) and [“Configuring a Controller” on page 1037](#).

Specifying Custom Radio Profile Setup Settings

IN THIS SECTION

- [Basic Settings for Custom Radio Profiles | 940](#)
- [RF Scanning Settings for Custom Radio Profiles | 944](#)
- [Auto Tune Settings for Custom Radio Profiles | 945](#)
- [Voice Configuration Settings for Custom Radio Profiles | 948](#)
- [802.11 Attributes Settings for Custom Radio Profiles | 949](#)
- [Adaptive Channel Planning Settings for Custom Radio Profiles | 950](#)
- [Snooping Mapping Settings for Custom Radio Profiles | 951](#)

The **Custom Setup** tab includes seven sets of parameters, listed at the left of the screen: **Basic Settings**, **RF Scanning**, **Auto Tune**, **Voice Configuration**, **802.11 Attributes**, **Adaptive Channel Planning**, and **Snooping Map**.

Enter the custom Radio profile settings described in these topics. Required settings in the topics are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Basic Settings for Custom Radio Profiles

Enter the Basic Settings for the custom Radio profile described in [Table 200](#).

Table 200: Radio Profile Custom Setup Basic Settings

Field	Action
Name and Description	<p>Values for these attributes are copied from the Quick Setup tab if they were completed there.</p> <p>If attributes were not copied, type a unique name that identifies the profile.</p> <p>Use up to 32 characters for wireless profile names. Profile names must not contain special characters or spaces. Note that profiles automatically created by Network Director as part of device discovery or out-of-band changes may contain the underscore (_) character.</p>
DFS Channel (default is disabled)	Select Enable to enable dynamic frequency selection (DFS) channels to meet regulatory requirements as assigned by the access point country code.
RFID (default is disabled)	Select Enable to enable tracking of mobile assets by using RFID tags.
U-APSD (default is disabled)	<p>Select Enable to enable Unscheduled automatic power save delivery (U-APSD), which buffers unicast packets on access points until clients request transmission. There are two methods for requesting buffered unicast packets from access point radios, either the Wi-Fi Multimedia (WMM) method (default) or the Spectralink Voice Priority (SVP) method.</p> <p>NOTE: Only clients that are using power save mode are affected by this setting.</p>
Countermeasures (default is None)	<p>Select which clients receive countermeasures, None, All or Rogue (only rogue devices). Countermeasures are packets sent by a radio to mitigate interfering devices.</p> <p>NOTE: Configuring the actual countermeasure actions is not supported in Network Director.</p>
Client TX Power Constraint (default is None)	<p>Select Link Balance to balance power between access points and clients on all client links. To accomplish this, access points transmit the maximum power setting to potential clients. You might want to limit client power in places like auditoriums, airports, and dorms where there are a high number of clients.</p>

WLAN Service Profiles

Table 200: Radio Profile Custom Setup Basic Settings (continued)

Field	Action
Enable Weight Queing	If you configure SSID medium time weights, you are guaranteeing a minimum service level to specific Service profiles on a radio. Medium time weights determine the relative transmit utilization of the radio between Service profiles. You can configure the weight from 1 to 100 with 100 as the sum of all configured weights.
Task: Select Existing WLAN Service profiles	Click Select under WLAN Service Profiles and then select any or all listed WLAN Service profiles to associate with this Radio profile. Radios will broadcast the SSIDs named in the WLAN Service profile(s).

Table 200: Radio Profile Custom Setup Basic Settings (continued)

Field	Action
Task: Create and Add a New WLAN Service profile	

Table 200: Radio Profile Custom Setup Basic Settings (continued)

Field	Action
	<ol style="list-style-type: none"> 1. Click Create under WLAN Service Profiles. The Create WLAN Service Profile for Radio Profile window opens. 2. Provide the following settings for a WLAN Profile: <ol style="list-style-type: none"> a. Profile Name—Type a unique name that identifies the profile. Use up to 32 characters for wireless profile names. Profile names must not contain special characters or spaces. Note that profiles automatically created by Network Director as part of device discovery or out-of-band changes may contain the underscore (_) character. b. Description c. Service Profile Type—Select 802.1X, Voice, Web Portal, Open Access or Custom. d. SSID—Provide an SSID to be broadcast to clients. 3. Indicate security settings for the WLAN Profile by selecting either RSN, WPA, or Static WEP. Provide this additional information: <ol style="list-style-type: none"> a. RSN—Select either AES (CCMP) or TKIP. b. WPA—Select either AES (CCMP) or TKIP. Select either 802.1X Authentication or PSK Authentication or both. c. Static WEP—Provide one or more WEP keys. 4. Indicate authentication settings for the WLAN Service profile. Enable either Configure Authentication Settings, or Select Existing Authentication. If you enabled Configure Authentication Settings, select either Create RADIUS Server or Select RADIUS Server. If you selected Create RADIUS Server, provide this information: <ol style="list-style-type: none"> a. RADIUS Server Address b. Secret If you enabled Select Existing Authentication, click Select, indicate one of the listed profiles, and then click OK. The Authentication Profile name appears in the Authentication Profile field in the Create WLAN Service Profile for Radio Profile

Table 200: Radio Profile Custom Setup Basic Settings (continued)

Field	Action
	<p>window.</p> <p>5. Indicate authorization settings for the WLAN Service profile. Select either Configure Authorization Settings, or Select Existing Authorization.</p> <p>If you enabled Configure Authorization Settings, provide either a VLAN Name or a VLAN Pool Name.</p> <p>If you enabled Select Existing Authorization, click Select, indicate one of the listed profiles, and then click OK. The Authorization Profile name appears in the Authorization Profile field in the Create WLAN Service Profile for Radio Profile window.</p> <p>6. Click OK.</p> <p>The selected WLAN Service profile appears on the list of WLAN Service Profiles on the Basic Setup tab.</p> <p>7. Click Done or, to reconfigure default advanced settings, click Advanced Setup.</p>

Next, click the **RF Scanning** tab for custom Radio profiles and follow the directions “[RF Scanning Settings for Custom Radio Profiles](#)” on page 944.

RF Scanning Settings for Custom Radio Profiles

Enter the RF Scanning settings for the custom Radio profile described in [Table 201](#).

Table 201: Radio Profile Custom Setup RF Scanning Settings

Field	Action
Mode (default is active)	Select Passive to make scanning detect only active WLAN devices. (Passive scanning cannot be turned off.) Select Active to make scanning detect all wireless devices, even if they are not currently transmitting.
Channel Scope (default is regulatory)	Select the range of access point radio channels to be scanned and audited— Regulatory (channels required by country of operation), All , or Operating (active channels).
CTS-to-Self (default is disabled)	Select Enable to have access points send clear-to-send packets to themselves on the new channel before switching an access point radio to another channel. This action confirms that both the send and receive are working on the access point before it attempts communication with other devices.
Spectral Scan (default is disabled)	Select Enable to enable access points to scan their coverage area for all electronic devices, not just wireless devices.

Next, click the **Auto Tune** tab for custom Radio profiles and follow the directions [“Auto Tune Settings for Custom Radio Profiles”](#) on page 945.

Auto Tune Settings for Custom Radio Profiles

Power tuning computation is performed on the access point itself without any help from the controller. Access points listen for nearby access points on the same channel and then adjust their power to provide good coverage while avoiding co-channel interference. For more information, see [“Understanding Auto Tune Power Policy for Wireless Radios”](#) on page 866.

Enter the Auto Tune settings for the custom Radio profile described in

Table 202: Radio Profile Custom Setup Auto Tune Settings

Field	Action
Version	Power Auto Tuning parameters differ, depending on which version of MSS you are using. Select either Power & Channel Tuning (Below MSS 8.0 version), Power & Channel Tuning (Below MSS 9.0 version), or Power Policy (MSS 9.0 and above versions).
Auto Tune Settings	
Version (default is Max Coverage) (Power Policy 9+)	<p>Only when you selected the version Power Policy for MSS Version 9.0 and newer, you can base your power policy on Max Coverage, Cell Parity, or Max Channel Capacity. Both Cell Parity and Max Channel Capacity have additional configuration options that become available when you select them.</p> <p>Cell Parity Power (2GHz and 5GHz)—When the Power Policy for MSS Version 9.0 and newer parameter Power Policy is set to Cell Parity, you can set the same power on all radios, based on the radio capability and regulation. You can configure per-band power levels and the system accommodates these levels as allowed by regulatory constraints. For an equally spaced access point deployment, this power policy is better suited as it will not compute transmit power at run time. However, for very dense deployments, this policy may cause co-channel interference. The power values for 2.4-GHz and 5-GHz are different and all radios of the same channel band are set to equal power levels as allowed by the hardware, the channel, and the country code. By default, the power level is the highest value that can be used in common by all radios under the profile.</p>
Tune Channel (default is enabled) (Power & Channel Tuning < 8.0)	Only when you are using MSS versions earlier than 8.0, Adaptive Channel Planner (ACP) dynamically assigns the radio operating channel so that the wireless network can efficiently adapt to the RF environment conditions. Dynamic assignment can be changed when significant changes are measured in the interference level or in the network topology. Eventually, Wi-Fi bandwidth is maximized and maintains the efficiency of communication over the wireless network. Tune Channel is enabled by default, but you can disable it. It is also overwritten if a static channel set is configured. When Tune Channel is disabled, channels on the access points are static and require manual intervention to change the channels.

Table 202: Radio Profile Custom Setup Auto Tune Settings (continued)

Field	Action
<p>Tune Transmit Power (default is disabled)</p> <p>(Power & Channel Tuning < 9.0, Power & Channel Tuning < 8.0)</p>	<p>Click Enable to enable automatic power tuning for radios. Access point radios are initially set to maximum power for the country's regulatory domain. Radios maintain maximum power, which is recommended, unless a conflict or other event causes them to change. For more information, see "Understanding Auto Tune Power Policy for Wireless Radios" on page 866.</p>
<p>Ignore Clients (default is disabled)</p> <p>(Power & Channel Tuning < 8.0)</p>	<p>Only when you are using MSS versions earlier than 8.0, the Ignore Clients option can allow MSS to change the channel on a radio even if the radio has active client sessions. When Ignore Clients is disabled, MSS does not change the channel unless there are no active client sessions on the radio.</p>
<p>Tune Channel Range (default is Lower Bands)</p> <p>(Power & Channel Tuning < 8.0)</p>	<p>Only when you are using MSS versions earlier than 8.0, you can limit the channels used to the Lower Bands, or you can indicate All Bands.</p>
<p>Channel Tuning Interval (default is 3600 seconds)</p> <p>(Power & Channel Tuning < 8.0)</p>	<p>Only when you are using MSS versions earlier than 8.0, you can change the interval between channel tuning opportunities. Every 3600 seconds, MSS examines the RF information gathered from the network and determines whether the channel needs to be changed to compensate for RF changes.</p>
<p>Channel Tuning Holddown (default is 900 seconds)</p> <p>(Power & Channel Tuning < 8.0)</p>	<p>Only when you are using MSS versions earlier than 8.0, you can change the Channel Tuning Holddown. Channel holddown avoids unnecessary channel changes due to highly transient RF changes, such as activation of a microwave oven.</p>
<p>TX Power Backoff Timer (default is 10 seconds)</p> <p>(Power & Channel Tuning < 8.0)</p>	<p>Only when you are using MSS versions earlier than 8.0, you can change the interval at which radios reduce power after temporarily increasing the power to maintain the minimum data rate for an associated client. At the end of each power-backoff interval, radios that temporarily increased their power reduce it by 1 dBm every 10 seconds. The power backoff continues in 1 dBm increments after each interval until the power returns to expected setting.</p>

Table 202: Radio Profile Custom Setup Auto Tune Settings (continued)

Field	Action
<p>Power Tuning Interval (default is 600 seconds)</p> <p>Power Policy 9+ for Max Channel Capacity Power Policy, Power & Channel Tuning < 9.0, Power & Channel Tuning < 8.0</p>	<p>Power Tuning Interval is the parameter for Auto Tune, no matter which version of MSS you are using, that sets the number of seconds between reevaluations of power. Power changes can take place only after an evaluation or when an anomaly occurs. You can change the wait interval between evaluations from the default 600 seconds.</p>
<p>Power Ramp Interval (default is 60 seconds)</p> <p>Power & Channel Tuning < 9.0, Power & Channel Tuning < 8.0</p>	<p>Power Ramp Interval is the rate at which power is increased or decreased on radios in a Radio profile until the optimum power level calculated by RF auto-tuning is reached. You can change the 1 dBm increment to increase and decrease in larger or smaller steps.</p>
<p>Minimum Power</p> <p>Power & Channel Tuning < 9.0, Power Policy 9+ for Max Channel Capacity Power Policy</p>	<p>Select a number from 1 through 24 to indicate the minimum power applied to a radio. The min-power setting places a floor under the power range that radios use when attempting to maintain power parity with each other. For instance, if the highest common power level is limited by a radio with a regulatory or hardware limit of 10 dB, and you set the min-power level to 12 dB, all radios capable of 12 dB are set to use 12dB even though it is higher than the highest common power level. If a minimum power level is configured that is higher than a configured maximum power setting, the minimum power level is rejected. If a minimum power level is configured that is higher than the upper limit of radios used by the Radio profile, the configuration is accepted, but a warning is issued.</p>
<p>Maximum Power</p> <p>(Power Policy 9+ for Max Channel Capacity Power Policy)</p>	<p>Only when you are using MSS 9.0 or newer with Max Channel Capacity as the Power Policy, you can set radios' maximum power to a value that depends on the radio model.</p>
<p>Power Density</p> <p>(Power Policy 9+ for Max Channel Capacity Power Policy)</p>	<p>Only when you are using MSS 9.0 or newer with Max Channel Capacity as the Power Policy, you can set power density to Low, Medium, or High.</p>

Next, click the **Voice** tab for custom Radio profiles and follow the directions “[Voice Configuration Settings for Custom Radio Profiles](#)” on page 948.

Voice Configuration Settings for Custom Radio Profiles

Enter the Voice settings for the custom Radio profile described in [Table 203](#).

TIP: If Voice configuration QoS mode is set to **SVP**, SVP is also given highest priority and the CAC settings still apply.

Table 203: Radio Profile Custom Setup Voice Settings

Field	Action
QoS Mode (default is WMM)	<p>Sending (WMM) and Receiving (Policing) Packets provides levels of service that vary with the type of data sent on an access point. By default, access points deliver the same service to all packets—you must configure admission control for CAC to work properly. For more information, see “Understanding Call Admission Control” on page 1068.</p> <p>Select WMM (default) quality-of-service mode to prioritize VoIP traffic, unless you are using Spectralink phones. In that case only, select SVP. In either case WMM CAC configurations apply to traffic. The only difference is that SVP also gets highest priority when SVP is selected. .</p> <p>NOTE: When SVP is selected instead of WMM, SVP voice transmission also has highest priority in addition to voice and video. The rest of the WMM QoS settings still apply. SVP is needed only for legacy equipment.</p>

Background

Background has the longest fixed wait time before forwarding queued packets, giving it lowest priority.

ACM Mode	Enable ACM Mode to enable background admission control mode for downstream traffic.
ACM Limit	Indicate a value from 1 through 100 for the ACM Limit to limit downstream background mode to a percentage of bandwidth.
ACM Policing	Enable ACM Policing to enable background mode for upstream traffic.

Best Effort

Best effort provides QoS by providing more resources than are needed by the network. This simple method provides high quality service to all packets on an IP backbone and is therefore frequently used on the network core.

Table 203: Radio Profile Custom Setup Voice Settings (continued)

Field	Action
	<p>Enable ACM Mode to enable best-effort admission control mode for downstream traffic.</p> <p>Indicate a value from 1 through 100 for the ACM Limit to limit downstream best-effort mode to a percentage of bandwidth.</p> <p>Enable ACM Policing to enable best-effort mode for upstream traffic.</p>

Voice

Voice transmission has the shortest fixed wait time before forwarding queued packets, giving it highest priority.

ACM Mode	Enable ACM Mode to enable voice admission control mode for downstream traffic.
ACM Limit	Indicate a value from 1 through 100 for the ACM Limit to limit downstream voice mode to a percentage of bandwidth.
ACM Policing	Enable ACM Policing to enable voice mode for upstream traffic.

Video

Video streaming has the shortest fixed wait time before forwarding queued packets, giving it highest priority.

Video 2 ACM Mode	Select Video 2 ACM Mode to enable video admission control mode for downstream traffic.
Video 2 ACM Limit	Type a value from 1 through 100 for the Video 2 ACM Limit to limit downstream video mode to a percentage of bandwidth.
Video 2 ACM Policing	Select Video 2 ACM Policing to enable video mode for upstream traffic.

Next, click the **802.11 Attributes** tab for custom Radio profiles and follow the directions [“802.11 Attributes Settings for Custom Radio Profiles”](#) on page 949.

802.11 Attributes Settings for Custom Radio Profiles

Enter the 802.11 Attributes settings for the custom Radio profile described in

Table 204: Radio Profile Custom Setup 802.11 Attributes Settings

Field	Action
802.11 Attributes	
Beacon Interval	Indicate the rate at which a radio advertises beacons SSID(s). The interval can be a value from 25 ms through 8191 ms. The default is 100 ms.

Table 204: Radio Profile Custom Setup 802.11 Attributes Settings (continued)

Field	Action
DTIM Interval	Number of times after every beacon that a radio sends a delivery traffic indication map (DTIM). An WLA access point sends the multicast and broadcast frames stored in its buffers to clients who request them in response to the DTIM. The DTIM interval applies to both the beacons and the unbeacons.
Maximum Receive Lifetime	Indicate the number of milliseconds that a frame received by a radio can remain in buffer memory. The time can be from 500 ms (0.5 second) through 250,000 ms (250 seconds). The default is 2000 ms (2 seconds).
Maximum Transmit Lifetime	Indicate the number of milliseconds a frame scheduled to be transmitted by a radio can remain in buffer memory. The time can be from 500 ms (0.5 second) through 250,000 ms (250 seconds). The default is 2000 ms (2 seconds).
Long Preamble Length	Indicate a long preamble length. An 802.11b/g radio generates unicast frames to send to a client with the specified preamble length. An 802.11b/g radio always uses a long preamble in beacons, probe responses, and other broadcast or multicast traffic.
Rate Enforcement	When data rate enforcement is enabled, clients transmitting at the disabled rates cannot associate with the access points.
RST Threshold	Indicate the maximum length a frame can be before a radio uses the Request-to-Send/Clear-to-Send (RTS/CTS) method to send the frame. The RTS/CTS method clears the air of other traffic to avoid corruption of the frame due to a collision with another frame. Default is 65535.
Fragment Threshold	Specify the longest length a frame can be before a radio transmits it by fragmenting it into multiple frames. The threshold can be a value from 256 through 2346. The default is 2346.
802.11n Attributes	
802.11n Channel Width	Indicate the width of an 802.11n channel—either 20 MHz or 40 MHz.

Next, click the **Adaptive Channel Planning** tab for custom Radio profiles and follow the directions “[Adaptive Channel Planning Settings for Custom Radio Profiles](#)” on page 950.

Adaptive Channel Planning Settings for Custom Radio Profiles

Adaptive Channel Planner makes channel tuning decisions based on feedback from RF scanning—for more information, see “[Understanding Adaptive Channel Planner](#)” on page 861 and “[Understanding Wireless Scanning](#)” on page 869.

Enter the Adaptive Channel Planning settings for the custom Radio profile described in

Table 205: Radio Profile Custom Setup Adaptive Channel Planning Settings

Field	Action
Task: Configuring 802.11b/g Channels	By default, 802.11bg radios use channels 1, 6, and 11 because these common settings prevent interference. To change the channels used, select a number and move it into the Available or Selected columns with the arrows. For more information about wireless channels, see “Understanding Wireless Radio Channels” on page 856 .
Task: Configuring 802.11a Channels	<p>You can configure the 802.11a radio on an access point to allow certain channels to be available or unavailable. If you select lower bands, MSS selects a channel from the lower eight bands in the 802.11a range of channels: 36, 40, 44, 48, 52, 56, 60, or 64. If you select all-bands, MSS selects a channel from the entire 802.11a range of channels: 36, 40, 44, 48, 52, 60, 64, 149, 153, 157, or 161.</p> <p>NOTE: When a controller learns that an access point 802.11a radio has detected radar on a channel, the auto-tune module immediately switches the radio to another channel. If radar is detected on a radio that has its auto-tune channel feature disabled, the radio goes out of service as required by the Federal Communications Commission (FCC) and European Telecommunications Standards Institute (ETSI).</p>

Next, click the **Snoop Map** tab for custom Radio profiles and follow the directions .[“Snooping Mapping Settings for Custom Radio Profiles” on page 951](#).

Snooping Mapping Settings for Custom Radio Profiles

Enter the Snooping Mapping settings for the custom Radio profile described in [Table 206](#).

Table 206: Radio Profile Custom Setup Snoop Map Settings

Field	Action
Task: Map Snooping to a Controller	<p>You must have an existing Snooping profile before you can map it to a Radio profile. To create a Snooping profile, see “Creating and Managing RF Snooping Filter Profiles” on page 1125.</p> <p>To map snooping to a Radio profile, click Select on the Snooping option of the Advanced Setup tab, select a profile from the list, and then click OK.</p>

What To Do Next

Next, assign this Radio profile to radios. For directions, see [“Assigning a Radio Profile to Radios” on page 952](#).

NOTE: Assigned settings from any profile, including this one, have lower priority than settings made directly to a controller or an access point. For more information, see [“Adding and Managing an Individual Access Point”](#) on page 1156 and [“Configuring a Controller”](#) on page 1037.

RELATED DOCUMENTATION

[Understanding the Network Director User Interface | 85](#)

[Creating and Managing a WLAN Service Profile | 1090](#)

[Understanding WMM Power Save and WLAN Client Battery Life | 859](#)

[Understanding Adaptive Channel Planner | 861](#)

[Understanding Auto Tune Power Policy for Wireless Radios | 866](#)

[Understanding the IEEE 802.11 Standard for Wireless Networks | 1076](#)

[Deploying Configuration to Devices | 1180](#)

[Network Director Documentation home page](#)

Assigning a Radio Profile to Radios

IN THIS SECTION

- [Selecting the Controllers for Radio Profile Assignment | 954](#)
- [Assigning the Radio Profile to Radios | 954](#)
- [Editing Radio Profile Assignments | 956](#)
- [What To Do Next | 956](#)


To operate, configured Radio profiles must be assigned to a controller, and then the controller must be deployed. This topic describes assigning a Radio profile to a controller.

NOTE: You must have a previously configured Radio profile to assign—if you do not, stop now and follow the directions [“Creating and Managing a Radio Profile”](#) on page 932.

The following sections describe how to use the Assign Radio Profile wizard and the Edit Assignments page.

NOTE: If a radio already has a profile assigned and you assign another profile to its controller, then the new Radio profile overrides the existing Radio profile on the radio.

To assign a Radio profile to controllers:

1. In the View pane, expand the **Wireless Network** and then select a controller or a controller cluster. If you select a controller cluster, all controllers in that cluster are selected. If you select a Custom Group, all members of that Custom Group are selected.
2. Click  in the Network Director banner.
3. Under Views, select one of these options: **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Datacenter View**, or **Topology View**.

4. In the Tasks pane, expand **Wireless**, expand **Profiles**, and then click **Radio**.

The Manage Radio Profiles page opens, displaying a list of all existing Radio profile assignments and options to add, edit, or delete Radio profiles.

5. Select the Radio profile you wish to assign and click **Assign**.

The Assign Radio Profile wizard guides you through three parts—Object Selection (controller selection), Radio Profile Assignment, and Review. You can either click the help icon (?) for information about these steps or read the next three sections in this topic.

NOTE: After you assign a Radio profile to controllers, you can modify your assignments by selecting the Radio profile from the Manage Radio Profiles page and clicking **Edit Assignments**.

The following sections describe how to use the Assign Radio Profile wizard and the Edit Assignments page.

Selecting the Controllers for Radio Profile Assignment

The Assign Radio Profile wizard guides you through Object Selection, the selection of controllers to be assigned this Radio profile.

Select wireless controllers from the network tree:

1. Select controllers from the list by placing a check mark in the corresponding box. If you select a group or cluster, all members are selected.

NOTE: Simply highlighting a controller will not work, a check mark must appear in the corresponding box or you will get a message in the next step that says "Please select at least one object type."

2. Click either **Profile Assignment** or **Next** when you have finished selecting the devices.

The next step of the wizard, Profile Assignment, appears.

Assigning the Radio Profile to Radios

Use the Profile Assignment step of the Assign Radio Profile wizard to select specific radios on the selected controller for assignment.

TIP: Any profile assignments or attributes you define during this procedure will replace the existing ones. Existing profile assignments are not shown.

To assign a Radio profile to controllers:

1. From the Assignments list, select a controller or a controller cluster by placing a check mark in the corresponding box.

NOTE: If Network Director fails to read the configuration of one or more devices after the device discovery, those devices are not displayed in the Assign Profile page and you will not be able to assign profiles to the devices. The Manage Jobs page in System mode displays details of the device discovery jobs. Use the information displayed on this page to take appropriate corrective steps to enable Network Director to reread the configuration of the failed device. For more information, see [“Discovering Devices in a Physical Network” on page 204](#).

After you assign a Radio profile to controllers, you can still modify your assignments by selecting the Radio profile from the Manage Radio Profiles page and clicking **Edit Assignments**.

2. Click **Assign > Assign to Radio**.

A list of access point radios on the selected controller is displayed

3. Select one or more radios from the list of radios, and then click **Assign**.

Under assignments, the controller, the access points, and the radios that are assigned this Radio profile are listed.

4. Click **Review** or **Next**.

The next step of the wizard, Review, appears.

5. Look over the Radio profile assignment. If it is correct, click **Finish**. If you need to make a change, click **Edit**, make the changes, and then click **Finish**.

After you click Finish, the Create Profile Assignments Job Details window opens, which reports on the status of the profile assignment job. If you have assigned the profile to a large number of objects, the profile assignment job can take some time to complete. Instead of waiting for the Job Details window to report job completion status, you can close it and check the details of the profile assignment job at a later time using the Manage Job task in System mode.

NOTE: If any assignment fails, the profile assignment job fails and none of the assignments are created. Check the details for the profile assignment job for information about why the assignment failed.

You must deploy the controller configuration for it to take effect. For directions to deploy the controller's configuration, see [“Deploying Configuration to Devices” on page 1180](#).

Editing Radio Profile Assignments

Radio profiles must be assigned to radios for them to function. You can, however, change these assignments to additional or different radios, even while the radios are operating.

To change a radio's Radio profile:

1. Select a Radio profile from the list Manage Radio Profiles.
2. Click **Edit**.

A list of the Radio profile's assignments is displayed.

3. Make any needed changes, and then click **Finish**.

After you click Finish, the Create Profile Assignments Job Details window opens, which reports on the status of the profile assignment job. If you have assigned the profile to a large number of objects, the profile assignment job can take some time to complete. Instead of waiting for the Job Details window to report job completion status, you can close it and check the details of the profile assignment job at a later time using the Manage Job task in System mode.

NOTE: If any assignment fails, the profile assignment job fails and none of the assignments are created. Check the details for the profile assignment job for information about why the assignment failed.

What To Do Next

Next, deploy the devices with assigned Radio profiles—see [“Deploying Configuration to Devices” on page 1180](#).

RELATED DOCUMENTATION

[Creating and Managing a Radio Profile | 932](#)

[Deploying Configuration to Devices | 1180](#)

[Understanding Radio Profiles | 879](#)

[Creating and Managing Wireless Auto AP Profiles | 980](#)

[Network Director Documentation home page](#)

Specifying Custom Radio Profile Setup Settings

IN THIS SECTION

- [Basic Settings for Custom Radio Profiles | 957](#)
- [RF Scanning Settings for Custom Radio Profiles | 961](#)
- [Auto Tune Settings for Custom Radio Profiles | 962](#)
- [Voice Configuration Settings for Custom Radio Profiles | 965](#)
- [802.11 Attributes Settings for Custom Radio Profiles | 967](#)
- [Adaptive Channel Planning Settings for Custom Radio Profiles | 968](#)
- [Snooping Mapping Settings for Custom Radio Profiles | 969](#)
- [What To Do Next | 969](#)

The **Custom Setup** tab includes seven sets of parameters, listed at the left of the screen: **Basic Settings**, **RF Scanning**, **Auto Tune**, **Voice Configuration**, **802.11 Attributes**, **Adaptive Channel Planning**, and **Snooping Map**.

Enter the custom Radio profile settings described in these topics. Required settings in the topics are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Basic Settings for Custom Radio Profiles

Enter the Basic Settings for the custom Radio profile described in [Table 200](#).

Table 207: Radio Profile Custom Setup Basic Settings

Field	Action
Name and Description	<p>Values for these attributes are copied from the Quick Setup tab if they were completed there.</p> <p>If attributes were not copied, type a unique name that identifies the profile.</p> <p>Use up to 32 characters for wireless profile names. Profile names must not contain special characters or spaces. Note that profiles automatically created by Network Director as part of device discovery or out-of-band changes may contain the underscore (_) character.</p>
DFS Channel (default is disabled)	Select Enable to enable dynamic frequency selection (DFS) channels to meet regulatory requirements as assigned by the access point country code.

Table 207: Radio Profile Custom Setup Basic Settings (continued)

Field	Action
RFID (default is disabled)	Select Enable to enable tracking of mobile assets by using RFID tags.
U-APSD (default is disabled)	Select Enable to enable Unscheduled automatic power save delivery (U-APSD), which buffers unicast packets on access points until clients request transmission. There are two methods for requesting buffered unicast packets from access point radios, either the Wi-Fi Multimedia (WMM) method (default) or the Spectralink Voice Priority (SVP) method. NOTE: Only clients that are using power save mode are affected by this setting.
Countermeasures (default is None)	Select which clients receive countermeasures, None , All or Rogue (only rogue devices). Countermeasures are packets sent by a radio to mitigate interfering devices. NOTE: Configuring the actual countermeasure actions is not supported in Network Director.
Client TX Power Constraint (default is None)	Select Link Balance to balance power between access points and clients on all client links. To accomplish this, access points transmit the maximum power setting to potential clients. You might want to limit client power in places like auditoriums, airports, and dorms where there are a high number of clients.
WLAN Service Profiles	
Enable Weight Queing	If you configure SSID medium time weights, you are guaranteeing a minimum service level to specific Service profiles on a radio. Medium time weights determine the relative transmit utilization of the radio between Service profiles. You can configure the weight from 1 to 100 with 100 as the sum of all configured weights.
Task: Select Existing WLAN Service profiles	Click Select under WLAN Service Profiles and then select any or all listed WLAN Service profiles to associate with this Radio profile. Radios will broadcast the SSIDs named in the WLAN Service profile(s).

Table 207: Radio Profile Custom Setup Basic Settings (continued)

Field	Action
Task: Create and Add a New WLAN Service profile	

Table 207: Radio Profile Custom Setup Basic Settings (continued)

Field	Action
	<p>1. Click Create under WLAN Service Profiles.</p> <p>The Create WLAN Service Profile for Radio Profile window opens.</p> <p>2. Provide the following settings for a WLAN Profile:</p> <ol style="list-style-type: none"> Profile Name—Type a unique name that identifies the profile. Use up to 32 characters for wireless profile names. Profile names must not contain special characters or spaces. Note that profiles automatically created by Network Director as part of device discovery or out-of-band changes may contain the underscore (_) character. Description Service Profile Type—Select 802.1X, Voice, Web Portal, Open Access or Custom. SSID—Provide an SSID to be broadcast to clients. <p>3. Indicate security settings for the WLAN Profile by selecting either RSN, WPA, or Static WEP. Provide this additional information:</p> <ol style="list-style-type: none"> RSN—Select either AES (CCMP) or TKIP. WPA—Select either AES (CCMP) or TKIP. Select either 802.1X Authentication or PSK Authentication or both. Static WEP—Provide one or more WEP keys. <p>4. Indicate authentication settings for the WLAN Service profile. Enable either Configure Authentication Settings, or Select Existing Authentication.</p> <p>If you enabled Configure Authentication Settings, select either Create RADIUS Server or Select RADIUS Server.</p> <p>If you selected Create RADIUS Server, provide this information:</p> <ol style="list-style-type: none"> RADIUS Server Address Secret <p>If you enabled Select Existing Authentication, click Select, indicate one of the listed profiles, and then click OK. The Authentication Profile name appears in the Authentication Profile field in the Create WLAN Service Profile for Radio Profile</p>

Table 207: Radio Profile Custom Setup Basic Settings (continued)

Field	Action
	<p>window.</p> <p>5. Indicate authorization settings for the WLAN Service profile. Select either Configure Authorization Settings, or Select Existing Authorization.</p> <p>If you enabled Configure Authorization Settings, provide either a VLAN Name or a VLAN Pool Name.</p> <p>If you enabled Select Existing Authorization, click Select, indicate one of the listed profiles, and then click OK. The Authorization Profile name appears in the Authorization Profile field in the Create WLAN Service Profile for Radio Profile window.</p> <p>6. Click OK.</p> <p>The selected WLAN Service profile appears on the list of WLAN Service Profiles on the Basic Setup tab.</p> <p>7. Click Done or, to reconfigure default advanced settings, click Advanced Setup.</p>

Next, click the **RF Scanning** tab for custom Radio profiles and follow the directions “[RF Scanning Settings for Custom Radio Profiles](#)” on page 944.

RF Scanning Settings for Custom Radio Profiles

Enter the RF Scanning settings for the custom Radio profile described in [Table 201](#).

Table 208: Radio Profile Custom Setup RF Scanning Settings

Field	Action
Mode (default is active)	Select Passive to make scanning detect only active WLAN devices. (Passive scanning cannot be turned off.) Select Active to make scanning detect all wireless devices, even if they are not currently transmitting.
Channel Scope (default is regulatory)	Select the range of access point radio channels to be scanned and audited— Regulatory (channels required by country of operation), All , or Operating (active channels).
CTS-to-Self (default is disabled)	Select Enable to have access points send clear-to-send packets to themselves on the new channel before switching an access point radio to another channel. This action confirms that both the send and receive are working on the access point before it attempts communication with other devices.

Table 208: Radio Profile Custom Setup RF Scanning Settings (*continued*)

Field	Action
Spectral Scan (default is disabled)	Select Enable to enable access points to scan their coverage area for all electronic devices, not just wireless devices.

Next, click the **Auto Tune** tab for custom Radio profiles and follow the directions [“Auto Tune Settings for Custom Radio Profiles”](#) on page 945.

Auto Tune Settings for Custom Radio Profiles

Power tuning computation is performed on the access point itself without any help from the controller. Access points listen for nearby access points on the same channel and then adjust their power to provide good coverage while avoiding co-channel interference. For more information, see [“Understanding Auto Tune Power Policy for Wireless Radios”](#) on page 866.

Enter the Auto Tune settings for the custom Radio profile described in

Table 209: Radio Profile Custom Setup Auto Tune Settings

Field	Action
Version	Power Auto Tuning parameters differ, depending on which version of MSS you are using. Select either Power & Channel Tuning (Below MSS 8.0 version), Power & Channel Tuning (Below MSS 9.0 version), or Power Policy (MSS 9.0 and above versions).
Auto Tune Settings	
Version (default is Max Coverage) (Power Policy 9+)	<p>Only when you selected the version Power Policy for MSS Version 9.0 and newer, you can base your power policy on Max Coverage, Cell Parity, or Max Channel Capacity. Both Cell Parity and Max Channel Capacity have additional configuration options that become available when you select them.</p> <hr/> <p>Cell Parity Power (2GHz and 5GHz)—When the Power Policy for MSS Version 9.0 and newer parameter Power Policy is set to Cell Parity, you can set the same power on all radios, based on the radio capability and regulation. You can configure per-band power levels and the system accommodates these levels as allowed by regulatory constraints. For an equally spaced access point deployment, this power policy is better suited as it will not compute transmit power at run time. However, for very dense deployments, this policy may cause co-channel interference. The power values for 2.4 GHz and 5 GHz are different and all radios of the same channel band are set to equal power levels as allowed by the hardware, the channel, and the country code. By default, the power level is the highest value that can be used in common by all radios under the profile.</p>

Table 209: Radio Profile Custom Setup Auto Tune Settings (continued)

Field	Action
<p>Tune Channel (default is enabled)</p> <p>(Power & Channel Tuning < 8.0)</p>	<p>Only when you are using MSS versions earlier than 8.0, Adaptive Channel Planner (ACP) dynamically assigns the radio operating channel so that the wireless network can efficiently adapt to the RF environment conditions. Dynamic assignment can be changed when significant changes are measured in the interference level or in the network topology. Eventually, Wi-Fi bandwidth is maximized and maintains the efficiency of communication over the wireless network. Tune Channel is enabled by default, but you can disable it. It is also overwritten if a static channel set is configured. When Tune Channel is disabled, channels on the access points are static and require manual intervention to change the channels.</p>
<p>Tune Transmit Power (default is disabled)</p> <p>(Power & Channel Tuning < 9.0, Power & Channel Tuning < 8.0)</p>	<p>Click Enable to enable automatic power tuning for radios. Access point radios are initially set to maximum power for the country's regulatory domain. Radios maintain maximum power, which is recommended, unless a conflict or other event causes them to change. For more information, see "Understanding Auto Tune Power Policy for Wireless Radios" on page 866.</p>
<p>Ignore Clients (default is disabled)</p> <p>(Power & Channel Tuning < 8.0)</p>	<p>Only when you are using MSS versions earlier than 8.0, the Ignore Clients option can allow MSS to change the channel on a radio even if the radio has active client sessions. When Ignore Clients is disabled, MSS does not change the channel unless there are no active client sessions on the radio.</p>
<p>Tune Channel Range (default is Lower Bands)</p> <p>(Power & Channel Tuning < 8.0)</p>	<p>Only when you are using MSS versions earlier than 8.0, you can limit the channels used to the Lower Bands, or you can indicate All Bands.</p>
<p>Channel Tuning Interval (default is 3600 seconds)</p> <p>(Power & Channel Tuning < 8.0)</p>	<p>Only when you are using MSS versions earlier than 8.0, you can change the interval between channel tuning opportunities. Every 3600 seconds, MSS examines the RF information gathered from the network and determines whether the channel needs to be changed to compensate for RF changes.</p>

Table 209: Radio Profile Custom Setup Auto Tune Settings (continued)

Field	Action
<p>Channel Tuning Holddown (default is 900 seconds)</p> <p>(Power & Channel Tuning < 8.0)</p>	<p>Only when you are using MSS versions earlier than 8.0, you can change the Channel Tuning Holddown. Channel holddown avoids unnecessary channel changes due to highly transient RF changes, such as activation of a microwave oven.</p>
<p>TX Power Backoff Timer (default is 10 seconds)</p> <p>(Power & Channel Tuning < 8.0)</p>	<p>Only when you are using MSS versions earlier than 8.0, you can change the interval at which radios reduce power after temporarily increasing the power to maintain the minimum data rate for an associated client. At the end of each power-backoff interval, radios that temporarily increased their power reduce it by 1 dBm every 10 seconds. The power backoff continues in 1 dBm increments after each interval until the power returns to expected setting.</p>
<p>Power Tuning Interval (default is 600 seconds)</p> <p>Power Policy 9+ for Max Channel Capacity Power Policy, Power & Channel Tuning < 9.0, Power & Channel Tuning < 8.0</p>	<p>Power Tuning Interval is the parameter for Auto Tune, no matter which version of MSS you are using, that sets the number of seconds between reevaluations of power. Power changes can take place only after an evaluation or when an anomaly occurs. You can change the wait interval between evaluations from the default 600 seconds.</p>
<p>Power Ramp Interval (default is 60 seconds)</p> <p>Power & Channel Tuning < 9.0, Power & Channel Tuning < 8.0</p>	<p>Power Ramp Interval is the rate at which power is increased or decreased on radios in a Radio profile until the optimum power level calculated by RF auto-tuning is reached. You can change the 1 dBm increment to increase and decrease in larger or smaller steps.</p>

Table 209: Radio Profile Custom Setup Auto Tune Settings (continued)

Field	Action
Minimum Power Power & Channel Tuning < 9.0, Power Policy 9+ for Max Channel Capacity Power Policy	Select a number from 1 through 24 to indicate the minimum power applied to a radio. The min-power setting places a floor under the power range that radios use when attempting to maintain power parity with each other. For instance, if the highest common power level is limited by a radio with a regulatory or hardware limit of 10 dB, and you set the min-power level to 12 dB, all radios capable of 12 dB are set to use 12dB even though it is higher than the highest common power level. If a minimum power level is configured that is higher than a configured maximum power setting, the minimum power level is rejected. If a minimum power level is configured that is higher than the upper limit of radios used by the Radio profile, the configuration is accepted, but a warning is issued.
Maximum Power (Power Policy 9+ for Max Channel Capacity Power Policy)	Only when you are using MSS 9.0 or newer with Max Channel Capacity as the Power Policy, you can set radios' maximum power to a value that depends on the radio model.
Power Density (Power Policy 9+ for Max Channel Capacity Power Policy)	Only when you are using MSS 9.0 or newer with Max Channel Capacity as the Power Policy, you can set power density to Low , Medium , or High .

Next, click the **Voice** tab for custom Radio profiles and follow the directions [“Voice Configuration Settings for Custom Radio Profiles”](#) on page 948.

Voice Configuration Settings for Custom Radio Profiles

Enter the Voice settings for the custom Radio profile described in [Table 203](#).

TIP: If Voice configuration QoS mode is set to **SVP**, SVP is also given highest priority and the CAC settings still apply.

Table 210: Radio Profile Custom Setup Voice Settings

Field	Action
QoS Mode (default is WMM)	<p>Sending (WMM) and Receiving (Policing) Packets provides levels of service that vary with the type of data sent on an access point. By default, access points deliver the same service to all packets—you must configure admission control for CAC to work properly. For more information, see “Understanding Call Admission Control” on page 1068.</p> <p>Select WMM (default) quality-of-service mode to prioritize VoIP traffic, unless you are using Spectralink phones. In that case only, select SVP. In either case WMM CAC configurations apply to traffic. The only difference is that SVP also gets highest priority when SVP is selected. .</p> <p>NOTE: When SVP is selected instead of WMM, SVP voice transmission also has highest priority in addition to voice and video. The rest of the WMM QoS settings still apply. SVP is needed only for legacy equipment.</p>
Background	
Background has the longest fixed wait time before forwarding queued packets, giving it lowest priority.	
ACM Mode	Enable ACM Mode to enable background admission control mode for downstream traffic.
ACM Limit	Indicate a value from 1 through 100 for the ACM Limit to limit downstream background mode to a percentage of bandwidth.
ACM Policing	Enable ACM Policing to enable background mode for upstream traffic.
Best Effort	
Best effort provides QoS by providing more resources than are needed by the network. This simple method provides high quality service to all packets on an IP backbone and is therefore frequently used on the network core.	
	<p>Enable ACM Mode to enable best-effort admission control mode for downstream traffic.</p> <p>Indicate a value from 1 through 100 for the ACM Limit to limit downstream best-effort mode to a percentage of bandwidth.</p> <p>Enable ACM Policing to enable best-effort mode for upstream traffic.</p>
Voice	
Voice transmission has the shortest fixed wait time before forwarding queued packets, giving it highest priority.	
ACM Mode	Enable ACM Mode to enable voice admission control mode for downstream traffic.

Table 210: Radio Profile Custom Setup Voice Settings (continued)

Field	Action
ACM Limit	Indicate a value from 1 through 100 for the ACM Limit to limit downstream voice mode to a percentage of bandwidth.
ACM Policing	Enable ACM Policing to enable voice mode for upstream traffic.
Video	
Video streaming has the shortest fixed wait time before forwarding queued packets, giving it highest priority.	
Video 2 ACM Mode	Select Video 2 ACM Mode to enable video admission control mode for downstream traffic.
Video 2 ACM Limit	Type a value from 1 through 100 for the Video 2 ACM Limit to limit downstream video mode to a percentage of bandwidth.
Video 2 ACM Policing	Select Video 2 ACM Policing to enable video mode for upstream traffic.

Next, click the **802.11 Attributes** tab for custom Radio profiles and follow the directions “[802.11 Attributes Settings for Custom Radio Profiles](#)” on page 949.

802.11 Attributes Settings for Custom Radio Profiles

Enter the 802.11 Attributes settings for the custom Radio profile described in

Table 211: Radio Profile Custom Setup 802.11 Attributes Settings

Field	Action
802.11 Attributes	
Beacon Interval	Indicate the rate at which a radio advertises beacons SSID(s). The interval can be a value from 25 ms through 8191 ms. The default is 100 ms.
DTIM Interval	Number of times after every beacon that a radio sends a delivery traffic indication map (DTIM). An WLA access point sends the multicast and broadcast frames stored in its buffers to clients who request them in response to the DTIM. The DTIM interval applies to both the beacons SSID and the unbeacons SSID.
Maximum Receive Lifetime	Indicate the number of milliseconds that a frame received by a radio can remain in buffer memory. The time can be from 500 ms (0.5 second) through 250,000 ms (250 seconds). The default is 2000 ms (2 seconds).

Table 211: Radio Profile Custom Setup 802.11 Attributes Settings (continued)

Field	Action
Maximum Transmit Lifetime	Indicate the number of milliseconds a frame scheduled to be transmitted by a radio can remain in buffer memory. The time can be from 500 ms (0.5 second) through 250,000 ms (250 seconds). The default is 2000 ms (2 seconds).
Long Preamble Length	Indicate a long preamble length. An 802.11b/g radio generates unicast frames to send to a client with the specified preamble length. An 802.11b/g radio always uses a long preamble in beacons, probe responses, and other broadcast or multicast traffic.
Rate Enforcement	When data rate enforcement is enabled, clients transmitting at the disabled rates cannot associate with the access points.
RST Threshold	Indicate the maximum length a frame can be before a radio uses the Request-to-Send/Clear-to-Send (RTS/CTS) method to send the frame. The RTS/CTS method clears the air of other traffic to avoid corruption of the frame due to a collision with another frame. Default is 65535.
Fragment Threshold	Specify the longest length a frame can be before a radio transmits it by fragmenting it into multiple frames. The threshold can be a value from 256 through 2346. The default is 2346.
802.11n Attributes	
802.11n Channel Width	Indicate the width of an 802.11n channel—either 20 MHz or 40 MHz.

Next, click the **Adaptive Channel Planning** tab for custom Radio profiles and follow the directions “[Adaptive Channel Planning Settings for Custom Radio Profiles](#)” on page 950.

Adaptive Channel Planning Settings for Custom Radio Profiles

Adaptive Channel Planner makes channel tuning decisions based on feedback from RF scanning—for more information, see “[Understanding Adaptive Channel Planner](#)” on page 861 and “[Understanding Wireless Scanning](#)” on page 869.

Enter the Adaptive Channel Planning settings for the custom Radio profile described in

Table 212: Radio Profile Custom Setup Adaptive Channel Planning Settings

Field	Action
Task: Configuring 802.11b/g Channels	By default, 802.11bg radios use channels 1, 6, and 11 because these common settings prevent interference. To change the channels used, select a number and move it into the Available or Selected columns with the arrows. For more information about wireless channels, see “Understanding Wireless Radio Channels” on page 856 .
Task: Configuring 802.11a Channels	<p>You can configure the 802.11a radio on an access point to allow certain channels to be available or unavailable. If you select lower bands, MSS selects a channel from the lower eight bands in the 802.11a range of channels: 36, 40, 44, 48, 52, 56, 60, or 64. If you select all-bands, MSS selects a channel from the entire 802.11a range of channels: 36, 40, 44, 48, 52, 60, 64, 149, 153, 157, or 161.</p> <p>NOTE: When a controller learns that an access point 802.11a radio has detected radar on a channel, the auto-tune module immediately switches the radio to another channel. If radar is detected on a radio that has its auto-tune channel feature disabled, the radio goes out of service as required by the Federal Communications Commission (FCC) and European Telecommunications Standards Institute (ETSI).</p>

Next, click the **Snoop Map** tab for custom Radio profiles and follow the directions [.“Snooping Mapping Settings for Custom Radio Profiles” on page 951](#).

Snooping Mapping Settings for Custom Radio Profiles

Enter the Snooping Mapping settings for the custom Radio profile described in [Table 206](#).

Table 213: Radio Profile Custom Setup Snoop Map Settings

Field	Action
Task: Map Snooping to a Controller	<p>You must have an existing Snooping profile before you can map it to a Radio profile. To create a Snooping profile, see “Creating and Managing RF Snooping Filter Profiles” on page 1125.</p> <p>To map snooping to a Radio profile, click Select on the Snooping option of the Advanced Setup tab, select a profile from the list, and then click OK.</p>

What To Do Next

Next, assign this Radio profile to radios. For directions, see [“Assigning a Radio Profile to Radios” on page 952](#).

NOTE: Assigned settings from any profile, including this one, have lower priority than settings made directly to a controller or an access point. For more information, see [“Adding and Managing an Individual Access Point”](#) on page 1156 and [“Configuring a Controller”](#) on page 1037.

RELATED DOCUMENTATION

WLAN Setup

IN THIS SECTION

- [Setting Up Wireless Radios](#) | 971
- [Creating a Radio Profile Using WLAN Setup](#) | 971
- [Assigning a Radio Profile to Radios By Using WLAN Setup](#) | 975
- [What To Do Next](#) | 976

WLAN Setup consists of creating a Radio profile and also assigning that Radio profile to radios. You can also accomplish this by [“Creating and Managing a Radio Profile”](#) on page 932 and then [“Assigning a Radio Profile to Radios”](#) on page 952.


This topic includes the following:

Setting Up Wireless Radios

To set up WLAN radios, use WLAN Setup to create a Radio profile, then assign the Radio profile to radios:

1. Under Views in the Network Director banner, select either **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.
3. In the Tasks pane, expand **Wireless**, expand **Tasks**, and then click **WLAN Setup**.
The WLAN Setup page opens, displaying the WLAN Setup wizard. The wizard consists of three parts, **Create Radio Profile**, **Assignment**, and **Summary**.
4. **Create Radio Profile** is selected by default. Complete the **Create Radio Profile** configuration as described in both the online help and in [“Creating a Radio Profile Using WLAN Setup” on page 971](#).
5. Click the **Assignment** section of the WLAN Setup wizard. Complete the Radio profile assignment as described in both the online help and in [“Assigning a Radio Profile to Radios By Using WLAN Setup” on page 975](#).
6. Click the **Summary** section of the WLAN Setup wizard. Make any needed changes, and then click **Finish**.

Creating a Radio Profile Using WLAN Setup

Complete the configuration described in [Table 214](#).

Table 214: WLAN Setup Radio Profile Configuration

Field	Description
Name	Type a unique name that identifies the profile. Use up to 32 characters for wireless profile names. Profile names must not contain special characters or spaces. Note that profiles automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character.
Description	Type a description for this Radio profile.

Table 214: WLAN Setup Radio Profile Configuration (*continued*)

Field	Description
-------	-------------

WLAN Service Profiles

Radio profiles must include a WLAN Service profile. You can either create a WLAN Service profile or use an existing WLAN Service profile for the Radio profile.

Table 214: WLAN Setup Radio Profile Configuration (continued)

Field	Description
Task: Create a WLAN Service profile for this Radio profile.	

Table 214: WLAN Setup Radio Profile Configuration (continued)

Field	Description
	<ol style="list-style-type: none"> 1. Click Create under WLAN Service Profiles. The Create WLAN Service Profile for Radio Profile window opens. Required fields include a red asterisk. 2. Provide these basic WLAN Profile settings: <ul style="list-style-type: none"> • Profile Name • Description • Service Profile Type: <ul style="list-style-type: none"> • 802.1X Service Profile • Voice Service Profile • Web Portal Service Profile • Open Access Service Profile • Custom Service Profile • SSID 3. Select Security Settings for the WLAN Service profile. The options are: <ul style="list-style-type: none"> • RSN (WPA2) with AES (CCMP) • RSN (WPA2) with TKIP • WPA with AES (CCMP) • WPA with TKIP • 802.1X Authentication • PSK Authentication with a Pre-shared Key. • Static WEP authentication with up to four WEP Keys. 4. Provide Authentication Settings for the WLAN Service profile. You can either Configure Authentication Settings or you can Select an Existing Authentication Profile from a list. To Configure Authentication Settings (default), either Create a Radius Server or Select a Radius Server. 5. Provide Authorization Settings for the WLAN Service profile. You can either Configure Authorization Settings or you can Select an Existing Authorization Profile from a list. To Configure Authorization Settings (default), select either a VLAN or a VLAN Pool. 6. Click OK.

Table 214: WLAN Setup Radio Profile Configuration (continued)

Field	Description
	The Create WLAN Service Profile for Radio Profile window closes and the WLAN Service profile is added to the list of WLAN Service Profiles.
Task: Select a WLAN Service profile for this Radio profile.	<ol style="list-style-type: none"> 1. Click Select under WLAN Service Profiles. The Select WLAN Service Profile for Radio Profile window opens, displaying a list of existing profiles. 2. Select one or more WLAN Service profiles from the list by placing a check mark next to the names. 3. Click OK. The Select WLAN Service Profile for Radio Profile window closes and the WLAN Service profile is added to the list of WLAN Service Profiles for this Radio profile.
Task: Remove a WLAN Service profile from this Radio profile.	Select one or more of the listed WLAN Service profiles and then click Remove .
Auto Tune Setting (default is MSS 8.0)	<p>Select the version of auto-tuning that corresponds to the version of the MSS operating system in use:</p> <ul style="list-style-type: none"> • MSS 9.0 and newer—Power Policy • MSS 8.0—Power & Channel Tuning • MSS 7.7—Power & Channel Tuning

Next, click **Assign Radios** and follow the directions [“Assigning a Radio Profile to Radios By Using WLAN Setup”](#) on page 975.

Assigning a Radio Profile to Radios By Using WLAN Setup

To assign a Radio profile created with WLAN Setup to radios:

1. Select a controller from the displayed tree.
2. Click **Select**.

The Assign Profile to Radios window opens with a list of radios.

3. Add radios to the controller by selecting one or more radios from the tree in the and then clicking **Assign**.

4. Select one or more radios by placing a check mark next to the name of the radio.
5. Click **Assign**.
The Assign Profile to Radios window closes and the assigned radios are added to the list on the Assign Radios page.
6. Click either **Next** or **Review**, make any changes, and then click **Finish**.

What To Do Next

Next, deploy the devices with assigned Radio profiles—see [“Deploying Configuration to Devices” on page 1180](#).

RELATED DOCUMENTATION

[Creating and Managing a Radio Profile | 932](#)

[Assigning a Radio Profile to Radios | 952](#)

[Understanding Radio Profiles | 879](#)

[Network Director Documentation home page](#)

Configuring Wireless Mesh and Bridging

IN THIS SECTION

- [Create a Mesh SSID and Radio Profile for Access Point Portal Radios | 977](#)
- [Create an SSID and Radio Profile for Access Point Mesh Radios | 977](#)
- [Configure the Mesh Access Points | 978](#)
- [Physically Set Up the Mesh Access Points | 979](#)
- [After the Mesh is Set Up | 979](#)
- [Make Any Further Changes to Mesh Access Points From the Switch | 980](#)

A wireless mesh is useful when you have a hard to reach area that needs network coverage. In this case, you can have one or more access point hops provide wireless service to the remote area. For more

information about wireless mesh and bridging, see [“Understanding Wireless Mesh” on page 894](#) and [“Understanding Wireless Bridging” on page 912](#).

Wireless Mesh is not supported by the Network Director 1.0 release.

This topic describes:

Create a Mesh SSID and Radio Profile for Access Point Portal Radios

Create a mesh SSID in a WLAN Service profile to be used by one of two radios on all mesh portal access points. This SSID is used for mesh link communications with the switch and other portals. It is not used for client associations. You must also have a unique Radio profile for mesh services.

1. Configure a WLAN Service profile for mesh services, giving the SSID a name recognizable as a mesh connection. See [“Creating and Managing a WLAN Service Profile” on page 1090](#) for directions. When creating the WLAN Service profile:
 - Enable **Mesh** on the Basic Settings tab. If this mesh will be a bridge, between buildings for example, also enable **Bridging**.
 - Do not associate an Authorization profile on the Basic Settings tab. Mesh does not work with this kind of authorization.
 - Select the security encryption **WPA2** with the **CCMP** cipher on the Security tab. Also select **PSK** authentication on the Security tab and provide a passphrase. The passphrase must be the same one configured on the mesh access points.
2. You must have a unique Radio profile for mesh services. Create a Radio profile, linking the mesh SSID WLAN that you created. For directions, see [“Creating and Managing a Radio Profile” on page 932](#). In the profile, disable auto channel. Do not enable auto-tune power.

You will assign this SSID to the access points once they are configured.

Create an SSID and Radio Profile for Access Point Mesh Radios

Radios that do not serve as portals use a regular WLAN Service profile that you would use on any other access point. You can use an existing WLAN Service profile for these mesh radios or create a new one following the directions [“Creating and Managing a WLAN Service Profile” on page 1090](#).

Mesh radios can also use a Radio profile that you use on other access points. Either use an existing Radio profile for these mesh radios or create a new one following the directions [“Creating and Managing a Radio Profile” on page 932](#). Assign the Radio profile to all radios that are not serving as mesh portals. For directions, see [“Assigning a Radio Profile to Radios” on page 952](#)

Configure the Mesh Access Points

You need at least one access point portal, and can optionally have more portals and mesh access points. The difference between a mesh portal access point and a mesh access point is that a portal dedicates one radio to passing traffic back and forth from the switch. For this reason, you must use dual-radio access points for all mesh portals so that one radio can be used for mesh link communications (using the SSID reserved for this purpose) while the other radio is used for client associations.

Configure the mesh access points while they are connected to the controller—you will untether them after they are configured, then place them in the mesh. Use these CLI commands to configure the mesh access points:

NOTE: For this release, you cannot configure mesh access points in Network Director. Mesh access points and bridging must be configured from the CLI.

LINK TO

To configure mesh access points:

1. Attach the access points to your controller, apply power, and allow the access point to boot as a regular access points.
2. Once the access point has booted, use the following CLI command to enable mesh services on the mesh access points.

```
set ap apnum boot-configuration mesh mode enable
```

3. Use the following CLI command to specify the pre-shared key on each access point—be sure that you used the same key you identified in the mesh SSID:

```
set ap apnum boot-configuration mesh {psk-phrase pass-phrase | psk-raw raw-pass}
```

When a pass-phrase is specified, it is converted into a raw hexadecimal key and stored in the access point boot configuration.

4. The communication link between a access point and a controller is divided into TAPA and CAPWAP packets. The TAPA packets contain control traffic information and the CAPWAP packets contain client data. Use the following CLI command to set the TAPA control channel timeout on the access point.

```
set ap apnum time-out
```

The default timeout is 10 seconds but you should increase the timeout depending on the length of the mesh link. If DFS is enabled, you might want to increase the timeout to 140 seconds to allow the radio to scan channels.

5. Use Network Director to assign the mesh SSID (WLAN Service profile) to one radio on each mesh portal access point. For directions, see [“Assigning a Radio Profile to Radios” on page 952](#).
6. Use Network Director to assign the regular SSID (WLAN Service profile) to the second radio on dual-radio access points and the single radio on single-radio access points. For directions, see [“Assigning a Radio Profile to Radios” on page 952](#).

NOTE: When using external antennas in conjunction with mesh configurations, enable mesh mode before configuring the external antenna. After adding and configuring the external antenna, reboot the access point.

Physically Set Up the Mesh Access Points

Disconnect all of the configured mesh access points from the controller and deploy them in the final location. Be sure a mesh portal access point (not just a mesh access point) is connected to the switch so the radio with the mesh SSID used for communications can connect the mesh to the switch. For a sample illustration of a mesh setup, see [“Understanding Wireless Mesh” on page 894](#).

NOTE: The mesh portal access point must be within radio range of any other mesh portal or mesh access point.

After the Mesh is Set Up

Mesh authentication uses the pre-shared key (PSK) information. If there are multiple mesh portals advertising the mesh SSID, the mesh access point selects the mesh portal with the strongest received signal strength indicator (RSSI) value. Once authentication is complete, the mesh access point searches for a switch using the identical control packet exchanges as do non-mesh access points on the network.

The mesh link is an authenticated encrypted radio link between mesh access points, and once the link is established, the mesh access point does not switch to another mesh portal unless the access point loses contact with the original mesh portal. When a mesh SSID is specified, the regulatory domain of the switch and the power restrictions are copied to the access point flash memory. This prevents the mesh access point from operating outside of regulatory limits after booting and before receiving a complete configuration

from the switch. Consequently, it is important that the regulatory and antenna information specified on the switch reflects the locale where the mesh access point is to be deployed, in order to avoid regulatory violations.

Make Any Further Changes to Mesh Access Points From the Switch

After the mesh access points are installed in a final location, and establish a connection to the mesh portal, you can do any further configuration from the switch CLI.

RELATED DOCUMENTATION

[Understanding Wireless Mesh | 894](#)

[Understanding Wireless Bridging | 912](#)

[Creating and Managing a WLAN Service Profile | 1090](#)

[Creating and Managing a Radio Profile | 932](#)

[Assigning a Radio Profile to Radios | 952](#)

[Network Director Documentation home page](#)

Creating and Managing Wireless Auto AP Profiles

IN THIS SECTION

- [Managing Wireless Auto AP Profiles | 981](#)
- [Creating a Wireless Auto AP Profile | 983](#)
- [Specifying Basic Settings for a Wireless Auto AP Profile | 984](#)
- [Specifying LLDP & Remote AP Settings for a Wireless Auto AP Profile | 985](#)
- [Specifying Advanced Settings for a Wireless Auto AP Profile | 987](#)
- [Reviewing a Wireless Auto AP Profile | 990](#)
- [What To Do Next | 991](#)

An Auto AP profile is a set of configurations applied to access points on a controller that have not been specifically configured. These access points were, instead, discovered by the controller. If you create an Auto AP profile, then assign and deploy the Auto AP profile to a controller, that controller applies the Auto AP profile to any access points that are discovered. If the access point is not seen by the controller at any

point, the Auto AP configured access point is removed from the list of access points—they have no persistent configuration on the controller. For more information about Auto AP profiles, see [“Understanding Auto AP Profiles” on page 883](#).

NOTE: A controller can have only one Auto AP profile assigned.

If you do not want an Auto-AP profile applied to access points, you can add and configure individual access points by [“Adding and Managing an Individual Access Point” on page 1156](#).

This topic describes:

Managing Wireless Auto AP Profiles

From the Manage Auto AP Profiles page, you can:

- Configure one Auto AP profile per controller by selecting a controller in the View pane, and then clicking **Add**. For directions, see [“Creating a Wireless Auto AP Profile” on page 983](#).
- Modify an existing Auto AP profile by selecting it and clicking **Edit**.
- **Assign** an existing Auto AP profile to a controller.
- Edit the assignment of an Auto AP profile to a controller by selecting the profile and clicking **Edit Assignment**.
- View the current assignments for a profile by selecting the profile and clicking **Details**.
- Delete an Auto AP profile by selecting a profile and clicking **Delete**.

TIP: You cannot delete a profile that is in use—that is, assigned to access points.

- Clone a profile by selecting a profile and clicking **Clone**.

[Table 215](#) describes the information provided about an Auto AP profile on the Manage Auto AP Profiles page.

Table 215: Manage Auto AP Profiles Fields

Field	Description
Profile Name	Name given to the Auto AP profile when the profile was created.
Device Family	Wireless (controller)
Description	Any description provided during creation of the Auto AP profile.

Table 215: Manage Auto AP Profiles Fields (continued)

Field	Description
Enable Auto AP	Indicates whether the Auto AP profile is enabled or disabled.
Assignment State	Indicates whether or not the Auto AP profile is deployed to a controller.
Creation Time	Date and time when this profile was created.
Last Updated Time	Date and time when this profile was last modified.
User Name	The username of the user who created or modified the profile.

TIP: Some columns might be hidden—this is configurable. To show or hide fields listed in the Manage Authorization Profiles table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

Creating a Wireless Auto AP Profile

In Network Director, you can automatically configure profiles for access points discovered by a controller. This Auto AP profile becomes the default configuration for all access points that are discovered but not configured on the controller. This way, all access points, even those not configured by you, have a configuration. The Auto AP profile will also act as a template if you convert an Auto AP to a configured, persistent access point.


At minimum, when creating an Auto AP, you must specify a profile name, and also provide the names of an associated VLAN Profile and a Radio profile. You can use default settings for the rest of the Auto AP profile, or you can optionally change or add Auto AP profile configurations such as:

- Radio Type
- Access points' bias for the controller
- Radio mode (enabled, disabled, or sentry) for each radio
- Radio profile for each radio
- LLDP settings
- Remote access point settings for outage duration and connection evaluation
- Blinking
- Data Security
- High latency mode
- Bonjour Profile
- Automatic firmware updates
- Local switching tunnel settings
- Associated VLAN profiles

To create an Auto AP profile:

1. Under Views, select either **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Datacenter View** or **Topology View**.

2. Click  in the Network Director banner.
3. In the Tasks pane, expand **Wireless**, expand **Profiles**, and then click **Auto AP**.

The Manage Auto AP Profiles page opens, displaying the list of currently configured Auto AP profiles.

4. Click **Add**.

The Create Auto AP Profile for Wireless Controllers (WLC) wizard opens. The wizard has four sections, Basic Settings (this one is selected), LLDP & Remote AP Settings, Advanced Settings, and Review.

5. Complete the sections of the Auto AP wizard described in both online help and [“Specifying Basic Settings for a Wireless Auto AP Profile” on page 984](#), [“Specifying LLDP & Remote AP Settings for a Wireless Auto AP Profile” on page 985](#), [“Specifying Advanced Settings for a Wireless Auto AP Profile” on page 987](#), and [“Reviewing a Wireless Auto AP Profile” on page 990](#).

6. Click **Finish**.

The Auto AP profile is added to the list of Auto AP profiles.

Next, assign the Auto AP configuration to one or more controllers, following the directions [“Assigning an Auto AP Profile to Controllers” on page 991](#).

Specifying Basic Settings for a Wireless Auto AP Profile

For an Auto AP in use (either enabled or in sentry mode), the required basic Auto AP profile settings are a name for the profile and a reference to a Radio profile for each radio—for the rest of the settings, you can use the default settings.

The basic settings for the Auto AP profile are described in [Table 216](#). Required settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 216: Wireless Auto AP Basic Settings

Field	Description
Profile Name	The Auto AP profile name. can consist of up to 32 characters with no special characters other than _ - (underscore or dash).
Description	Optionally provide up to 255 characters.
Enable Auto AP (default is enabled)	Clear this check box to disable Auto AP without deleting the Auto AP profile.
Radio Type (default is 802.11g)	Radio type applied to radios with no other specific configuration. Select 802.11a , 802.11b , or 802.11g . For more explanation, see “Understanding the IEEE 802.11 Standard for Wireless Networks” on page 1076 .

Table 216: Wireless Auto AP Basic Settings (continued)

Field	Description
Bias (default is High)	Select High , Low , or Low-Sticky bias. Bias matters only for access points indirectly connected to the controller through an intermediate layer 2 or layer 3 network. An access point always attempts to boot on AP port 1 first, and if a controller is directly attached on AP port 1, the access point boots from there regardless of the bias settings. Set the access point's bias to High or Low for the controllers in this profile. Alternately, set the bias to Low-Sticky to have the access point continue to use the current controller for the active data link even if another controller configured with high bias becomes available. For more explanation, see “Understanding Access Point Bias for Controllers” on page 852 .
Radio Settings	
If an access point has two radios, both Radio 1 and Radio 2 are displayed and can be reconfigured from the defaults.	
Radio Type (default is 802.11g for Radio 1 and 802.11a for Radio 2)	Radio type applied to each radio that has no other specific configuration. Select either 802.11a , 802.11b , or 802.11g .
Radio Mode (default is Disabled)	Radio mode defines the functioning state of the radios in this profile. Select Sentry or Enabled to enable the radios. Select Disabled to disable the radios. TIP: A radio in sentry mode scans for interference but does not transmit user traffic.
Radio Profile	A controller needs a Radio profile if radio mode is set to either Enabled or Sentry. Provide the name of an associated Radio profile. For more explanation, see “Understanding Radio Profiles” on page 879 . To create a Radio profile, see “Creating and Managing a Radio Profile” on page 932 . TIP: To avoid failure during deployment, select a Radio Profile if the Radio Mode is enabled for any of the radios.

Click either **Next** or **LLDP & Remote AP Settings** to continue configuring the Auto AP. Follow the directions [“Specifying LLDP & Remote AP Settings for a Wireless Auto AP Profile” on page 985](#) to complete that section of the wizard.

Specifying LLDP & Remote AP Settings for a Wireless Auto AP Profile

To configure the LLDP & Remote AP Settings for the Auto AP profile, enter the settings described in [Table 217](#).

Table 217: Wireless Auto AP LLDP & Remote AP Settings

Field	Description
LLDP Settings	
<p>Link Layer Discovery Protocol (LLDP) is a Layer 2 protocol that enables a network device to advertise its identity and capabilities on the local network. These attributes are used to discover neighbor devices. The attributes contain type, length, and value descriptions and are referred to as type-length-values (TLVs). Devices that support LLDP use TLVs to receive and send information to neighboring devices. Details such as configuration information, device capabilities, and device identity can be advertised using this protocol.</p>	
LLDP Mode (default is TX)	<p>Select either transmit (TX) or None to specify the LLDP operational mode on access points. Link Layer Discovery Protocol (LLDP) is a link layer protocol used by network devices for advertising their identity, capabilities, and neighbors on a WLAN.</p> <p>NOTE: LLDP and LLDP-MED cannot operate simultaneously on a network. By default, a network device sends only LLDP packets until LLDP-MED packets are received from an endpoint device, after which the network device sends out LLDP-MED packets until it receives LLDP packets.</p>
LLDP-MED Mode (default is enabled)	<p>LLDP-MED, an extension to LLDP that operates between endpoint devices such as IP phones, and network devices, such as switches, is enabled by default. Specifically, it provides support for voice over IP (VoIP) applications and provides additional TLVs for the capabilities discover, network policy, Power over Ethernet (PoE), and inventory management. Options are enabled and disabled.</p> <p>NOTE: LLDP and LLDP-MED cannot operate simultaneously on a network. By default, a network device sends only LLDP packets until LLDP-MED packets are received from an endpoint device, after which the network device sends out LLDP-MED packets until it receives LLDP packets.</p>
Power via MDI (default is disabled)	<p>Select Enable to have LLDP-MED use the medium dependent interface (MDI) power management TLV for power over Ethernet (PoE).</p>
Enable Inventory (default is disabled)	<p>Select Enable Inventory to have LLDP-MED use the Inventory TLV, enabling you to track inventory for network devices, including their manufacturer, software and hardware versions, and serial or asset number. Options for inventory are enabled and disabled if you have enabled LLDP-MED.</p>

Enable Remote AP Settings

Remote AP Settings, enabled by default, apply to access points connected by a WAN link to a central network. If the WAN link becomes unavailable, then the remote sites with access points remain active and continue to provide connectivity to wireless clients, using the settings indicated here.

Table 217: Wireless Auto AP LLDP & Remote AP Settings (*continued*)

Field	Description
Outage Duration (default is 0 hours)	<p>Select the amount of time in hours that remote sites with access points should remain active and continue to provide connectivity to wireless clients if the WAN link becomes unavailable.</p> <p>If Enable Remote AP Settings is selected, outage duration indicates the number of hours that access points remain in outage mode once they lose connection with a controller. The valid times are from 0 to 120 hours (5 days). The default setting, zero, means access points stay in outage mode indefinitely.</p>
Connection Evaluation Period (default is 300 seconds)	<p>Select the amount of time in seconds that a keep-alive interval of pings is sent to detect an active link for remote access points. The default value is 300 seconds and the range is 5 through 86400 seconds.</p> <p>When Enable Remote AP Settings is selected and an outage occurs, a periodic timer sends discovery messages to the primary access manager (PAM) to determine when the controller is available on the network again. This timer, called an evaluation timer, is configurable and can be used as a hold-down timer to confirm detection of the WAN outage and as a mechanism to detect when the connection is restored.</p>

Click either **Next** or **Advanced Settings** to continue configuring the Auto AP. Follow the directions “[Specifying Advanced Settings for a Wireless Auto AP Profile](#)” on page 987 to complete that section of the wizard.

Specifying Advanced Settings for a Wireless Auto AP Profile

Advanced Auto AP settings include broadcast settings, client types, and voice configuration. The only required configuration here is the name of a VLAN profile to use with the Auto AP profile. [Table 218](#) lists the advanced settings and how to configure them.

Table 218: Wireless Auto AP Profile Advanced Settings

Field	Action
Blink (disabled by default)	Enable LED blink mode on access points to make them easy to identify. When blink mode is enabled on supporting models, the health and radio LEDs alternately blink green and amber. When blink mode is enabled on an AP2750, the 802.11a LED blinks on and off.
Data Security (disabled by default)	Enable data security to configure access points with data path encryption.

Table 218: Wireless Auto AP Profile Advanced Settings (*continued*)

Field	Action
High Latency Mode (disabled by default)	Enable high latency for this automatic configuration.
Bonjour Profile	To make the automatic configuration include Bonjour, enable this option.
Firmware Update (enabled by default)	Disable this option if you do not want the controller to update access points with firmware upgrades.
LED Mode (default is Auto)	Select the LED behavior for access points in this profile, Auto , Static , or Off . Auto is standard behavior. Static LEDs do not flash when traffic is on the network but all other behavior is standard. Selecting Off turns LEDs off.
AP Communication Timeout (default is 25 seconds)	Indicate the length of time in seconds that an access point waits for communication from a controller before timing out.

Table 218: Wireless Auto AP Profile Advanced Settings (continued)

Field	Action				
<p>Enable Local Switching (default is enabled)</p>	<p>Local switching causes traffic that was tunneled to a controller to be locally switched by access points. Local switching includes the following additional settings:</p> <table border="1" data-bbox="505 432 1421 1591"> <thead> <tr> <th data-bbox="505 432 967 1591">WLA Local Switch VLAN Profile</th> <th data-bbox="967 432 1421 1591"></th> </tr> </thead> <tbody> <tr> <td data-bbox="505 432 967 1591"></td> <td data-bbox="967 432 1421 1591"> <p>When a VLAN profile is applied to an access point, traffic for that VLAN is locally switched by the access point—that is, traffic directly exits the access point into the VLAN instead of tunneling back to a controller for switching.</p> <p>If Local Switching is enabled, you can apply a VLAN profile to access points to use with local switching:</p> <ol style="list-style-type: none"> 1. Click Add. 2. Select an existing VLAN profile from the list. 3. Click OK. <p>The VLAN profile is now associated with this Auto AP profile and will be used to locally switch traffic between access points. The VLAN profile name, VLAN name, device family, VLAN ID and device description are added to the list of VLAN profiles.</p> <p>If Local Switching is enabled and a VLAN has been added here, you can select a VLAN profile from the list, and then click Clear to remove the VLAN from local switching in this Auto AP profile.</p> </td> </tr> </tbody> </table>	WLA Local Switch VLAN Profile			<p>When a VLAN profile is applied to an access point, traffic for that VLAN is locally switched by the access point—that is, traffic directly exits the access point into the VLAN instead of tunneling back to a controller for switching.</p> <p>If Local Switching is enabled, you can apply a VLAN profile to access points to use with local switching:</p> <ol style="list-style-type: none"> 1. Click Add. 2. Select an existing VLAN profile from the list. 3. Click OK. <p>The VLAN profile is now associated with this Auto AP profile and will be used to locally switch traffic between access points. The VLAN profile name, VLAN name, device family, VLAN ID and device description are added to the list of VLAN profiles.</p> <p>If Local Switching is enabled and a VLAN has been added here, you can select a VLAN profile from the list, and then click Clear to remove the VLAN from local switching in this Auto AP profile.</p>
WLA Local Switch VLAN Profile					
	<p>When a VLAN profile is applied to an access point, traffic for that VLAN is locally switched by the access point—that is, traffic directly exits the access point into the VLAN instead of tunneling back to a controller for switching.</p> <p>If Local Switching is enabled, you can apply a VLAN profile to access points to use with local switching:</p> <ol style="list-style-type: none"> 1. Click Add. 2. Select an existing VLAN profile from the list. 3. Click OK. <p>The VLAN profile is now associated with this Auto AP profile and will be used to locally switch traffic between access points. The VLAN profile name, VLAN name, device family, VLAN ID and device description are added to the list of VLAN profiles.</p> <p>If Local Switching is enabled and a VLAN has been added here, you can select a VLAN profile from the list, and then click Clear to remove the VLAN from local switching in this Auto AP profile.</p>				
<p>Tunnel Affinity (default is 4)</p>					

Table 218: Wireless Auto AP Profile Advanced Settings (continued)

Field	Action
	<p>If Local Switching is enabled, tunnel affinity enables access points with Local Switching enabled to create and terminate client VLAN tunnels—a VLAN is not required on these access points. Affinity is the capacity that an AP has for tunnels. An AP can have affinity for 0 (none) to 10 tunnels, with the default of 4 tunnels. Zero indicates that the access point is not used as a tunnel endpoint.</p>
	<p>AP Tunneling (default is disabled)</p> <p>If Local Switching is enabled, AP Tunneling configures tunneling on the access points in this profile. If a client connects to an access point that has local switching enabled on a VLAN, and the VLAN does not exist in the VLAN profile, then the client connects in overlay mode. This setting enables clients to still connect if the tunnel limits (set by the Tunnel Affinity value) are reached. If a client cannot connect, an appropriate error message is recorded in the event log.</p>

Click either **Next** or **Review** to continue configuring the Auto AP. Follow the directions [“Reviewing a Wireless Auto AP Profile” on page 990](#) to complete that section of the wizard.

Reviewing a Wireless Auto AP Profile

Before you create the Auto AP profile, check the configuration:

1. Review your Auto AP service profile selections by scrolling to the bottom of the screen.
2. Make any needed changes by clicking **Edit** in the appropriate section. Make the changes, and then return to the Review tab.
3. Click **Finish**. The Auto AP profile appears in the Auto AP profile list.

What To Do Next

Next, assign the Auto AP configuration to one or more controllers, following the directions “[Assigning an Auto AP Profile to Controllers](#)” on page 991.

RELATED DOCUMENTATION

[Understanding Auto AP Profiles](#) | 883

[Adding and Managing an Individual Access Point](#) | 1156

[Understanding the IEEE 802.11 Standard for Wireless Networks](#) | 1076

[Understanding Access Point Bias for Controllers](#) | 852

[Understanding Radio Profiles](#) | 879

[Creating and Managing a Radio Profile](#) | 932

[Assigning an Auto AP Profile to Controllers](#) | 991

[Converting Automatically Discovered Access Points to Manually Configured Access Points](#) | 1248

[Network Director Documentation home page](#)

Assigning an Auto AP Profile to Controllers

IN THIS SECTION

- [Controller Selection for Auto AP Assignment](#) | 992
- [Assigning the Auto AP Profile to Controllers](#) | 993
- [Reviewing and Assigning the Auto AP Profile Configuration](#) | 994
- [Editing Auto AP Profile Assignments](#) | 994
- [What To Do Next](#) | 995

For configured Auto AP profiles to take effect, they must be assigned to a controller, and then the controller’s configuration must be deployed. This topic describes assigning an Auto AP profile to either a controller, or a controller cluster. You must have a completed Auto AP profile to complete this task—for information about creating an Auto AP profile, see “[Creating and Managing Wireless Auto AP Profiles](#)” on page 980. For a description of Auto AP profiles and how they work, see “[Understanding Auto AP Profiles](#)” on page 883.

This topic describes how to use the Assign Auto AP Profile wizard and the Edit Assignments page.

NOTE: If a controller already has an Auto AP profile assigned and you assign another Auto AP profile, then the new profile overrides the existing profile on the controller.

To assign an Auto AP profile to a controller:

1. In any view in the View pane, expand the Wireless Network and then select a wireless object. You can select one controller or a group of controllers. The Auto AP profile is applied to the selected controllers, and any unconfigured access point connected to those controllers will use the Auto AP profile.

2. Click  in the Network Director banner.

3. Click **Auto AP** under Profile and Configuration Management in the Tasks pane.
The Manage Auto AP Profiles page opens with a list of existing Auto AP profiles.

4. Select the Auto AP profile you want to assign and click **Assign**.

The Assign Auto AP Profile wizard opens. It consists of three parts, Device Selection (controller selection), Auto AP Profile Assignment, and Review. You can click the help icon (?) to get help on these steps or you can follow the directions [“Controller Selection for Auto AP Assignment” on page 992](#) and [“Assigning an Auto AP Profile to Controllers” on page 991](#).

NOTE: After you assign an Auto AP profile to controllers, you can modify your assignments by selecting the Auto AP profile from the Manage Auto AP Profiles page and clicking **Edit Assignments**.

The following sections describe how to use the Assign Auto AP Profile wizard and the Edit Assignments page.

Controller Selection for Auto AP Assignment

The first step of the Assign Auto AP Profile wizard is selecting the controllers that will assign the Auto AP profile to unconfigured access points. You can assign an Auto AP profile to a single controller or to controllers associated with a domain or cluster.

To select an object from the network tree:

1. Expand the tree and then select a controller from the list by placing a check mark in the corresponding box. Any controllers below the selected object in the hierarchy are automatically selected.

NOTE: Simply highlighting an object does not select the object. A check mark must appear in the corresponding box or you will get a message in the next step that says **Please select at least one object.**

2. Click either **Profile Assignment** or **Next** when you have finished selecting the objects.

The next step of the wizard, Profile Assignment, appears. Follow the directions [“Assigning the Auto AP Profile to Controllers” on page 993](#) to complete Profile Assignment.

Assigning the Auto AP Profile to Controllers

Use the Profile Assignment step of the Assign Auto AP Profile wizard to the controllers for Auto AP profile assignment.

TIP: Controllers can have only one Auto AP. Any Auto AP profile assignments you define for a controller during this procedure will replace the existing ones. Existing profile assignments are not displayed.

To assign an Auto AP profile to one or more controllers:

1. From the Assignments list, select one or more controllers by placing a check mark in the corresponding box of the object containing controllers.

NOTE: If Network Director fails to read the configuration of one or more devices after device discovery, those devices are not displayed in the Assign Profile page and you will not be able to assign profiles to the devices. The Manage Jobs page in System mode displays details of the device discovery jobs. Use the information displayed on this page to take appropriate corrective steps to enable Network Director to reread the configuration of the failed device. For more information, see [“Discovering Devices in a Physical Network” on page 204](#).

2. Click either **Assign to Device** or **Assign to Cluster**, depending on whether you selected a controller or a controller cluster.

The third column, Assigned To, reflects either Device or Cluster.

3. Click **OK** to close the Job Details window.

The Auto AP profile now appears in the Manage Auto AP Profiles list with the status *Pending Deployment*.

4. To delete an existing assignment, select an assignment from the list and then click **Remove**. To see all existing assignments, click **View Assignments**.
5. Click either **Review** or **Next**.

The next step of the wizard, Review, appears. For directions, see [“Reviewing and Assigning the Auto AP Profile Configuration” on page 994](#).

Reviewing and Assigning the Auto AP Profile Configuration

Review the Auto AP profile assignment. If you do not need to make changes, click **Finish**. If you need to make a change, click **Edit**, make the changes, and then click **Finish**.

After you click Finish, the Create Profile Assignments Job Details window opens, which reports on the status of the profile assignment job. If you have assigned the profile to a large number of controllers, the profile assignment job can take some time to complete. Instead of waiting for the Job Details window to report job completion status, you can close it and check the details of the profile assignment job at a later time using the Manage Job task in System mode.

NOTE: If any assignment fails, the profile assignment job fails and none of the assignments are created. Check the details for the profile assignment job for information about why the assignment failed.

Editing Auto AP Profile Assignments

For Auto AP profiles to function, they must be assigned to either controllers or controller clusters. You can, however, change these assignments to additional or different controllers, even while the controllers are operating.

To change a controller’s Auto AP profile:

1. Select an Auto AP profile from the Manage Auto AP profiles list.
2. Click **Edit**.

A list of the Auto AP profile’s assignments is displayed.

3. Make any needed changes, and then click **Finish**.

After you click Finish, the Create Profile Assignments Job Details window opens, which reports on the status of the profile assignment job. If you have assigned the profile to a large number of objects, the profile assignment job can take some time to complete. Instead of waiting for the Job Details window to report job completion status, you can close it and check the details of the profile assignment job at a later time using the Manage Job task in System mode.

NOTE: If any assignment fails, the profile assignment job fails and none of the assignments are created. Check the details for the profile assignment job for information about why the assignment failed.

What To Do Next

Next, deploy the configuration changes you made to the controller or controller cluster—for directions, see [“Deploying Configuration to Devices” on page 1180](#).

RELATED DOCUMENTATION

[Understanding Auto AP Profiles | 883](#)

[Creating and Managing Wireless Auto AP Profiles | 980](#)

[Converting Automatically Discovered Access Points to Manually Configured Access Points | 1248](#)

[Network Director Documentation home page](#)

Understanding Bonjour

IN THIS SECTION

- [Why Would I Use Bonjour? | 996](#)
- [What Is Zero Configuration? | 996](#)
- [How Do I Configure Bonjour on a Juniper Networks Network Using Network Director? | 996](#)

Bonjour is Apple's implementation of Zero-configuration networking (Zeroconf), a group of technologies that includes service discovery, address assignment, and hostname resolution. Bonjour locates devices such as printers, other computers, and the services that those devices offer on a local network by using multicast Domain Name System (mDNS) service records. Bonjour's Zero Configuration (Zeroconf) provides a way to configure and browse for services over IP networks without knowing a service's name or IP address. Using Bonjour, users can connect on a network to share files, printers, and Internet connections without configuring subnet masks or DNS servers.

Why Would I Use Bonjour?

Bonjour is useful when you have a BYOD policy that enables users to use programs such as iTunes, iChat, Adobe Systems Creative Suite 3, Proteus, Adium, Fire, Pidgin, Skype, Vine Server, or TiVo. Bonjour also comes bundled with some third-party applications, such as Adobe's Photoshop CS3 suite, to take advantage of Zeroconf technology. College campuses are a typical user of Bonjour.

What Is Zero Configuration?

Zeroconf Neighborhood Explorer for Windows is the windows version of Bonjour that enables Apple products such as iPhones[®], iPods, iTouch, and Apple TV use on the network. Services such as AirPlay and AirPrint can also be deployed on wireless networks. Zero configuration IP networking enables these users to find printers and share files.

How Do I Configure Bonjour on a Juniper Networks Network Using Network Director?

Since services provided by Bonjour can consume considerable resources, you will want to limit the services to named VLANs. A typical Bonjour configuration sequence is:

- [Creating and Managing mDNS Profiles on page 997](#)
- [Creating and Managing an mDNS VLAN List on page 1002](#)
- [Assigning an mDNS Profile to Devices on page 1001](#)

RELATED DOCUMENTATION

[Creating and Managing mDNS Profiles | 997](#)

[Creating and Managing an mDNS VLAN List | 1002](#)

[Assigning an mDNS Profile to Devices | 1001](#)

[Network Director Documentation home page](#)

Creating and Managing mDNS Profiles

IN THIS SECTION

- [Managing mDNS Profiles | 997](#)
- [Creating an mDNS Profile | 999](#)
- [Specifying Settings for an mDNS Profile | 999](#)
- [What to Do Next | 1001](#)

mDNS Zero configuration IP networking enables users to find printers, network resources, or music sharing on a network. If you are running mDNS on your network, users can instantly find printers, or a friend's network game or music device, and share those files with someone else. These mDNS services are available in Network Director—Apple TV, Internet printer, or Digital Auto Access Protocol (iTunes).

For more information, see [“Understanding Bonjour” on page 995](#).

Managing mDNS Profiles

From the Manage mDNS Profile page, you can:

- Create a new mDNS Profile by clicking **Add**. For directions, see [“Creating an mDNS Profile” on page 999](#).
- Modify an existing mDNS Profile by selecting it and clicking **Edit**.
- Assign an existing mDNS Profile by selecting it and clicking **Assign**. For directions, see [“Assigning an mDNS Profile to Devices” on page 1001](#).
- Edit an existing mDNS Profile assignment by selecting it and clicking **Edit Assign**.
- View information about an mDNS Profile by either clicking the profile name or by selecting the profile and clicking **Details**.
- Delete an mDNS Profile by selecting site name and clicking **Delete**.

TIP: You cannot delete a profile that is in use. To see the current state of a profile, select the site name and click **Details**.

- Clone an existing mDNS Profile by selecting it and clicking **Clone**.

Table 219 describes the information provided about mDNS Profiles on the Manage mDNS Profiles page. This page lists all mDNS Profiles defined for your network, regardless of the scope you selected in the network view.

Table 219: mDNS Profile Information

Field	Description
Profile Name	Name assigned when the profile was created.
Description	Any description added during creation of the mDNS Profile.
Assignment State	Displays the assignment state of the profile. A profile can be: <ul style="list-style-type: none"> • Unassigned—When the profile is not assigned to any object. • Deployed—When the profile is assigned and is deployed from Deploy mode. • Pending Deployment—When the profile is assigned, but not yet deployed in the network.
Creation Time	Date and time when the profile was created.
Last Updated Time	Date and time when the profile was last modified.
User Name	The username of the person who created or modified the profile.


TIP: All columns might not be displayed. To show or hide fields listed in the table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

Creating an mDNS Profile

To create an mDNS Profile to provide Apple TV, Internet printing, and/or iTunes support, follow these steps:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.
3. In the Tasks pane, expand **Wireless**, expand **Profiles**, and then click **mDNS Profile**.
The Manage mDNS Profile page appears, displaying a list of currently configured mDNS Profiles and their attributes.
4. Click **Add**.
The Create mDNS Profile page appears.
5. Provide the configuration indicated in [“Specifying Settings for an mDNS Profile” on page 999](#).
6. Click **Done**.
7. Limit the resources consumed by various mDNS services by [“Creating and Managing an mDNS VLAN List” on page 1002](#).
8. Assign the mDNS Profile following the directions in [“Assigning an mDNS Profile to Devices” on page 1001](#).

Specifying Settings for an mDNS Profile

mDNS provides a method for Apple devices to discover services on a local area network. In Network Director, Apple TV, Internet printing, and/or iTunes support can be added. You should limit use of mDNS to specific VLANs to prevent interference with the rest of the network—for directions to do this, see [“Creating and Managing an mDNS VLAN List” on page 1002](#).

Provide the settings described in [Table 220](#) to create an mDNS Profile. For more information about mDNS, see [“Understanding Bonjour” on page 995](#).

Table 220: Specifying Settings for an mDNS Profile

Field	Description
Profile Name	<p>Type a unique name that identifies the profile.</p> <p>Use up to 32 characters for the name. Profile names must not contain special characters or spaces. Note that profiles automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character.</p>
Description	Type up to 256 characters.
Task: Add a rule to the mDNS Profile	<p>To add a rule to this mDNS Profile:</p> <p>TIP: You can add multiple rules.</p> <ol style="list-style-type: none"> Click Add under mDNS Profile Rules. The Create mDNS Profile Rule window opens. Provide the following settings for the mDNS rule: <ul style="list-style-type: none"> ● Host Name Glob—Indicate a host for the mDNS Profile. An asterisk like this * (the default) is a wildcard meaning all hosts. ● Service Name—Select one or more services for mDNS, either Apple TV, Internet printer, or Digital Auto Access Protocol (iTunes): <ul style="list-style-type: none"> ● _airplay._tcp—Apple TV ● _ipp._tcp—Internet printer ● _daap._tcp—Digital Auto Access Protocol (iTunes) ● Service Type—Select either Discover (default) or Advertise. ● VLAN Scope—Select either Global or Local. <ul style="list-style-type: none"> ● Local—Service information is only relevant on the local VLAN. ● Global—Service information is relevant beyond the local VLAN. <p>You can also type a service name with the following pattern <code>_<service>.<protocol></code></p> Click OK. The mDNS rule is added to the list of mDNS Profile Rules on the Create mDNS Profile page.
Task: Edit a rule in the mDNS Profile	Select a rule from the mDNS Profile Rules list, and then click Edit . Make changes, and then click OK .
Task: Remove a rule from the mDNS Profile	Select a rule from the mDNS Profile Rules list, and then click Delete . Click OK .

What to Do Next

An mDNS Profile must be assigned to specific devices— see “[Assigning an mDNS Profile to Devices](#)” on [page 1001](#).

RELATED DOCUMENTATION

[Creating and Managing an mDNS VLAN List | 1002](#)

[Understanding Bonjour | 995](#)

[Assigning an mDNS Profile to Devices | 1001](#)

[Network Director Documentation home page](#)

Assigning an mDNS Profile to Devices


Add support for Apple TV, Internet printing, and/or iTunes to a device by assigning an mDNS Profile to the device.

You must have an existing mDNS profile to assign—to create one, see “[Creating and Managing mDNS Profiles](#)” on [page 997](#).

To assign an mDNS Profile to devices:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  **Build** in the Network Director banner.
3. In the View pane, select a controller.
4. In the Tasks pane, expand **Wireless**, expand **Profiles**, and then select **mDNS**.
The Manage mDNS Profile page appears, displaying the list of currently configured mDNS Profiles.
5. Select an mDNS Profile from the list, and then click **Assign**.

The Assign mDNS Profile wizard opens. The wizard consists of three sections, **Device Selection**, **Profile Assignment**, and **Review**.

6. From **Device Selection**, select one or more devices for mDNS Profile assignment.

7. Click either **Next** or **Profile Assignment**.

The Assignments list is displayed with the selected devices.

8. Select any or all devices from the Assignments list and then click **Assign to AP**.

The Assign Profile to Managed Nodes window opens, displaying a list of access points on the device.

9. Select any or all listed access points, and then click **Assign**.

The Assign Profile to Managed Nodes window closes and the new assignment is added to the Assignments list.

10. Click either **Next** or **Review**.

Devices with mDNS assignments are listed.

11. To make changes, click **Edit**, make the changes, and then click the **Review** section of the wizard again.

12. Click **Finish**.

RELATED DOCUMENTATION

[Creating and Managing mDNS Profiles | 997](#)

[Understanding Bonjour | 995](#)

[Network Director Documentation home page](#)

Creating and Managing an mDNS VLAN List

IN THIS SECTION

- [Managing an mDNS VLAN List | 1003](#)
- [Creating an mDNS VLAN List | 1003](#)
- [Specifying mDNS VLAN Members | 1004](#)

mDNS provides a general method for Apple devices to set up a network without any configuration. You should limit the use of mDNS to specific VLANs to prevent interference with the rest of the network. By default, a list of VLANs is determined by the cache entries returned by the controller and by the sessions on the network looking for a particular type of service. You can optionally specify a list of VLANs to limit access to certain resources. For example, by placing Apple TV (_airplay service) on a separate VLAN and Internet Printer (_ipp service) on another VLAN, you can limit those bandwidth-hungry services.

This topic describes:

Managing an mDNS VLAN List

From the Manage mDNS VLAN List page, you can:

- Create a new wireless mDNS VLAN List by clicking **Add**. For directions, see [“Creating an mDNS VLAN List” on page 1003](#).
- Modify an existing mDNS VLAN List by selecting it and clicking **Edit**.
- View information about an mDNS VLAN List, including the interfaces it is associated with, by clicking the profile name or by selecting the profile and clicking **Details**.
- Delete an mDNS VLAN List by selecting a profile and clicking **Delete**.

TIP: You cannot delete profiles that are in use—that is, assigned to objects or used by other profiles. To see the current assignments for an mDNS VLAN List, select the mDNS VLAN List and click **Details**.

- Clone a VLAN list by selecting a profile and clicking **Clone**.


Creating an mDNS VLAN List

Create an mDNS VLAN list to limit access to certain mDNS resources. A VLAN list works with an mDNS Profile, limiting mDNS applications to specified VLANs. Also see [“Creating and Managing mDNS Profiles” on page 997](#).

To create an mDNS VLAN list:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.
3. In the Tasks pane, expand **Wireless**, expand **Profiles**, and then click **mDNS Profile**.
The Manage mDNS Profile page appears, displaying a list of currently configured mDNS Profiles and their attributes.
4. Click **Add**.
The Create VLAN List page appears.
5. Specify the settings as described in both the online help and in [“Specifying mDNS VLAN Members” on page 1004](#).
6. Click **Done** to save the Authorization profile.
The system saves the Authorization profile and displays the Manage Authorization Profiles page. Your new or modified Authorization profile is listed in the table of Authorization profiles.

Specifying mDNS VLAN Members

mDNS services such as Apple TV, Internet printer, or Digital Auto Access Protocol (iTunes) can consume extensive resources. To limit use of mDNS services to named VLANs:

1. Provide a name and description for the mDNS VLAN List.

2. Select existing VLANs to add to the list:
 - a. Click **VLAN** under *Select VLAN and/or VLAN Pool*.

The Select VLAN window opens, displaying a list of configured VLANs and VLAN pools.
 - b. Select one or more VLANs by placing a check mark next to the name.
 - c. Click **OK**.

The selected VLAN pools are added to the list of VLANs and VLAN pools.

3. Select existing VLAN pools to add to the list:
 - a. Click **VLAN Pool** under *Select VLAN and/or VLAN Pool*.

The Select VLAN Pool window opens, displaying a list of configured VLAN Pools.
 - b. Select one or more VLAN pools by placing a check mark next to the name.
 - c. Click **OK**.

The selected VLAN pools are added to the list of VLANs and VLAN pools.
 - d. Click **Done**.

RELATED DOCUMENTATION

[Creating and Managing mDNS Profiles | 997](#)

[Understanding Bonjour | 995](#)

[Assigning an mDNS Profile to Devices | 1001](#)

[Network Director Documentation home page](#)

Creating and Managing Local Switching Profiles

IN THIS SECTION

- [Managing Local Switching Profiles | 1006](#)
- [Creating a Local Switching Profile | 1008](#)

- Specifying Local Switching Profile Settings | 1008
- What To Do Next | 1010

With local packet switching, packets switch directly from access points to the wired network instead of passing through a controller. When local switching is enabled, the client VLAN is directly accessible through the wired interface on the access points and packets are switched directly to and from this interface. Access points configured to perform local packet switching can perform better than access points tunneling data through a controller. Local switching is disabled by default.

TIP: An access point can be configured to switch packets for some VLANs locally and tunnel packets for other VLANs through the controller switch.

Use the Local Switching page to create new Local Switching Profiles and manage existing Local Switching Profiles.

Managing Local Switching Profiles

From the Manage Local Switching page, you can:

- Create a new Local Switching Profile by clicking **Add**. For directions, see [“Creating a Local Switching Profile” on page 1008](#).
- Modify an existing Local Switching Profile by selecting it and clicking **Edit**.
- Assign a Local Switching Profile to access points by selecting the profile and clicking **Assign**. For directions, see [“Assigning a Local Switching VLAN Profile to Existing Access Points” on page 1011](#).
- Edit an existing Local Switching Profile by selecting it and clicking **Edit Assign**.
- View information about a Local Switching Profile by either clicking the profile name or by selecting the profile and clicking **Details**.
- Delete a Local Switching Profile by selecting site name and clicking **Delete**.

TIP: You cannot delete a profile that is in use. To see the current state of a profile, select the site name and click **Details**.

- Clone a Local Switching Profile by selecting a profile and clicking **Clone**.

Table 221 describes the information provided about Local Switching Profiles on the Manage Switching Profiles page. This page lists all Local Switching Profile defined for your network, regardless of the scope you selected in the network view.

Table 221: Local Switching Profile Information

Field	Description
Name	Type a unique name that identifies the profile. Use up to 32 characters for wireless profile names. Profile names must not contain special characters or spaces. Note that profiles automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character.
Description	Add any description here.
VLAN Members	VLANs that belong to this Local Switching Profile
Assignment State	Displays the assignment state of the profile. A profile can be: <ul style="list-style-type: none"> • Unassigned—When the profile is not assigned to any object. • Deployed—When the profile is assigned and is deployed from Deploy mode. • Pending Deployment—When the profile is assigned, but not yet deployed in the network.
Creation Time	Date and time when the profile was created.
Last Updated Time	Date and time when the profile was last modified.
User Name	The username of the person who created or modified the profile.

TIP: All columns might not be displayed. To show or hide fields listed in the table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

Creating a Local Switching Profile

To create a Local Switching Profile, follow these steps:

1. Under Views, select either **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.

3. Click **Add**.

The Create Local Switching page opens.

4. Provide the local switching settings listed in [“Specifying Local Switching Profile Settings” on page 1008](#).

5. Click **Done**.

Specifying Local Switching Profile Settings

Local Switching Profiles consist of a list of VLANs whose members perform local packet switching. To add a VLAN to a Local Switching Profile, complete the tasks in [Table 222](#).

Table 222: Adding VLANs to a Local Switching Profile

Field	Action
Name and Description	<p>Type a unique name that identifies the profile.</p> <p>Use up to 32 characters for wireless profile names. Profile names must not contain special characters or spaces. Note that profiles automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character.</p>

Table 222: Adding VLANs to a Local Switching Profile (*continued*)

Field	Action
Task: Add existing VLANs to this local switching profile	<ol style="list-style-type: none"><li data-bbox="675 338 1458 464">1. Click Select. The Select VLAN Members window opens, displaying a list of all VLANs with device family WLC currently available in Network Director.<li data-bbox="675 506 1458 537">2. Select one or more VLANs from the list.<li data-bbox="675 590 1458 663">3. You can edit the VLAN Mode, tagged/untagged property, and VLAN tag values for the selected VLAN<li data-bbox="675 705 1458 789">4. Click OK. The VLANs are added to the list of Selected VLAN Members.

Table 222: Adding VLANs to a Local Switching Profile (*continued*)

Field	Action
Task: Add the names of VLANs not yet managed by Network Director	<ol style="list-style-type: none"> <li data-bbox="675 338 1458 422">1. Click Add. The Add VLAN Members window opens. <li data-bbox="675 474 1458 506">2. Type the name of a VLAN on another controller in this domain. <li data-bbox="675 558 1458 621">3. When the optional tag-value is set, it is used as the 802.1Q tag for the VLAN. Enable Tagged and provide a Tag Value. <li data-bbox="675 674 1458 1178">4. Select a VLAN Mode: <ul style="list-style-type: none"> <li data-bbox="716 726 1458 936">● Local Switching—Use VLANs in the profile for packet transfer. By default, the VLAN mode is local-switching. When a client connects to an access point with local switching enabled and assigned a VLAN, if the VLAN is part of the VLAN profile, the session is locally switched. However, if a VLAN is not part of the VLAN profile, the session uses overlay mode. <li data-bbox="716 957 1458 1020">● Overlay—Use AP to AP tunneling to create and terminate client VLAN tunnels. This way, a VLAN is not required on every access point. <li data-bbox="675 1041 1458 1178">5. Optionally add an mDNS Profile to this VLAN to enable Apple TV, Internet printing, or iTunes use on the VLAN. Click Select next to mDNS Profile, select a profile from the list, and then click OK. The name of the profile appears in the mDNS Profile field. <li data-bbox="675 1230 1458 1388">6. Click OK. The Add VLAN Members window closes and the VLAN(s) are added to the list of selected VLAN members on the Create Local Switching Profile page. <li data-bbox="675 1440 1458 1598">7. Click Done. The Create Local Switching Profile page closes and the Manage Local Switching Profiles pages is displayed with the new profile added to the Manage Local Switching Profiles list.

What To Do Next

Assign VLAN Switching Profiles to access points—see [“Assigning a Local Switching VLAN Profile to Existing Access Points”](#) on page 1011.

RELATED DOCUMENTATION

[Understanding Local Switching on Access Points | 907](#)

[Assigning a Local Switching VLAN Profile to Existing Access Points | 1011](#)

[Creating and Managing Remote Site Profiles | 1014](#)

[Creating and Managing VLAN Pools | 535](#)

[Network Director Documentation home page](#)

Assigning a Local Switching VLAN Profile to Existing Access Points


Local switching uses VLANs to switch packets directly from access points to the wired network instead of passing through a controller. When one of those VLANs, indicated in a Local Switching Profile, is applied to an access point, that access point traffic is locally switched. This topic explains how to assign a Local Switching Profile to existing access points. Before you assign access points to Local Switching Profiles, you must first create those profiles indicating the VLAN members—for directions, see [“Creating and Managing Local Switching Profiles” on page 1005](#). You can also assign a Local Switching Profile when you add a new access point—see [“Assigning a Local Switching Profile During Access Point Configuration” on page 1013](#).

TIP: When applying a Local Switching Profile to an access point causes traffic previously tunneled through a controller to be locally switched, the access point’s client sessions are terminated, and the clients must re-associate with the access point. This also happens if an access point stops using a Local Switching Profile.

To assign a Local Switching Profile to existing access points:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.
3. From the Tasks pane, expand **Wireless**, expand **Profiles**, and then select **Local Switching VLAN**.

The Manage Local Switching VLAN Profiles page is displayed with a list of all current Local Switching Profiles.

4. Select one of the listed Local Switching VLAN Profiles and then click **Assign**.

The Assign Local Switching VLAN Profile wizard opens with three sections, Device Selection (default selection), Profile Assignment, and Review.

5. In Device Selection, select controllers by choosing at least one device for assignment. If you select Wireless Network, all controllers are selected.

6. Click either **Profile Assignment** in the wizard or **Next**.

The Profile Assignment section of the wizard opens with all selected controllers listed under Assignments.

7. Assign the Local Switching VLAN Profile to access points by selecting one or more controllers from the Assignments list and then clicking **Assign to AP**.

The Assign Profile to Managed Nodes window opens.

8. Select individual access points from the Managed Node Name list and then click **Assign**.

The access points are added to the Assignments list under their controller.

9. Click either **Review** in the wizard or **Next**.

10. Make any needed changes by clicking **Edit**.

The wizard reverts to Profile Assignment where you can make changes and then click either **Review** in the wizard or **Next** again to return to the review.

11. Click **Finish**.

The message *Assignments in progress. Please Wait.* is displayed until assignments are complete. The assignment is then added to the Manage Local Switching Profiles page.

TIP: Local Switching VLAN Profiles can also be assigned to existing access points (see [“Assigning a Local Switching VLAN Profile to Existing Access Points” on page 1011](#)), and included in Remote Site Profiles (see [“Creating and Managing Remote Site Profiles” on page 1014](#)).

[Understanding Local Switching on Access Points | 907](#)

[Creating and Managing Local Switching Profiles | 1005](#)

[Network Director Documentation home page](#)

Assigning a Local Switching Profile During Access Point Configuration

Local switching uses VLANs to switch packets directly from access points to the wired network instead of passing through a controller. When one of those VLANs, indicated in a Local Switching Profile, is applied to an access point, that access point traffic is locally switched. This topic explains how to assign a Local Switching Profile to access points. Before you assign access points to Local Switching Profiles, you must first create those profiles indicating the VLAN members—for directions, see [“Creating and Managing Local Switching Profiles” on page 1005](#).

There are two ways to assign a Local Switching Profile to an access point. You can create a Local Switching Profile before you configure your access points—then, as you configure access points, you add the existing Local Switching Profile. These directions are for that method—the directions assume that a Local Switching Profile exists but access point configuration does not exist.

TIP: You can also configure access points first, then create a Local Switching Profile, and then assign the Local Switching Profile to the configured access points—for directions for this method, see [“Assigning a Local Switching VLAN Profile to Existing Access Points” on page 1011](#).

TIP: When applying a Local Switching Profile to an access point causes traffic previously tunneled through a controller to be locally switched, the access point’s client sessions are terminated, and the clients must re-associate with the access point. This also happens if an access point stops using a Local Switching Profile.

You must have a configured Local Switching Profile for this procedure—for directions, see [“Creating and Managing Local Switching Profiles” on page 1005](#). Then, follow the directions to

RELATED DOCUMENTATION

[Creating and Managing Local Switching Profiles | 1005](#)

[Assigning a Local Switching VLAN Profile to Existing Access Points | 1011](#)

Creating and Managing Remote Site Profiles

IN THIS SECTION

- [Managing Remote Site Profiles | 1014](#)
- [Creating a Remote Site Profile | 1016](#)
- [Specifying Remote Site and Intrusion Detection Logging Settings | 1016](#)
- [What To Do Next | 1024](#)

Networks often include a central network site with controllers and one or more remote sites with only access points. The central and remote sites are connected by a WAN link. If you have configured a Remote Site Profile and the WAN link becomes unavailable, then the remote sites remain active and continue to provide connectivity to wireless clients.

Use the Manage Remote Sites page to create new Remote Site Profiles and manage existing Remote Site Profiles.

This topic describes:

Managing Remote Site Profiles

From the Manage Remote Sites page, you can:

- Create a new Remote Site Profile by clicking **Add**. For directions, see [“Creating a Remote Site Profile” on page 1016](#).
- Modify an existing Remote Site Profile by selecting it and clicking **Edit**.
- Assign an existing Remote Site Profile by selecting the profile and clicking **Assign**. For directions, see [“Assigning Remote Site Profiles to Access Points” on page 1024](#). You can also assign a remote site profile when you add an access point to Network Director—see [“Adding and Managing an Individual Access Point” on page 1156](#).
- Change an existing Remote Site Profile assignment by selecting it and clicking **Edit Assignment**.
- View information about a Remote Site Profile by either clicking the profile name or by selecting the profile and clicking **Details**.
- Delete a Remote Site Profile by selecting site name and clicking **Delete**.

TIP: You cannot delete sites that are in use. To see the current state of a site, select the site name and click **Details**.

- Clone a Remote Site Profile by selecting a profile and clicking **Clone**.

Table 223 describes the information provided about remote sites on the Manage Remote Sites page. This page lists all Remote Site Profiles defined for your network, regardless of the scope you selected in the network view.

Table 223: Remote Site Profile Information

Field	Description
Remote Site Name	Name of the Remote Site Profile
Country Code	Country Code of the Remote Site Profile. If no country code was selected, the Remote Site Profile uses the controller's country information.
Backup SSID Mode	When enabled , SSIDs configured to be available only during an outage (backup SSIDs) are enabled at the Remote Site during an outage.
Local Switching Profile	Local Switching Profile used for local switching at the remote site. To create a Local Switching Profile, see “Creating and Managing Local Switching Profiles” on page 1005. TIP: Access points' assigned Local Switching Profiles take precedence over the Remote Site's assigned Local Switching Profile.
Assignment State	Displays the assignment state of the profile. A profile can be: <ul style="list-style-type: none"> • Unassigned—When the profile is not assigned to any object.. • Deployed—When the profile is assigned and is deployed from Deploy mode. • Pending Deployment—When the profile is assigned, but not yet deployed in the network.
Creation Time	Date and time when the profile was created.
Last Updated Time	Date and time when the profile was last modified.
User Name	The username of the person who created or modified the profile.


TIP: All columns might not be displayed. To show or hide fields listed in the table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

Creating a Remote Site Profile

To create a remote site, follow these steps:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  **Build** in the Network Director banner.
3. Click **Add** on the Manage Remote Sites page.
The Create Remote Site page opens.
4. On the Create Remote Site page, provide information listed in [“Specifying Remote Site and Intrusion Detection Logging Settings” on page 1016](#).
5. Click **Done**.

The remote site is created and added to the list on the Manage Remote Sites page.

Specifying Remote Site and Intrusion Detection Logging Settings

To create a Remote Site profile, configure the settings in described in [Table 224](#).

Table 224: Remote Site Profile Settings

Field	Description
Remote Site Settings	

Table 224: Remote Site Profile Settings (continued)

Field	Description
Remote Site Name	Type a unique name that identifies the profile. Use up to 32 characters for wireless profile names. Profile names must not contain special characters or spaces. Note that profiles automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character.
Country Code (default is None)	Select a two letter Country Code for the Remote Site profile. If you do not select a country code, the Remote Site profile uses the controller's country information.
Local Switching Profile (default is None)	Indicate a Local Switching profile to be used for local switching at the remote site—the default is None . To create a Local Switching profile, see “Creating and Managing Local Switching Profiles” on page 1005 . TIP: Access points' configured Local Switching profiles take precedence over the Remote Site's Local Switching profile.
Path MTU (default is 0)	Maximum Transmission Unit, which is the size (in bytes) of the largest protocol data unit that the layer can pass onward for the WAN link—default is zero.
Backup SSID Mode (default is disabled)	Specific SSIDs can be made available either only during an outage or during normal operation. All SSIDs configured to be available only during an outage are backup SSIDs. When Backup SSID Mode is enabled, backup SSIDs are enabled at the Remote Site during an outage.
Intrusion Detection Logging	

Table 224: Remote Site Profile Settings (continued)

Field	Description
Enable Intrusion Detection Logging	Enables Intrusion Detection Logging for the Remote Site and displays the remaining settings.
	Server IP Address: IP address of the log server
	Server Port: Port used on the log server
	Severity Filter: Select one of the security filter options: <ul style="list-style-type: none"> • Emergency • Alert • Critical • Error (default) • Warning • Notice • Info • Debug all

Advanced Settings

Expand **Advanced Settings** to see RADIUS Attributes and Venue Names.

RADIUS Attributes

RADIUS Server Group (None by default)	Select an existing Radius Server Group from the list—the rest of the RADIUS Server Group settings then become available. You must have an existing Radius Server Group—to create one, see “Creating and Managing RADIUS Profiles” on page 339 .
	RADIUS Timeout: Length of time in seconds that MSS waits for a RADIUS server to respond when an authentication or authorization request is sent. Min : -1, Default: -1, Max: 65535
	RADIUS Retransmit Count: Retransmit RADIUS authentication after this number of seconds. Min : -1, Default: -1, Max: 100
	RADIUS Deadtime: Unavailable time for server after no response—after this length of time, a server is again a candidate to receive requests. Min : -1, Default: -1, Max: 1440
	RADIUS NAS ID: Name of the RADIUS client that originated the access request.

Table 224: Remote Site Profile Settings (continued)

Field	Description
Enable WLC Polling (default is disabled)	Enables controller polling after the outage time has elapsed. If the WAN link to the controller is established, the access point is not rebooted. Disabled by default.
Persistent Configuration (default is disabled)	When an access point reboots after the WAN outage time elapses, the access point polls the controller for its configuration. If Persistent Configuration is enabled, the access point reboots with the configuration it received from the controller before the outage occurred. Disabled by default.

Venue Names (802.11u)

802.11u is an IEEE standard that supports WLAN Internet working with external networks. This standard details the concept of offloading network traffic from cellular carriers to a Wi-Fi network to reduce traffic on expensive 3G/4G networks. Current solutions for hotspots require client devices to manually identify and select the local network as well authenticate to it. This service might offer varying levels of security, bandwidth capability, and quality. The Juniper Networks wireless LAN (WLAN) solution supports Wi-Fi certified passpoint (Hotspot) requirements and can seamlessly onboard Wi-Fi client devices at Hotspot deployments that enables both mobile operators and Multiple System Operators (MSOs) to also offload mobile data traffic onto Wi-Fi Hotspots.

Venue Group (default is 0)	Venue group is a number that identifies a general category of venues. The venue name is used by the Access Network Query Protocol (ANQP) to identify the Access Point. ANQP helps mobile clients learn more about the network before forming a connection. The meanings of the numbers are described in the directions below. Alpha-numeric Min Length :1 Max Length: 255
Venue Type (default is 0)	User defined Venue type from 0 through 11. The meanings of the numbers are described in the directions below.

Table 224: Remote Site Profile Settings (continued)

Field	Description
Task: Add a Venue to a Remote Site Profile	

Table 224: Remote Site Profile Settings (continued)

Field	Description
	<p>To add a venue to a Remote Site profile:</p> <ol style="list-style-type: none"> 1. Select a Venue Group number for the venue. The venue group number indicates that the venue is: <ul style="list-style-type: none"> ● 0—Unspecified ● 1—Assembly ● 2—Business ● 3—Educational ● 4—Factory and Industrial ● 5—Institutional ● 6—Mercantile ● 7—Residential ● 8—Storage ● 9—Utility and Miscellaneous ● 10—Vehicular ● 11—Outdoor 2. Select a Venue Type number for the venue. The venue type number is combined with the venue group code (listed above) to more specifically define the venue. In the following list, the venue group code is listed first, followed by the venue type code and a description. <ul style="list-style-type: none"> ● 0-0: Unspecified ● 1-0: Unspecified assembly ● 1-1: Arena ● 1-2: Stadium ● 1-3: Passenger Terminal (for example, an airport, bus, ferry, or train station) ● 1-4: Amphitheater ● 1-5: Amusement Park ● 1-6: Place of Worship ● 1-7: Convention Center ● 1-8: Library ● 1-9: Museum ● 1-10: Restaurant ● 1-11: Theater ● 1-12: Bar ● 1-13: Coffee Shop ● 1-14: Zoo or Aquarium

Table 224: Remote Site Profile Settings (*continued*)

Field	Description
	<ul style="list-style-type: none"> ● 1-15: Emergency Coordination Center ● 2-0: Unspecified Business ● 2-1: Doctor or Dentist office ● 2-2: Bank ● 2-3: Fire Station ● 2-4: Police Station ● 2-6: Post Office ● 2-7: Professional Office ● 2-8: Research and Development Facility ● 2-9: Attorney Office ● 3-0: Unspecified Educational ● 3-1: School, Primary ● 3-2: School, Secondary ● 3-3: University or College ● 4-0: Unspecified Factory and Industrial ● 4-1: Factory ● 5-0: Unspecified Institution ● 5-1: Hospital ● 5-2: Long-Term Care Facility (for example, a nursing home, hospice, etc.) ● 5-3: Alcohol and Drug Rehabilitation Center ● 5-4: Group Home ● 5-5: Prison or Jail ● 6-0: Unspecified Mercantile ● 6-1: Retail Store ● 6-2: Grocery Market ● 6-3: Automotive Service Station ● 6-4: Shopping Mall ● 6-5: Gas Station ● 7-0: Unspecified Residential ● 7-1: Private Residence ● 7-2: Hotel or Motel ● 7-3: Dormitory ● 7-4: Boarding House ● 8-0: Unspecified Storage ● 9-0: Unspecified Utility and Miscellaneous

Table 224: Remote Site Profile Settings (continued)

Field	Description
	<ul style="list-style-type: none"> ● 10-0: Unspecified Vehicular ● 10-1: Automobile or Truck ● 10-2: Airplane ● 10-3: Bus ● 10-4: Ferry ● 10-5: Ship or Boat ● 10-6: Train ● 10-7: Motor Bike ● 11-0: Unspecified Outdoor ● 11-1: Muni-mesh Network ● 11-2: City Park ● 11-3: Rest Area ● 11-4: Traffic Control ● 11-5: Bus Stop ● 11-6: Kiosk <p>3. Click Add under Venue Names (802.11u).</p> <p>A check mark appears next to Venue Name and the cursor moves to the Venue Name column.</p> <p>4. Type a name for the venue in the Venue Name field.</p> <p>5. Select a Venue Language Code to define the language used in the Venue Name field. The Language Code field is a two or three character language code selected from ISO-639.</p> <p>6. Click Update.</p> <p>The Venue Name and Venue Code that you just added become part of the Venue Names list.</p>
<p>Task: Remove a Venue from a Remote Site Profile</p>	<p>To remove a venue from a Remote Site profile:</p> <ol style="list-style-type: none"> 1. Select a venue from the Venue Names list. 2. Click Remove. <p>The venue name disappears from the Venue Names list.</p>

Table 224: Remote Site Profile Settings (continued)

Field	Description
WAN Attributes	
WAN Uplink Speed	Limit the WAN uplink speed in kbps. Range is 0-2147483647. Default : 0 (no limit)
WAN Downlink Speed	Limit the WAN downlink speed in kbps. Range is 0-2147483647. Default : 0 (no limit)

What To Do Next

Assign the Remote Sites to access points by following the directions in [“Assigning Remote Site Profiles to Access Points”](#) on page 1024.

NOTE: Assigned settings from any profile, including this one, have lower priority than settings made directly to a controller or an access point. For more information, see [“Adding and Managing an Individual Access Point”](#) on page 1156 and [“Configuring a Controller”](#) on page 1037.

RELATED DOCUMENTATION

[Assigning Remote Site Profiles to Access Points | 1024](#)

[Network Director Documentation home page](#)

Assigning Remote Site Profiles to Access Points


Remote Sites are assigned to access points either from the Manage Remote Sites page (directions are described here) or during creation of an access point—for those directions, see [“Adding and Managing an Individual Access Point”](#) on page 1156.

Before you begin, you need at least one configured Remote Site profile. For directions, see [“Creating and Managing Remote Site Profiles”](#) on page 1014.

To assign a Remote Site profile to one or more access points:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.
3. Select a Remote Site profile from the list and then click **Assign**.
The Assign Remote Site wizard opens, displaying a list of wireless devices. The wizard consists of three sections, Device Selection (selected), Profile Assignment, and Review.
4. From Device Selection, select one or more wireless devices from the list. If you select a device such as a controller, the entire Wireless Network, or My Network, all devices below that selection are also selected.
5. Click either **Profile Assignment** or **Next**.
The devices you selected are listed on the Profile Assignment page of the wizard.
6. Select one of the listed devices for assignment and then click **Assign to AP** to see the access points on that device.
The Assign Profile to Managed Nodes window opens, displaying a list of access points found on the selected device.
7. Select one or more access points from the list and then click **Assign**.
The Assign Profile to Managed Nodes window closes.
8. Click either **Review** or **Next**.
The assignments are listed in the Assignments window.
9. To make any changes in the Review section of the wizard, Click **Edit** and then make the changes.
10. Click **Finish**.

The Create Profile Assignment Job Details window opens. If the job is 100% complete, and the listed status is SUCCESS, the Remote Site profile is now listed to the access points.

11. Click **OK**.

The Remote Site profile name appears on the Manage Remote Sites page with the assignment state **Pending Deployment**. Deploy the pending Remote Site profile following the directions in [“Deploying Configuration to Devices”](#) on page 1180.

NOTE: Assigned settings from any profile, including this one, have lower priority than settings made directly to a controller or an access point. For more information, see [“Adding and Managing an Individual Access Point”](#) on page 1156 and [“Configuring a Controller”](#) on page 1037.

RELATED DOCUMENTATION

| [Network Director Documentation home page](#)

Creating and Managing RF Detection Profiles

IN THIS SECTION

- [Managing RF Detection Profiles | 1027](#)
- [Creating an RF Detection Profile | 1028](#)
- [Specifying RF Detection Profile Classification Settings | 1029](#)
- [What To Do Next | 1032](#)

In addition to sending and receiving data, radios provide RF detection, locating and tracking other electronic device signals on the network. When active scan is enabled in a Radio profile, the radios with an RF Detection profile actively scan other channels in addition to the data channel that is currently in use. Active scan operates on both enabled radios and disabled radios. For more information about scanning, see [“Understanding Wireless Scanning”](#) on page 869.

TIP: A radio in sentry mode is a dedicated scanner (no data transmission) providing better RF detection because the radio spends more time scanning each channel.

You must indicate how to classify the information gathered—to do this in Network Director, you create RF Detection profiles. You can set rules for devices to be classified as rogues, blocklisted devices, SSIDs, and friendly neighbor devices. You can also specifically add a device to the Rogues list, Blocklist, SSID list, or Neighbor list to classify the device yourself.

Managing RF Detection Profiles

From the Manage RF Detection page, you can:

- Create a new RF Detection profile by clicking **Add**. For directions, see [“Creating an RF Detection Profile” on page 1028](#).
- Modify an existing RF Detection profile by selecting it and clicking **Edit**.
- Assign a RF Detection profile to access points by selecting the profile and clicking **Assign**. For directions, see [“Assigning RF Detection Profiles to Controllers” on page 1033](#).
- Edit an existing RF Detection profile by selecting it and clicking **Edit Assignment**.
- Delete a RF Detection profile by selecting site name and clicking **Delete**.

TIP: You cannot delete a profile that is in use. To see the current state of a profile, select the site name and click **Details**.

- Clone a RF Detection profile by selecting a profile and clicking **Clone**.

[Table 225](#) describes the information provided about RF Detection profiles on the Manage Switching profiles page. This page lists all RF Detection profile defined for your network, regardless of the scope you selected in the network view.

Table 225: RF Detection Profile Information

Field	Description
Profile Name	Unique name, assigned when the profile was created, that identifies the profile.
WLA Signature Enabled	Access points' WLA signatures can be disabled or enabled. When a WLA signature is enabled, MSS can detect an attempts to spoof management packets from the access point.

Table 225: RF Detection Profile Information (*continued*)


Field	Description
WLA Signature	Access points' WLA signatures are a set of bits in a management frame sent by an WLA as an identifier to MSS.
Dynamic Blacklist Timeout Enabled	Are devices automatically removed from the blacklist after a certain period of time? Enabled means they are automatically removed at some point, while disabled means they are never automatically removed.
Dynamic Blacklist Timeout	When a dynamic (automatic) blacklist has a timeout parameter for length of time on a blacklist, this value indicates the number of seconds before a device is removed from the blacklist.
Assignment State	Displays the assignment state of the profile. A profile can be: <ul style="list-style-type: none"> ● Unassigned—When the profile is not assigned to any object.. ● Deployed—When the profile is assigned and is deployed from Deploy mode. ● Pending Deployment—When the profile is assigned, but not yet deployed in the network.
Creation Time	Date and time when the profile was created.
Last Updated Time	Date and time when the profile was last modified.
User Name	The username of the person who created or modified the profile.

Creating an RF Detection Profile

To create a RF Detection profile for wireless devices, follow these steps:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.
3. In the Tasks pane, expand **Wireless**, expand **Profiles**, and then click **RF Detection**.

The Manage RF Detection Profiles page appears, displaying the list of currently configured RF Detection profiles.

4. Click **Add** on the Manage RF Detection Profiles page.

The Create RF Detection Profile page opens.

5. On the Create RF Detection Profile for Wireless page, provide the rule settings and individual configurations listed in [“Specifying RF Detection Profile Classification Settings” on page 1029](#).

6. Click **Done**.

The Create RF Detection Profile page closes and the RF Detection profile is added to the list on the Manage RF Detection Profiles page.

Specifying RF Detection Profile Classification Settings

Specify the RF Classification settings described in [Table 226](#).

Table 226: RF Detection Settings

Field	Description
Profile Name	Type a unique name, up to 32 characters, that identifies the profile. Profile names must not contain special characters or spaces. Note that profiles automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character.
Description	Type up to 256 characters.
WLA Signature (disabled by default)	An access point's WLA signature, a set of bits in a management frame sent by an access point as an identifier to MSS, can be disabled or enabled. When WLA signature is enabled, MSS can detect attempts to spoof management packets from the access point.
Dynamic Blacklist Timeout (enabled by default)	Number of seconds that a blocklisted client stays on the blacklist. The default is 300 seconds.

Table 226: RF Detection Settings (continued)

Field	Description
RF Classification Rules	
<p>Expand RF Classification Rules to see these settings. Eight RF classification rules interpret data gathered by RF detection, and then classify detected devices according to the rules. Some rules cannot be changed—for example, a device recognized as a rogue is also classified as a rogue. However, about half of the rules have options—for example, if a device is recognized as ad-hoc (not using an access point), you can elect to ignore it or to classify it as a rogue. To make changes to the rules, expand this RF Classification Rules section.</p>	
<p>NOTE: Any individual classification you do takes precedence over the rules.</p>	
<p>For more information about classification of RF data, see</p>	
Device is on rogue list	<p>Rule: Classify as Rogue Rule cannot be altered. For more information about rogue devices, see “Understanding Rogue Clients” on page 923.</p>
AP is part of Mobility Domain	<p>Rule: Classify as Member Rule cannot be altered.</p>
Device is on Neighbor list	<p>Rule: Classify as Neighbor Rule cannot be altered. For more information about mobility domains, see “Understanding Mobility Domains” on page 1051.</p>
SSID is not in your network	<p>Rule: Classify as Neighbor Rule cannot be altered.</p>
SSID has been determined to be an SSID Masquerade.	<p>Rule: Select either Classify as Rogue or Skip Test Classification (ignore). For an explanation of this situation, see “Understanding an SSID Masquerade” on page 926.</p>
Client Destination (DST) MAC address has been seen in the network	<p>Rule: Select either Classify as Rogue or Skip Test Classification (ignore).</p>
Device is on ad-hoc device list.	<p>Rule: Select either Classify as Rogue or Skip Test Classification (ignore). For an explanation of this situation, see “Understanding Ad-Hoc Networks” on page 927.</p>
When no other classification has been made, use this Default rule.	<p>Rule: Select either Classify as Suspect, Classify as Rogue or Skip Test Classification (ignore).</p>

Table 226: RF Detection Settings (continued)

Field	Description
<p>RF Classification Parameters: Classify Devices as Rogues</p>	
<p>You can define any device, by indicating its MAC address, as a rogue. This individual classification takes precedence over any classification done by the rules.</p>	
<p>Task: Add a rogue device to the Rogues List</p>	<ol style="list-style-type: none"> 1. Click Add under Rogues. The Create Rogue Device Entry window opens. 2. Enter a MAC address in the Create Rogue Device Entry window. 3. Click OK. The Create Rogue Device Entry window closes and the MAC address now appears in the MAC Address list under Rogues.
<p>RF Classification Parameters: Blocklist Devices</p>	
<p>You can blocklist any device, by indicating its MAC address. This individual classification takes precedence over any classification done by the rules.</p>	
<p>Task: Blocklist a device.</p>	<ol style="list-style-type: none"> 1. Click Add under Black List. The Create Black List Entry window opens. 2. Enter a MAC address in the Create Black List Entry window. 3. Click OK. The Create Black List Entry window closes and the MAC address now appears in the MAC Address list under Black List.
<p>RF Classification Parameters: Classify SSID as Known</p>	
<p>You can define any detected SSID as known. This individual classification takes precedence over any classification done by the rules.</p>	

Table 226: RF Detection Settings (continued)

Field	Description
Task: Add an SSID to the Known SSIDs list	<ol style="list-style-type: none"> 1. Click Add under SSIDs. The Create Known SSID Entry window opens. 2. Type the name of the SSID in the Create Known SSID Entry window. 3. Click OK. The Create Known SSID Entry window closes and the SSID now appears in the SSID list under SSIDs.

RF Classification Parameters: Classify Devices as Neighbors

You can define any device, by indicating its MAC address, as a neighbor. You can also define devices from any vendor as neighbors by using the Vendor OUI. This classification takes precedence over any classification done by the rules.

Task: Add a friendly device to the Neighbors List	<ol style="list-style-type: none"> 1. Click Add under Neighbors. The Create Neighbor Device Entry window opens. 2. Enter a MAC address in the Create Neighbor Device Entry window. The address is six bytes (xx:x:xx:xx:xx:xx) for a device or it can refer to a Vendor OUI in which case it is 3 bytes (xx:xx:xx). 3. Click OK. The Create Neighbor Device Entry window closes and the MAC address now appears in the MAC Address list under Neighbors.
---	---

What To Do Next

Assign the RF Detection Filter profile to an access point following the directions in [“Assigning RF Detection Profiles to Controllers”](#) on page 1033.

RELATED DOCUMENTATION

[Assigning RF Detection Profiles to Controllers](#) | 1033

[Network Director Documentation home page](#)

Assigning RF Detection Profiles to Controllers

An RF Detection profile can provide configuration for any of these devices:

- Rogue devices—Identify a device as a rogue by adding the MAC address to the rogues list. All automatic actions taken against rogues will be applied to devices on this list.
- Neighbor devices—To identify friendly devices, such as non-Juniper Networks access points in your network or a neighboring network, add the MAC address to the neighbors list.
- Devices you want blocklisted—Identify devices you want banned from the network by adding the MAC address to the blocklist.
- SSIDs you want blocklisted—Identify SSIDs you want banned from advertising on the network by adding the MAC address to the SSID list.

Before you begin, you need at least one configured RF Detection profile. For directions, see [“Creating and Managing RF Detection Profiles” on page 1026](#).

To assign an RF Detection profile to one or more controllers:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.

3. In the Tasks pane, expand **Wireless**, expand **Profiles**, and then click **RF Detection**.

The Manage RF Detection Profiles page appears, displaying the list of currently configured RF Detection profiles.

4. Select a RF Detection profile from the list and then click **Assign**.

The Assign RF Detection Profile wizard opens, displaying a list of wireless devices. The wizard consists of three sections, Device Selection (selected), Profile Assignment, and Review.

5. From Device Selection, select one or more controllers or controller clusters from the list. If you select **Wireless Network**, all controllers below that selection are also selected.

TIP: Controller clusters are created in mobility domains—see [“Creating a Mobility Domain for Wireless LAN Controllers”](#) on page 1053.

6. Click either **Profile Assignment** or **Next**.

The controllers you selected are listed on the Profile Assignment page of the wizard in the Assignments list.

7. Select one or more of the listed controllers for assignment and then click either **Assign to Device** for a controller or **Assign to Cluster** for a controller cluster,

The Assigned To column of the Assignments list reflects the new assignment.

8. Click either **Review** or **Next**.

The assignments are listed in the Assignments window.

9. To make any changes in the Review section of the wizard, Click **Edit** and then make the changes.

10. Click **Finish**.

The Create Profile Assignment Job Details window opens. If the job is 100% complete, and the listed status is SUCCESS, the RF Detection profile is now listed to the controller(s).

11. Click **OK**.

The RF Detection profile name appears on the Manage RF Detection Profile page with the assignment state **Pending Deployment**. Deploy the pending RF Detection profile following the directions in [“Deploying Configuration to Devices”](#) on page 1180.

RELATED DOCUMENTATION

[Creating and Managing RF Detection Profiles | 1026](#)

[Deploying Configuration to Devices | 1180](#)

[Network Director Documentation home page](#)

Configuring Link Layer Discovery Protocol (LLDP) on an Access Point

Link Layer Discovery Protocol (LLDP) is a link layer protocol used by network devices to advertise identity, capabilities, and neighbors. It also provides additional TLVs for capabilities discovery, network policy, Power over Ethernet (PoE), and inventory management.

Link Layer Discovery Protocol-Media Endpoint Discovery (LLDP-MED) is an extension to LLDP that operates between endpoint devices such as IP phones to provide support for voice over IP (VoIP) applications. LLDP-MED endpoints determine the capabilities of a connected device and whether those capabilities are enabled.

TIP: LLDP and LLDP-MED cannot operate simultaneously on a network. By default, access points send only LLDP packets until LLDP-MED packets are received from an endpoint device. The access point then sends out LLDP-MED packets until it receives LLDP packets.

For information about LLDP, see [“Understanding LLDP and LLDP-MED” on page 930](#).

To configure LLDP on the access point from the **LLDP** tab of the Add AP process, complete the configuration described in [Table 227](#).

Table 227: LLDP Settings for an Access Point

Field	Description
LLDP Mode (enabled by default)	<p>Link Layer Discovery Protocol (LLDP) on an access point is enabled for transmission by default, which means that access point transmits its identity, capabilities, and neighbors with LLDP. If you do not want the access point to use LLDP, disable LLDP by selecting Disable instead of transmit (TX).</p> <p>NOTE: You can enable both LLDP Mode and LLDP-MED Mode, but only one can operate at a time. By default, network devices send only LLDP packets until LLDP-MED packets are received from an endpoint device. The network device then sends out LLDP-MED packets until it receives LLDP packets.</p>

Table 227: LLDP Settings for an Access Point (continued)

Field	Description
LLDP-MED Mode (enabled by default)	<p>Media Endpoint Discovery is an enhancement of LLDP that is disabled by default. To have the access point transmit its identity, capabilities, and neighbors with the LLDP-MED protocol, enable LLDP-MED Mode and complete the two configuration options for LLDP-MED.</p> <p>NOTE: You can enable both LLDP Mode and LLDP-MED Mode, but only one can operate at a time. By default, network devices send only LLDP packets until LLDP-MED packets are received from an endpoint device. The network device then sends out LLDP-MED packets until it receives LLDP packets.</p> <hr/> <p>Power via MDI: Enable Power via Media Dependent Interface (MDI) to have the access point also convey power information, such as the type of power, power priority, and the amount of power required by the device. Information is collected on the Ethernet interface.</p> <hr/> <p>Inventory: Enable Inventory to have the access point also transmit detailed inventory information to a controller. Inventory information includes hardware revision, firmware version, software version, serial number, manufacturer name, model name, and asset ID.</p>

TIP: You can also configure LLDP on a controller—see [“Configuring a Controller”](#) on page 1037.

Next, click the **Radio 1** tab to configure the access point’s radio information. For directions, see [“Specifying Access Point Radio Settings”](#) on page 1168.

RELATED DOCUMENTATION

[Understanding LLDP and LLDP-MED | 930](#)

[Configuring a Controller | 1037](#)

Configuring Wireless Controllers

IN THIS CHAPTER

- [Configuring a Controller | 1037](#)

Configuring a Controller

IN THIS SECTION

- [Configuring System Information for a Controller | 1038](#)
- [Configuring Link Layer Discovery Protocol \(LLDP\) on a Controller | 1039](#)
- [Configuring IP Services on a Controller | 1040](#)
- [Configuring DSCP CoS Mapping on a Controller | 1042](#)
- [Configuring RF Auto-tuning on a Controller | 1043](#)
- [Configuring Load Balancing on a Controller | 1044](#)
- [Configuring AAA 802.1X on a Controller | 1045](#)
- [Configuring AAA RADIUS on a Controller | 1047](#)
- [Configuring AAA LDAP on a Controller | 1049](#)
- [What To Do Next | 1050](#)

Controllers can be configured from Network Director only when they are selected in the View pane.

To configure a controller in Network Director:

1. Under Views, select one of these options: **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.

3. Select a controller or controller cluster in the leftmost pane.

When a controller is selected, **WLC System Settings** is added to the list of Key Tasks in the Tasks pane.

NOTE: You only see the option **WLC System Settings** after you select a controller or controller cluster.

4. In the Tasks pane, expand **Key Tasks** and then click **WLC System Settings**.

The System Settings page opens, displaying three tabs—**System**, **Wireless**, and **AAA**. Make any needed changes under the appropriate sections on each tab:

- **System Configuration**—make changes by following the directions in [“Configuring System Information for a Controller” on page 1038](#), [“Configuring Link Layer Discovery Protocol \(LLDP\) on a Controller” on page 1039](#), [“Configuring IP Services on a Controller” on page 1040](#), and [“Configuring DSCP CoS Mapping on a Controller” on page 1042](#).
- **Wireless Configuration**—make changes by following the directions in [“Configuring RF Auto-tuning on a Controller” on page 1043](#) and [“Configuring Load Balancing on a Controller” on page 1044](#).
- **AAA Configuration**—make changes by following the directions in [“Configuring AAA 802.1X on a Controller” on page 1045](#), [“Configuring AAA RADIUS on a Controller” on page 1047](#), and [“Configuring AAA LDAP on a Controller” on page 1049](#).

5. Click **Done**.

The controller is now configured but not deployed—you must deploy the controller to activate the configuration. For more information, see [“Deploying Configuration to Devices” on page 1180](#).

Configuring System Information for a Controller

System information is received from the selected controller’s or controller cluster’s current controller configuration. You can change some of this information as indicated in [Table 228](#). This information is part of the System Settings for a controller.

Table 228: Basic System Information for Controllers

Field	Directions
Basic Information	
IP Address	No change can be made to information received from the selected controller.
Model	No change can be made to information received from the controller.
Serial Number	No change can be made to information received from the controller.
OS Version	No change can be made to information received from the controller.
Contact	You can change this value. Type a contact name for the controller.
Location	You can change this value. Type a location for the controller.
Prompt	You can change this value. Type the prompt that will appear in the CLI of the controller.
Application Detection	
Enable mDNS Detection if available	mDNS is only available on access points that support the feature when MSS 9.1 or later is the operating system. Check to enable multicast Domain Name System (mDNS) host name resolution service to support Apple mDNS.

Configuring Link Layer Discovery Protocol (LLDP) on a Controller

Link Layer Discovery Protocol (LLDP) is a link layer protocol used by network devices to advertise identity, capabilities, and neighbors. Link Layer Discovery Protocol-Media Endpoint Discovery (LLDP-MED) is an extension to LLDP that operates between endpoint devices such as IP phones to provide support for voice over IP (VoIP) applications.

TIP: LLDP and LLDP-MED cannot operate simultaneously on a network.

For information about LLDP, see [“Understanding LLDP and LLDP-MED” on page 930](#). To configure LLDP on an access point, see [“Configuring Link Layer Discovery Protocol \(LLDP\) on an Access Point” on page 1035](#).

To configure LLDP on a controller, provide the LLDP settings described in [Table 229](#). The settings are located under the controller’s or controller cluster’s System Settings LLDP option in Network Director.

Table 229: LLDP Settings for Controllers

Field	Directions
Enable LLDP (default is enabled)	Checked by default to enable the Link Layer Discovery Protocol on the controller. If you disable it, all TLVs (system capabilities, system name, system description) on the controller are discarded.
Transmission Interval (default is 30 seconds)	Specify the LLDP advertisement interval in seconds. The range is 5 to 32786 seconds and the default value is 30 seconds. LLDP frames can be sent earlier if local changes affect any of the selected TLVs (system capabilities, system name, system description).
Hold Time (default is 120 seconds)	Specify the length of time that a controller retains LLDP information before discarding it. The range is 0 to 65535 seconds with a default value of 120 seconds. We recommend a value four times the transmission interval value.
Reinitialization Delay (default is 2 seconds)	Configure the delay time, in seconds, before LLDP is initialized on any port. The range is 2 to 5 seconds with the default value of 2 seconds.
Transmit Delay (default is 2 seconds)	Specify the length of time between LLDP frame transmissions. The range is 1 to 8192 seconds with the default value of 2 seconds. The transmit delay value limits the rate at which local changes affect LLDP frames—frames sent advertise only the most recent changes. For example, if you change the controller name every 5 minutes, this triggers the network to send new LLDP advertisements. By setting the transmit-delay parameter, you can limit the rate at which new LLDP advertisements are sent on the network. LLDP frames are sent at each tx-interval number of seconds—If local changes occur then the frames are sent earlier, but not less than the value indicated here.
System Capabilities (default is enabled)	System capabilities is a TLV. When checked (default), this information is broadcast on the network.
System Name (default is enabled)	System name is a TLV. When checked (default), this information is broadcast on the network.
System Description (default is enabled)	System description is a TLV. When checked (default), this information is broadcast on the network.

Configuring IP Services on a Controller

To configure IP Services on a controller, provide the IP routes, IP aliases, and ARP settings described in [Table 230](#). The settings are located under the controller or controller cluster's System Settings IP Services option in Network Director.

Table 230: IP Routes, IP Aliases, and ARP Settings for Controllers

Field	Directions
-------	------------

Static Routes

A static route is an explicit route from a controller to a host. Static routes do not expire but they are removed by a software reboot.

<p>Task: Add a Static Route to a Controller</p>	<p>To add a static route to the selected controller or controller cluster:</p> <ol style="list-style-type: none"> Click Add under Static Routes. The Create Static Route window opens. Provide these static route settings: <ul style="list-style-type: none"> Default Route—Check this option to make this static route the default. Destination—Type the IP address of the static route destination host. Gateway—Type the IP address of the host's static route gateway. Metric—Select a number between 0 and 2,147,483,647 to indicate the cost of a route to a router. Lower-cost routes are preferred over higher-cost routes. When you add multiple routes to the same destination, MSS groups the routes together and orders them from lowest cost at the top of the group to highest cost at the bottom of the group. If you add a new route that has the same destination and cost as a previous route, MSS places the new route at the top of the group of routes with the same cost. Click OK. The static route is added to the Static Routes list on the System Settings page.
---	---

IP Aliases

An alias is a string that represents an IP address. After configuring the alias, you can refer to a controller using the alias name instead of the IP address. Aliases take precedence over DNS assignments. When you enter a hostname, MSS checks for an alias with that name first, then uses DNS to resolve the name.

Table 230: IP Routes, IP Aliases, and ARP Settings for Controllers (*continued*)

Field	Directions
Task: Add an IP Alias to a Controller	<p>To add an IP alias to the selected controller or controller cluster:</p> <ol style="list-style-type: none"> Click Add under IP Aliases. The Create IP Alias window opens. Provide these IP alias settings: <ul style="list-style-type: none"> Host Name—Alias name for the controller Host IP Address—IP address of the controller Click OK. The IP alias is added to the IP Alias list on the System Settings page.

ARP Settings

The Address Resolution Protocol (ARP) table maps IP addresses to MAC addresses. An ARP entry enters the table either automatically, or having been specifically configured with the parameters in this section.

Aging Time	The aging timeout specifies how long a dynamic entry can remain unused before the software removes the entry from the ARP table. The default aging timeout is 1200 seconds (20 minutes). The aging timeout does not affect the local entry, static entries, or permanent entries.
Task: Add an ARP Entry to a Controller	<p>To add an IP alias to the selected controller or controller cluster:</p> <ol style="list-style-type: none"> Click Add under ARP Settings. The Create ARP Entry window opens. Provide these ARP entry settings: <ul style="list-style-type: none"> IP Address—Type the IP address of the ARP Entry. MAC Address—Type the MAC address of the ARP Entry. Click OK. The ARP entry is added to the ARP Entries list on the System Settings page.

Configuring DSCP CoS Mapping on a Controller

Differentiated Services Code Point (DSCP) is a field in IPv4 and IPv6 packet headers. Class of Service (CoS) is a set of priority levels in quality-of-service (QoS) configurations. CoS-to-DSCP and DSCP-to-CoS mappings are used to evaluate packets for scheduling and assigning packets to one of the QoS queues. The result

is Layer 2 and Layer 3 classification and marking of traffic to help provide end-to-end quality-of-service (QoS) throughout a network. The settings to configure DSCP on a controller or controller cluster are located under the controller's System Settings IP Services option in Network Director.

Configure either **DSCP to CoS** or **CoS to DSCP** by selecting options from the list and then clicking **Done**.

Configuring RF Auto-tuning on a Controller

RF auto-tuning automatically makes channel tuning decisions for access point radios on the basis of the RF data gathered by access points. Configuring power auto-tuning a controller means that all associated radios that have not been otherwise configured will be auto-tuned. Use this feature to configure auto-tuning on the selected controller or controller cluster.

TIP: You can also configure radios with auto-tuning by creating a Radio profile and assigning that profile to radios. See [“Creating and Managing a Radio Profile” on page 932](#).

RF auto-tuning is configured for a controller or controller cluster under the **Wireless** tab of the selected controllers' System Settings. Provide the RF auto-tuning settings described in [Table 231](#) to configure the selected controller or controller cluster.

Table 231: RF Auto-Tune Settings for Controllers

Field	Directions
802.11b/g	
Enable (default is enabled)	When checked (default), 802.11b/g power auto-tuning is implemented for the controller.
Interference Domain Threshold (default is 85 radios)	Define an interference domain threshold, which is the maximum set of radios in a mobility domain that can interfere with each other. The default is 85 radios.
Convergence Delay (default is 60 minutes)	Indicate the length of the delay between beginning calculations of new configuration. The default value is 60 minutes with a possible range of 0 to 10080 minutes.
Week Day	Schedule auto-tuning for a specific day.
Hour	Schedule auto-tuning for a specific hour of the indicated day,
Minute	Schedule auto-tuning for a specific minute of the indicated hour.

Table 231: RF Auto-Tune Settings for Controllers (*continued*)

Field	Directions
802.11a	
Enable (default is enabled)	When checked (default), 802.11b/g channel auto-tuning is implemented for the controller.
Interference Domain Threshold (default is 85 radios)	Define an interference domain threshold, which is the maximum set of radios in a mobility domain that can interfere with each other. The default is 85 radios.
Convergence Delay (default is 60 minutes)	Indicate the length of the delay between beginning calculations of new channel plans. The default value is 60 minutes with a range of 0 to 10080 minutes.
Week Day	Schedule auto-tuning for a specific day
Hour	Schedule auto-tuning for a specific hour of the indicated day,
Minute	Schedule auto-tuning for a specific minute of the indicated hour.

Configuring Load Balancing on a Controller

Load balancing distributes a workload across multiple wireless radios to achieve optimal utilization, maximize throughput, minimize response time, and avoid overload. Radios with heavy client loads are made less visible to new clients, causing those clients to associate with radios with a lighter load. For more information, see [“Understanding Load Balancing for Wireless Radios” on page 1070](#).

Load balancing is configured for a controller or controller cluster under the **Wireless** tab of the selected controllers’ System Settings. See [Table 232](#) for a description of the possible load-balancing settings for controllers and controller clusters.

Table 232: Load-balancing Settings for Controllers

Field	Directions
Enable (default is enabled)	RF Load-balancing is enabled by default on controllers. You can disable it by removing this check mark.

Table 232: Load-balancing Settings for Controllers (*continued*)

Field	Directions
Load Balancing Strictness (default is low)	You can specify how strictly MSS attempts to load-balance across radios on the controller. When low strictness is specified (default), heavily loaded access point radios are less visible and steer clients to less-busy radios, while ensuring that clients are not denied service even if all the WLA radios in the group are heavily loaded. When maximum strictness is specified, a radio with the maximum client load is invisible to new clients, and clients attempt to connect to other radios. In the event that all the radios in the group reach maximum client load, then no new clients can connect to the network.
Preferred Band (default is 5-GHz)	If a client supports both the 802.11a and 802.11b/g bands, you can configure MSS to steer the client to a less-busy radio for the purpose of load-balancing. This band-preference option makes access points with two radios attempt to hide one of the radios from a client with the purpose of steering the client to the other radio. Select 5GHz , 2.4GHz , or None for the preferred band.

Configuring AAA 802.1X on a Controller

IEEE 802.1X network users are authenticated when they identify themselves with a credential. Authentication can be passed through to RADIUS, performed locally on the controller, or partially completed by the controller. Assign authentication settings for the selected controller or controller cluster as described in [Table 233](#).

Table 233: 802.1X Authentication Settings for Controllers

Field	Directions
802.1X Settings	
System Authentication Control (default is enabled)	When checked (default) enables 802.1X authentication for all wired authentication ports on the controller.
Retransmit Timeout (default is 5 seconds)	Specify the number of seconds before retransmitting an Extensible Authentication Protocol over LAN (EAPoL) packet. Default is 5 seconds.
Authentication Server Timeout (default is 30 seconds)	Indicate number of seconds before the controller times out a request to a RADIUS server. The default is 30 seconds.
Key Transmit (default is enabled)	When checked (default), encryption key information is sent to the client after authentication in EAPoL-Key PDUs. The controller sends EAPoL key messages after successfully authenticating the client and receiving authorization attributes for the client. If the client is using dynamic WEP, the EAPoL key messages are sent immediately after authorization.

Table 233: 802.1X Authentication Settings for Controllers (*continued*)

Field	Directions
Reauthentication Attempts (default is 2)	Number of reauthentication attempts the controller makes before the client becomes unauthorized. The default number of reauthentication attempts is 2. You can specify from 1 to 10 attempts.
Bonded Period (default is 0)	Specify the number of seconds (default is 0) that MSS retains session information for an authenticated computer while waiting for the 802.1X client on the computer to start (re)authentication for the user. Normally, the Bonded Auth period needs to be set only if the network has Bonded Auth clients that use dynamic WEP, or use WEP-40 or WEP-104 encryption with WPA or RSN. These clients can be affected by the 802.1X reauthentication parameter or the RADIUS Session-Timeout parameter.
Quiet Period Timeout (default is 60 seconds)	Indicate the number of seconds the controller is unresponsive to a client after a failed authentication. The default is 60 seconds. The acceptable range is from 0 to 65,535 seconds.
Supplicant Timeout (default is 30 seconds)	Indicate the interval of time before the controller retransmits an 802.1X-encapsulated EAP request to a client. If both the RADIUS and supplicant timeouts are set, MSS uses the shorter of the two. If the RADIUS session-timeout attribute is not set, MSS uses this timeout value, by default 30 seconds.
Maximum Requests (default is 2)	Indicate the maximum number of times (0 to 10) an EAP request is transmitted to the client before timing out the authentication session. The default is 2 attempts.
Reauthentication and Reauthentication Period (default is 3600 seconds)	Reauthentication of 802.1X wireless clients is enabled on the controller by default. By default, the controller waits 3600 seconds (1 hour) between authentication attempts. You can change the defaults. You can also disable reauthentication.
Handshake Timeout (default is 2 seconds)	Indicate the timeout for the 4-way handshake and the 802.1X group-key handshake on the controller. The default value is 2000 milliseconds (2 seconds) with a range of 20-5000 milliseconds.

WEP Key Rolling

For dynamic WEP, MSS dynamically generates keys for broadcast, multicast, and unicast traffic. MSS generates unique unicast keys for each client session and periodically regenerates (rotates) the broadcast and multicast keys for all clients. You can change or disable the broadcast or multicast rekeying interval.

Table 233: 802.1X Authentication Settings for Controllers (*continued*)

Field	Directions
WEP Key Rolling and WEP Key Rolling Period (default is 1800 seconds)	Change or disable the key rotation interval. The default interval for rotating the WEP key is 1800 seconds.
TKIP/CCMP Key Rolling	
Unicast Key Rolling and Unicast Key Rolling Period (default is 300 seconds)	Change or disable the broadcast TKIP and CCMP key rotation interval. The default interval for TKIP and CCMP key rotation is 300 seconds.
Multicast Key Rolling and Multicast Key Rolling Period (default is 300 seconds)	Change or disable the multicast key rotation interval. The default interval for rotating the multicast key is 300 seconds.

Configuring AAA RADIUS on a Controller

Remote Authentication Dial In User Service (RADIUS) is a networking protocol that provides centralized authentication, authorization, and accounting (AAA) management for computers to connect and use a network service. You can configure a single controller or controller cluster with RADIUS settings under the AAA tab of System Settings in Network Director.

TIP: For more information, see [“Understanding Central Network Access Using RADIUS and TACACS+” on page 335](#). To configure a RADIUS profile that can be assigned to devices, see [“Creating and Managing RADIUS Profiles” on page 339](#).

Assign AAA RADIUS settings for the selected controller or controller cluster as described in [Table 234](#).

Table 234: AAA RADIUS Settings for Controllers

Field	Directions
RADIUS Default Settings	
Timeout (default is 5 seconds)	Adjust the length of time (default is 5 seconds) that elapses with no connection before Network Director gives an unreachable RADIUS server error.
Retry Count (default is 3 times)	Adjust the number of times (default is 3) that a controller retries connection with a RADIUS server after the connection is dropped or refused.

Table 234: AAA RADIUS Settings for Controllers (*continued*)

Field	Directions
Dead Time (default is 5 seconds)	Indicate the number of seconds before Network Director checks a RADIUS server that was previously unresponsive. Default is five seconds.
Key	Indicate a password string that is the shared secret that the controller uses to authenticate to the RADIUS server. No default.
Use MAC as Password (default is disabled)	Check to set the password to the controller's MAC address. If you enable Use MAC As Password, then the Authorization Password field becomes unavailable.
Authorization Password	Provide a password if you are not using the MAC address as the password—there is no default.
MAC Address Format (default is hyphens)	Indicate the format of the MAC address (no default) that will be sent as a password, either Hyphens , Colons , One hyphen or Raw . For descriptions and examples of these formats, see “Creating and Managing RADIUS Profiles” on page 339 .
Authentication Protocol (default is PAP)	<p>Select PAP, CHAP, MSCHAP-V2, or None to determine an authentication protocol for the RADIUS server. These authentication protocols work as follows:</p> <ul style="list-style-type: none"> • PAP stands for Password Authentication Protocol and is used by Point to Point Protocols to validate users before allowing them access to server resources. Almost all network operating system remote servers support PAP. However, PAP transmits unencrypted ASCII passwords over the network and is therefore not secure. Use it as a last resort when the remote server does not support the stronger authentication. • CHAP stands for Challenge Handshake Authentication Protocol and authenticates a user or network host to an authenticating entity. CHAP provides protection against replay attacks by the peer through the use of an incrementally changing identifier and of a variable challenge-value. CHAP requires that both the client and server know the plaintext of the secret password—it is never sent over the network. CHAP provides better security than PAP does. • MSCHAP—stands for Microsoft's implementation of the Challenge Handshake Authentication Protocol version 2 on the router for password-change support. This feature provides users accessing a router the option of changing the password when the password expires, is reset, or is configured to be changed at the next login. The MS-CHAP variant does not require either peer to know the plaintext of the secret password. MSCHAP-V2 is used as an authentication option with RADIUS servers used for Wi-Fi security using the WPA-Enterprise protocol.

Table 234: AAA RADIUS Settings for Controllers (*continued*)

Field	Directions
RADIUS DAC Settings	
Dynamic Authorization Client (DAC) is a dynamic RADIUS extension that enables administrators supporting a RADIUS server to disconnect a user and change the authorization attributes of an existing user session.	
RADIUS DAC Port (default is 3799)	Port used for changing the authorization attributes of an existing user session. Default port is 3799.

Configuring AAA LDAP on a Controller

Lightweight Directory Access Protocol (LDAP) is an Internet protocol that e-mail and other programs use to look up information from a server.

TIP: You can also create LDAP Profiles and assign them to devices from Network Director—for directions, see [“Creating and Managing LDAP Profiles” on page 345](#).

Assign LDAP settings for the selected controller or controller cluster as described in [Table 235](#).

Table 235: Setting AAA LDAP Settings for Controllers

Field	Directions
Authentication Port (default is 389)	The default LDAP authentication port is 389. You can change the port number by using the up and down arrows.
Timeout (default is 5 seconds)	Adjust the length of time (default is 5 seconds) that elapses with no connection before Network Director gives an unreachable LDAP server error. You can change this value by using the up and down arrows to 1 through 90 seconds.
Dead Time (default is 5 seconds)	Indicate the number of seconds an LDAP server is unresponsive before it is marked as unavailable. Default is five seconds.
Bind Mode (default is simple bind)	When an LDAP session is created (LDAP client connects to a server) the authentication state of the session is set to anonymous. BIND mode establishes the authentication state for a session and sets the LDAP protocol version. The default is Simple bind —you can change this to SASL-MD5 . With Simple bind, the users’ credentials are sent to the LDAP Directory Service in clear text. With SASL-MD5 bind, the users’ credentials are encrypted.

Table 235: Setting AAA LDAP Settings for Controllers (*continued*)

Field	Directions
MAC Address Format (default is hyphens)	Indicate the format of the MAC address that will be sent as a password, either Hyphens , Colons , One hyphen , Raw , or None . None , means that the MAC address is stated in a single stream (for example, 12ae53ef5676), with no subgrouping of the numbers. Hyphens indicates hyphen separation, for example, 12-ae-53-ef-56-76), Colons indicates colon separation, for example, 12:ae:53:ef:56:76).
Base DN	The top level of the LDAP directory tree is the base, referred to as the <i>base DN</i> . Enter a base domain name, for example, DC=eng, DC=Juniper Networks, or DC=com. This string indicates where to load users and groups.
Prefix DN (default is cn)	AD or NT domains use the NetBIOS prefix domain name. Default is cn .

What To Do Next

The controller is now reconfigured, but the changes are not yet being used for operation. Deploy the controller by following the directions in [“Deploying Configuration to Devices” on page 1180](#).

RELATED DOCUMENTATION

[Deploying Configuration to Devices | 1180](#)

[Network Director Documentation home page](#)

Configuring Wireless Mobility and Network Domains

IN THIS CHAPTER

- [Understanding Mobility Domains | 1051](#)
- [Creating a Mobility Domain for Wireless LAN Controllers | 1053](#)
- [Configuring Security Settings for a Mobility Domain | 1056](#)
- [Creating Network Domains for Wireless LAN Controllers | 1058](#)

Understanding Mobility Domains

The Mobility Domain is a virtual grouping of Wireless LAN Controllers (WLCs) that enable roaming and maintain user session information to provide for seamless roaming through identity-based networking. It also provides resiliency and scalability. Common configuration is shared dynamically between all members of the cluster and provides a single point of configuration for all members.

A Mobility Domain enables users to roam geographically across the system while maintaining data sessions and VLAN or subnet membership, including IP address, regardless of connectivity to the network backbone. When users move from one area of a building or campus to another, client associations with servers or other resources do not change. When users access a controller in a Mobility Domain, they become members of the VLAN designated through their authorized identity. If a native VLAN is not present on the accessed controller, the controller forms a tunnel to another controller in the Mobility Domain that includes the native VLAN.

In a Mobility Domain, one controller acts as a primary seed device, and distributes information to all other controllers defined in the Mobility Domain. Otherwise, the seed controller operates as any other Mobility Domain member.

The clustering feature between the controllers ensures smooth mobility across an entire wireless network. With clustering, you can create logical groups of controllers and access points, which share network and user information for continuous and uninterrupted support.

RELATED DOCUMENTATION

Creating a Mobility Domain for Wireless LAN Controllers | **1053**

[Network Director Documentation home page](#)

Creating a Mobility Domain for Wireless LAN Controllers

A Mobility Domain enables users to roam geographically across the system while maintaining data sessions and VLAN or subnet membership, including IP address, regardless of connectivity to the network backbone. As users move from one area of a building or campus to another, client associations with servers or other resources remains the same.

The clustering functionality ensures mobility across an entire wireless network. With clustering, you can effortlessly create logical groups of controllers and access points, which share network and user information in a proactive manner for continuous and uninterrupted support.

You can create a mobility domain using the Create Mobility Domain window from the Network Director user interface.

A Mobility Domain group consists of a primary seed controller that contains the IP address of all controller members. A secondary controller might be optionally assigned. You must specify the primary seed while creating the mobility domain. A mobility domain cannot exist without the primary seed. You must explicitly configure only one controller per domain as the primary seed. All other controllers in the domain receive their Mobility Domain information from the seed.

You can optionally specify a secondary seed in a Mobility Domain. The secondary seed provides redundancy for the primary seed switch in the Mobility Domain. If the primary seed becomes unavailable, the secondary seed assumes the role of the seed controller. This allows the Mobility Domain to continue functioning if the primary seed becomes unavailable. When the primary seed fails, the remaining members form a Mobility Domain, with the secondary seed taking over as the primary seed controller.

TIP: Even though selecting a secondary seed for mobility domain is optional, Network Director does not assign a secondary seed by default.

Enabling cluster for the mobility domain enables cluster mode for both primary and secondary seed. To include the other domain controllers in cluster, set the cluster mode for the domain controllers using the Cluster Mode column in the Controllers in Mobility Domain table. Only after enabling cluster for the mobility domain can the cluster mode for the other mobility domain controllers be set.

1. Enter a name for the mobility domain. The name can contain alphanumeric characters (up to 32 character length) and all special characters except space.

This is a mandatory field.

2. To select a primary seed, click **Select** next to **Primary Seed**. The Select Primary Seed for Mobility Domain window is displayed. Select a controller from the list of host names.

This is a mandatory field.

- To select a secondary seed, click **Select** next to **Secondary Seed**. The Select Secondary Seed for Mobility Domain window is displayed. Select a controller from the list of host names.

This is an optional field. However, If you do not select a secondary seed, the domain will be nonoperational during failure.

To delete an already specified secondary seed, click the **Clear** button next to Secondary Seed.

- Select the **Enable Cluster** check box if you want to enable clustering for both primary and secondary seeds.

The *Cluster Mode* column in the Controllers in Mobility Domain table is displayed only if you select the Enable Cluster check box.

- To add controllers to the mobility domain, click **Add** from the Controllers in Mobility Domain table. The Select Controllers for Mobility Domain window is displayed. Select the controllers from the table. You can select more than one controller.

To enable cluster mode for the mobility domain members, select **enable** in the Cluster Mode column from the Controllers in the Mobility Domain table.

TIP: The table displays the list of available devices. The devices you had selected as the primary and secondary seeds will not be displayed again. Also, if you select a device and click **OK**, then the selected device is not displayed in the available Select Controllers for Mobility Domain window.

- Click **Done** to create the mobility domain with the specified configuration. The message: **Domain saved successfully** is displayed.
- If you want to remove any controllers from the mobility domain, select the controller from the Controllers in Mobility Domain table and click **Delete**.

You can edit or delete the mobility domain that you created.

To edit a mobility domain:

- Select **Wireless Network** from My Network in Logical View pane and select the mobility domain that you want to edit.

Network Director lists the mobility domains that you created earlier.

- Select **Edit Mobility Domain** from Domain Management in the Tasks pane. The Edit Mobility Domain page for the selected mobility domain is displayed.

3. Follow the tasks described above in the Creating Mobility Domain procedure.

To delete a mobility domain:

1. Select **Wireless Network** from My Network in Logical View pane and select the mobility domain that you want to delete.
2. Select **Delete Mobility Domain** from Domain Management in the Tasks pane. The following confirmation message is displayed:

```
If mobility domain has cluster please disable cluster for the domain and deploy the changes and then delete mobility domain. Do you want to continue?
```

3. Click **Yes** if you want to delete the mobility domain and the mobility domain doesn't contain a cluster. Click **No** if you want to delete the mobility domain but the domain contains a cluster.

To delete a mobility domain that contains a cluster:

- a. Select **Edit Mobility Domain** from Domain Management in the Tasks pane.
- b. Disable clustering by unselecting **Enable Cluster** and clicking **Done**.
- c. Deploy the changed configuration on the affected wireless LAN controllers in Deploy mode.
- d. Now delete the mobility domain by selecting **Delete Mobility Domain** again and click **Yes** in response to the confirmation message.

RELATED DOCUMENTATION

[Understanding Mobility Domains | 1051](#)

[Creating Network Domains for Wireless LAN Controllers | 1058](#)

[Network Director Documentation home page](#)

Configuring Security Settings for a Mobility Domain

You can enhance security on your network by enabling WLC-WLC security. WLC-WLC security encrypts management traffic exchanged by WLC switches in a Mobility Domain.

When WLC-WLC security is enabled, management traffic among WLC switches in the Mobility Domain is encrypted using AES. The keying material is dynamically generated for each session and passed among switches using configured public keys.

To configure security settings for a mobility domain:

1. From the View pane, select the mobility domain for which you want to configure the security settings.
2. From the Tasks pane, select **Domain Management > Manage WLC - WLC Security**. The Manage WLC - WLC Security page opens.
3. Do one of the following:
 - Select **Enable** from the Security Mode list to enable security for the mobility domain.
 - Select **Disable** to disable security for the mobility domain. This is the default mode.
4. If you select to enable security, Network Director displays a table listing all the WLCs in the selected mobility domain. [Table 236](#) lists the fields that are displayed in the Manage WLC - WLC Security page.

Table 236: Wireless LAN Controller Details

Field	Description
WLC Controller	Host name of the WLC.
Role	The role of the WLC in the mobility domain. A WLC can be a primary controller, a secondary controller, or a member.
Public Key	The public key of the WLC.
Last Fetched Time	The time when the public key was last obtained from the WLC.

5. If you want to use the public keys from the WLC, do the following:
 - a. On the Mobility System Software (MSS™) command line interface (CLI) on each of the domain controller, run the **crypto generate key domain 128** command. This command generates an RSA public-private encryption key pair that is required for a Certificate Signing Request (CSR) or a self-signed certificate.

- b. Click **Retrieve Keys from Controllers** in the Manage WLC - WLC Security page. Network Director retrieves the public keys for all the domain controller WLCs and lists them in the Public Key column.
6. Click **Done** to save the changes that you made to the security settings.

RELATED DOCUMENTATION

[Understanding Mobility Domains | 1051](#)

[Creating a Mobility Domain for Wireless LAN Controllers | 1053](#)

[Network Director Documentation home page](#)

Creating Network Domains for Wireless LAN Controllers

A Network Domain enables the functionality of Mobility Domains to be extended over a multiple-site installation. In a Network Domain, one or more controllers act as a seed device. A Network Domain seed stores information about all of the VLANs on the Network Domain members. The Network Domain seeds share this information to create an identical database on each seed.

Each Network Domain member maintains a TCP connection to one of the seeds. When a Network Domain member needs information about a VLAN in a remote Mobility Domain, the member consults a connected Network Domain seed. If the seed has information about the remote VLAN, it responds with the controller's IP address to the VLAN. A VLAN tunnel is then created between the seed controller and the remote controller.

When there are multiple Network Domain seeds in an installation, a Network Domain member connects to the seed with the highest configured affinity. If that seed is unavailable, the Network Domain member connects to the seed with the next highest affinity.

You can configure multiple seeds in a Network Domain. If you configure multiple Network Domain seeds, a member consults the seed with the highest configured affinity. To configure multiple seeds in the same Network Domain (for example, a seed on each physical site in the Network Domain), you must establish a peer relationship among the seeds.

If you configure multiple controllers as seed devices in a Network Domain, the controllers establish a peer relationship to share information about the VLANs configured on the member devices of the network to create identical VLAN databases. Sharing information in this way provides redundancy in case one of the seed peers becomes unavailable.

In a Network Domain, a member controller consults a seed controller to determine a user VLAN membership in a remote Mobility Domain. You can create a network domain using the Create Network Domain page from the Network Director user interface.

1. Enter a name for the network domain. The name can contain up to 16 characters including alphanumeric values and all special characters except space.
2. Add controllers in the network domain. To add controllers:
 - Click **Add** from Controllers in Network Domain table. The Select Controllers for Network Domain page appears.
 - Select the devices or controllers that you want add to the network domain by selecting the check boxes. You can select more than one device.
 - Click **OK**. The selected devices appear in the Controllers in Network Domain table.
3. Select the role for each device or controller from the **Role** column in the Controllers in Network Domain table. The available roles are:

- Seed
- Member
- Seed and Member

TIP: You must select at least one device as the seed for the network domain. In a Network Domain, at least one seed device must be aware of each member device. The seed is required to maintain an active TCP connection with each member controllers. To configure a controller as a member of a Network Domain, you must specify one or more Network Domain seeds.

4. To delete a controller from the Controllers in Network Domain table, select the controller by clicking the check box and click **Delete**.

You can edit or delete the network domain that you created.

To edit a network domain:

1. Select **Wireless Network** from My Network in Logical View pane and select the network domain that you want to edit.

Network Director lists the network domains that you created earlier.

2. Select **Edit Network Domain** from Domain Management in the Tasks pane. The Edit Network Domain page for the selected network domain is displayed.

3. Follow the tasks described above in the Creating Network Domain procedure.

To delete a network domain:

1. Select **Wireless Network** from My Network in Logical View pane and select the network domain that you want to delete.

Network Director lists the mobility domains that you created earlier.

2. Select **Delete Network Domain** from Domain Management in the Tasks pane. The message: **Deleting domain "domain-name". Do you want to continue?** is displayed.

3. Click **Yes** if want to delete the network domain.

[RELATED DOCUMENTATION](#)

[Understanding Mobility Domains | 1051](#)

[Creating a Mobility Domain for Wireless LAN Controllers | 1053](#)

[Network Director Documentation home page](#)

Configuring WLAN Service (SSIDs)

IN THIS CHAPTER

- Understanding the Network Terms SSID, BSSID, and ESSID | 1061
- Understanding Network Director SSID Configuration Using Profiles | 1064
- Understanding Call Admission Control | 1068
- Understanding Load Balancing for Wireless Radios | 1070
- Understanding Wireless Encryption and Ciphers | 1072
- Understanding the IEEE 802.11 Standard for Wireless Networks | 1076
- Understanding PSK Authentication | 1079
- Understanding WLAN Service Profiles | 1081
- Creating and Managing a WLAN Service Profile | 1090
- Understanding Voice Clients and Voice Traffic | 1120
- Configuring a Voice SSID with Network Director | 1124
- Creating and Managing RF Snooping Filter Profiles | 1125
- Assigning RF Snooping Filter Profiles to Access Points | 1132

Understanding the Network Terms SSID, BSSID, and ESSID

IN THIS SECTION

- An SSID is the Name of a Network | 1062
- BSSIDs Identify Access Points and Their Clients | 1063
- An ESS Consists of BSSs | 1064

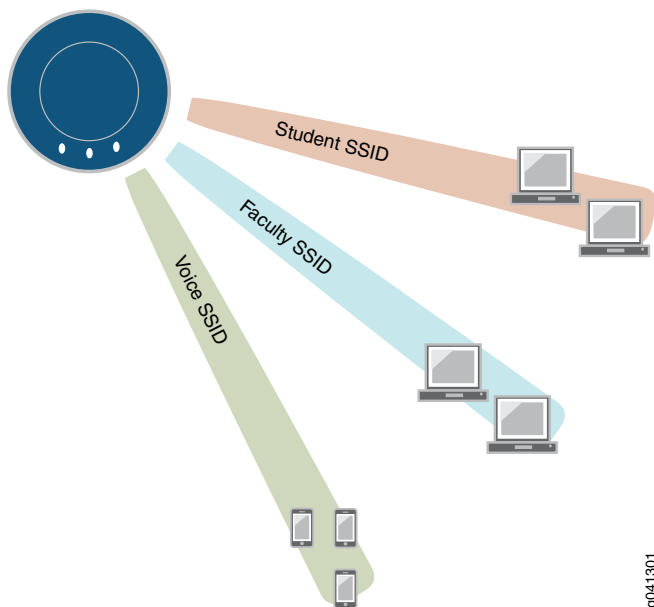
The terms BSSID, ESSID, and SSID are all used to describe sections of a wireless network (WLAN)—the three terms have slightly different meanings. As a wireless user you are concerned only with the broadcast SSIDs that let you connect to a wireless network. As an administrator, you also need to keep track of BSSIDs and, to a lesser degree, ESSIDs.

This topic describes:

An SSID is the Name of a Network

Because multiple WLANs can coexist in one airspace, each WLAN needs a unique name—this name is the service set ID (SSID) of the network. Your wireless device can see the SSIDs for all available networks—therefore, when you click a wireless icon, the SSIDs recognized by device are listed. For example, suppose your wireless list consists of three SSIDs named Student, Faculty, and Voice. This means that an administrator has created three WLAN Service profiles and, as part of each WLAN service profile, provided the SSID name Student, Faculty, or Voice. (For directions to create a WLAN Service profile, see [“Creating and Managing a WLAN Service Profile”](#) on page 1090.)

Figure 43: Radios can have up to 32 SSIDs



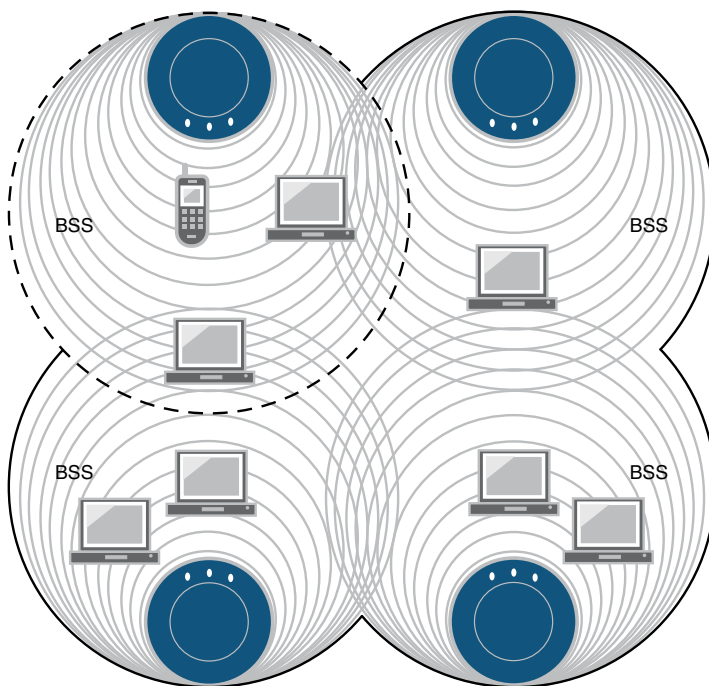
As a WLAN user, you are concerned only with the SSIDs. You select one from the list on your laptop or other device, provide your username and a password, and use the SSID. You might not have access to all SSIDs—the authentication and access privileges are usually different for different WLANs and their associated SSIDs.

BSSIDs Identify Access Points and Their Clients

Packets bound for devices within the WLAN need to go to the correct destination. The SSID keeps the packets within the correct WLAN, even when overlapping WLANs are present. However, there are usually multiple access points within each WLAN, and there has to be a way to identify those access points and their associated clients. This identifier is called a basic service set identifier (BSSID) and is included in all wireless packets.

Figure 44: Each Access Point has its Own BSS

BSS+BSS+BSS+BSS=ESS



BSSID = AP MAC address
SSID = name of network

9041300

As a user, you are usually unaware of which basic service set (BSS) you currently belong to. When you physically move your laptop from one room to another, the BSS you use can change because you moved from the area covered by one access point to the area covered by another access point, but this does not affect the connectivity of your laptop.

As an administrator, you are interested in the activity within each BSS. This tells you what areas of the network might be overloaded, and it helps you locate a particular client. By convention, an access point's MAC address is used as the ID of a BSS (BSSID). Therefore, if you know the MAC address, you know the BSSID—and, because all packets contain the originator's BSSID, you can trace a packet. This works fine for an access point with one radio and one WLAN configured.

Most often, there are different BSSIDs on an access point for each WLAN configured on a radio. If you have an access point with 2 radios and 32 WLANs configured on each, you would have 64 BSSIDs plus the base access point BSSID. To accommodate the multiple BSSIDs, each access point is assigned a unique block of 64 MAC addresses. Each radio has 32 MAC addresses and supports up to 32 service set identifiers (SSIDs), with one MAC address assigned to each SSID as a basic service set identification (BSSID). All MAC addresses for an access point are assigned based on the base MAC address of the access point.

NOTE: The access point MAC address block is listed on a label on the back of the access point.

To view a list of SSIDs for a network, look at the list of WLAN Service Profiles in Network Director.

Ad-Hoc Networks Do Not Have a MAC Address

Every BSS needs a BSSID, and using the access point's MAC address works fine most of the time. However, an ad-hoc network, a network that forwards traffic from node to node, has no access point. When a BSS does not have a physical access point, in an ad-hoc network for example, the network generates a 48-bit string of numbers that looks and functions just like a MAC address, and that BSSID goes in every packet.

An ESS Consists of BSSs

An extended basic service set (ESS) consists of all of the BSSs in the network. For all practical purposes, the ESSID identifies the same network as the SSID does. The term SSID is used most often.

RELATED DOCUMENTATION

[Creating and Managing a WLAN Service Profile | 1090](#)

[Network Director Documentation home page](#)

Understanding Network Director SSID Configuration Using Profiles

IN THIS SECTION

- [Configuring an SSID with Network Director | 1065](#)
- [SSID Access Information | 1066](#)
- [SSID Authentication Information | 1067](#)
- [SSID WLAN Information | 1067](#)

- SSID Radio Information | 1067
- What Do I Do When all of Profiles are Complete? | 1067

Service set identifiers (SSIDs) uniquely identify a set of profile configurations with names up to 32 characters long. SSIDs are broadcast to clients by access points—the SSID name appears in the list of available wireless networks on clients' laptops or other devices. For a wireless user, connecting to an SSID is as easy as selecting a name from the list. However, creating that SSID involves more steps.

This topic describes:

Configuring an SSID with Network Director

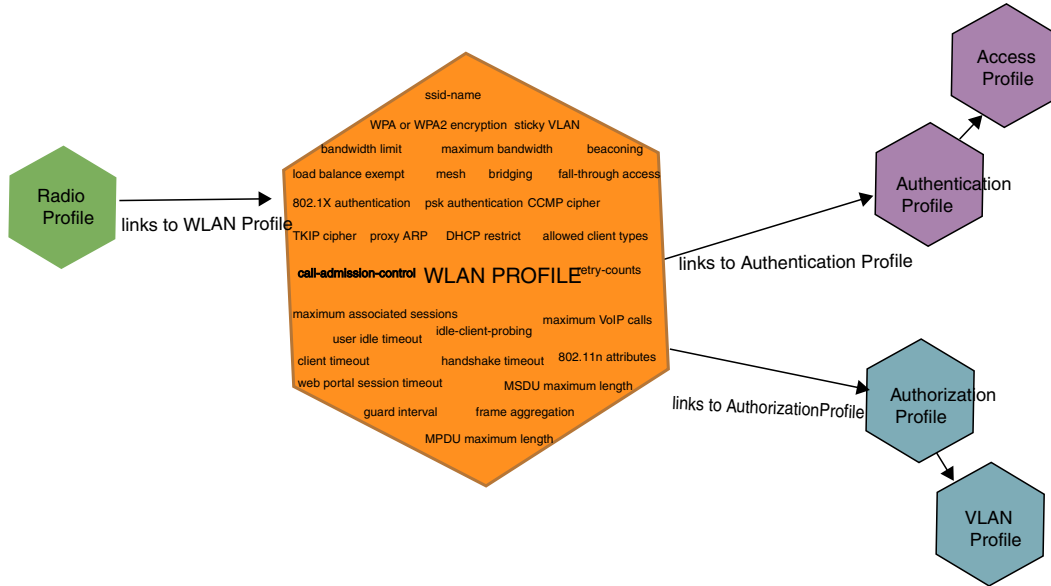
Configuring an SSID is more complicated than connecting to an SSID. You can think of the SSID as a number of puzzle pieces (called profiles) that work together to form the SSID. You typically associate all of the following profiles with each SSID:

- VLAN profile
- Access information in an Access profile
- Authentication information in an Authentication profile
- Authorization information in an Authorization profile
- WLAN configuration in a WLAN Service profile
- Radio configuration in a Radio profile

Figure 45 illustrates the connections between profiles.


NOTE: Some profiles are optional, depending on your configuration.

Figure 45: SSID Information is Located in These Profiles



Another way to think of the relationship between profiles is shown in Figure 46. You build the needed profiles, starting with profiles at the far right of the chart and working your way to the left. When the Radio profile is assigned to radios discovered by the controller, the profile configuration is complete.

Figure 46: Relationship Between Profiles Needed to Create an SSID



SSID Access Information

Access profiles include details about authentication methods and accounting servers used by an SSID. Radius, LDAP, and local authentication are the supported authentication methods and RADIUS and LDAP are the supported accounting methods. For more information about Access profiles, see “Understanding Access Profiles” on page 351 or “Creating and Managing Access Profiles” on page 352.

SSID Authentication Information

Authentication profiles include the authentication method and authentication parameters used for SSID client authentication. Available authentication methods are 802.1X (dot1x), MAC-RADIUS, Captive Portal, and Last Resort. 802.1X is the default authentication method for all device types. For more information about Authentication profiles, see [“Understanding Authentication Profiles” on page 381](#). To create, edit, or delete an Authentication profile, see [“Creating and Managing Authentication Profiles” on page 383](#).

SSID WLAN Information

You name SSIDs in the WLAN Service profile, in addition to configuring the encryption method (WPA or WPA2), and beaconing. Typically, you also indicate VLAN use, SSID bandwidth limits, radio load balancing, proxy ARP, DHCP restrictions, and permitted client types.

Call admission control, retry counts, guard interval, frame aggregation, MSDU, MPDU, client bandwidth limits, MTU, idle client probing, data rates, and WMM Power Save are also configured in the WLAN Service profile.

For more information, see [“Understanding WLAN Service Profiles” on page 885](#) or [“Creating and Managing a WLAN Service Profile” on page 1090](#).

SSID Radio Information

Radio-profiles are used to form groups of radios and define which wireless services are broadcast by the group. Additional radio related parameters are also configured at this level.

For more information, see [“Understanding Radio Profiles” on page 879](#) or [“Creating and Managing a Radio Profile” on page 932](#).

What Do I Do When all of Profiles are Complete?

If you are using Network Director, radios only actually broadcast an SSID after you assign a Radio profile to a controller (see [“Assigning a Radio Profile to Radios” on page 952](#), then deploy the controller configuration (see [“Deploying Configuration to Devices” on page 1180](#).) At this time, all associated profiles also go into effect.

RELATED DOCUMENTATION

[Understanding Access Profiles | 351](#)

[Creating and Managing Access Profiles | 352](#)

[Understanding Radio Profiles | 879](#)

[Creating and Managing a Radio Profile | 932](#)

[Understanding WLAN Service Profiles | 885](#)

[Creating and Managing a WLAN Service Profile | 1090](#)

[Understanding the Network Terms SSID, BSSID, and ESSID | 1061](#)

[Understanding Authentication Profiles | 381](#)

[Creating and Managing Authentication Profiles | 383](#)

[Deploying Configuration to Devices | 1180](#)

[Network Director Documentation home page](#)

Understanding Call Admission Control

IN THIS SECTION

- [What Radios Does Call Admission Control Affect? | 1069](#)
- [How Do I Configure Call Admission Control? | 1070](#)

Call admission control (CAC) regulates the addition of new voice over IP (VoIP) sessions on access point radios, guaranteeing a higher quality of service to a fixed number of voice clients by limiting the number of calls. Call admission control is configured in the WLAN Service profile using either Wi-Fi Multimedia (WMM) or, for older equipment, the legacy SpectraLink Voice Priority (SVP). CAC is used in the call setup step and applies only to real-time media traffic as opposed to data traffic.

NOTE: SVP is a legacy technology—Spectralink's new phones use WMM.

You can turn off call admission control (set it to none), or you can enable either session-based admission control or VoIP-based admission control:

- **None** (default)
- **Session-based:** Session-based call admission control limits the number of sessions (14 by default) on one WLAN Service profile. When used in conjunction with QoS, a small number of clients (0 through 100) are given priority access, while clients on other WLAN Service profiles are forced to use best effort. Configure session-based call admission control when an SSID will handle both data and voice.
- **Voip-call:** Call admission control limits only the number of voice sessions on one WLAN Service profile. VoIP-call is designed to be used with a voice-only Radio profile and SSID.

NOTE: We recommend creating voice-only SSIDs, either WMM VoIP-call or SVP VoIP call, for best performance.

This topic describes:

What Radios Does Call Admission Control Affect?

Call admission control is enabled in a WLAN Service profile and affects radios using the Radio profile associated with the WLAN Service profile (SSID).

NOTE: To ensure voice quality, do not map other WLAN Service profiles to a Radio profile you plan to use for dedicated voice traffic or use QoS to force all packets on other WLAN Service profiles to use best-effort priority.

How Do I Configure Call Admission Control?

Call admission control is enabled in a WLAN Service profile (SSID) and configured in the associated Radio profile.

For information about enabling call admission control, see [“Creating and Managing a WLAN Service Profile” on page 1090](#). For information about configuring call admission control, see [“Creating and Managing a Radio Profile” on page 932](#).

RELATED DOCUMENTATION

[Creating and Managing a WLAN Service Profile | 1090](#)

[Creating and Managing a Radio Profile | 932](#)

[Network Director Documentation home page](#)

Understanding Load Balancing for Wireless Radios

IN THIS SECTION

- [Why Would I Need Load Balancing? | 1071](#)
- [When Would I Avoid Load Balancing? | 1071](#)
- [How Load Balancing Works | 1071](#)
- [Can I Group Access Points for Load Balancing? | 1071](#)
- [Where Do I Configure Load Balancing in Network Director? | 1072](#)

Load balancing distributes a workload across multiple entities, in this case wireless radios, to achieve optimal utilization, maximize throughput, minimize response time, and avoid overload.

RF load-balancing for access points has the ability to reduce network congestion over an area by distributing client sessions across access point radios with overlapping coverage. With load-balancing, you can ensure that all access points on the network handle a proportionate share of wireless traffic, and that no single access point gets overloaded. Load balancing of access points is enabled by default in WLAN Service profiles—that means that all access points using a WLAN Service profile are load-balanced.

This topic describes:

Why Would I Need Load Balancing?

As new clients arrive, load-balancing distributes the clients among access points such that the access points share the client load. The wireless operating system, MSS, encourages clients to associate with the least loaded (by client count) access point, so that clients are well distributed across access points. This way, no one access point is more overloaded and there is less interruption of wireless services on the network. For example, in an auditorium or lecture hall, there might be a large number of clients in a relatively small amount of space. While a single access point might be sufficient for providing an RF signal to the entire area, more access points might be required to deliver enough aggregate bandwidth for all of the clients. When additional access points are installed in the room, RF load-balancing spreads clients evenly across the access points, increasing the available aggregate bandwidth by increasing the number of access points. Without load-balancing, you could have multiple access points in a classroom with all clients associated to just one of the access points.

When Would I Avoid Load Balancing?

We do not recommend load-balancing for low-latency applications such as voice or live-streaming (unbuffered) video. Load balancing is not advisable for voice transmission because load-balancing can impact roam times, which can impact voice quality for roaming clients. Load balancing can also make streaming video jittery with dropped frames.

How Load Balancing Works

The wireless operating system, MSS, balances the client load by adjusting how access points are perceived by clients advertising the same SSID. The capacity of an access point handling new clients is compared to other access points in the SSID. As new clients arrive, MSS encourages them to associate with the least loaded (by client count) access points, such that clients are well distributed across access points. By default, MSS only encourages clients to associate with an access point if there are access points available with capacity to accept more clients. Clients are never prevented from associating with an access point if it is the only one available.

Can I Group Access Points for Load Balancing?

You can optionally place access point radios into load-balancing groups. When two or more access point radios are placed in the same load-balancing group, MSS assumes that they have exactly the same coverage area, and attempts to distribute the client load across them equally. The AP radios do not have to be on the same controller. A balanced set of AP radios can span multiple controllers in a mobility domain. When you have grouped APs for load-balancing, you can also indicate how strictly the balancing is enforced, low, medium, high, and maximum enforcement.

Where Do I Configure Load Balancing in Network Director?

Load balancing is enabled and configured in multiple places for several kinds of load-balancing. The load-balancing option discussed here applies to wireless access points and radios.

NOTE: Load balancing configured in Access profiles applies to the RADIUS servers configured within that profile. For directions, see [“Creating and Managing Access Profiles” on page 352](#).

RF load-balancing for access points is enabled by default in Juniper Networks WLAN mobility domains, affecting all radios within that domain. You can, however, disable load-balancing in a WLAN Service profile. If you want a WLAN’s SSID on access points to be exempt from load-balancing, you can indicate that a WLAN Service profile is *Load Balancing Exempt*—see the directions for Network Director [“Creating and Managing a WLAN Service Profile” on page 1090](#).

You can assign a single access point to a specific load-balance group when you add that access point to a controller—see [“Adding and Managing an Individual Access Point” on page 1156](#). You can also configure load-balancing on a controller—see [“Configuring a Controller” on page 1037](#).

RELATED DOCUMENTATION

[Creating and Managing a WLAN Service Profile | 1090](#)

[Understanding Access Profiles | 351](#)

[Creating and Managing Access Profiles | 352](#)

[Configuring a Controller | 1037](#)

[Network Director Documentation home page](#)

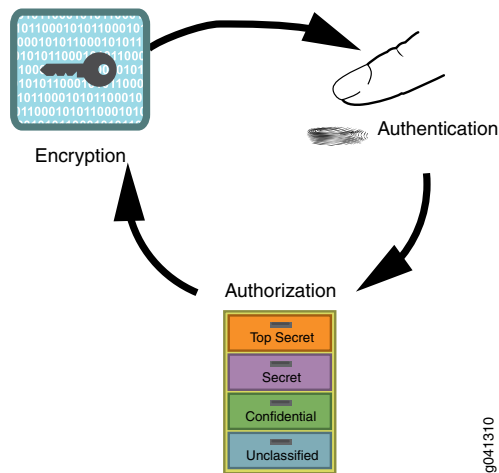
Understanding Wireless Encryption and Ciphers

IN THIS SECTION

- [Wired Equivalent Privacy \(WEP\) was the Original Wireless Encryption | 1073](#)
- [WPA Encryption Replaced WEP | 1073](#)
- [WPA2 Is the Strongest Encryption Available | 1074](#)
- [Security Ciphers for WPA and WPA2 | 1074](#)
- [Which Encryption Method Should I Use? | 1075](#)

Wireless network security relies on a combination of encryption, authentication, and authorization to provide maximum protection for a WLAN. Encryption is focused on protecting the information within a session, reading information in a data stream and altering it to make it unreadable to users outside the network. This topic discusses encryption.

Figure 47: Network Security Is Provided by Encryption, Authentication, and Authorization



Juniper Networks access points support all three standard types of wireless access point-client encryption: the legacy encryption Wired Equivalent Privacy (WEP), Wi-Fi Protected Access (WPA), and WPA2 (also called RSN). Encryption type is configured in WLAN Service profiles under the Security Settings tab. For information about applying encryption, see [“Creating and Managing a WLAN Service Profile” on page 1090](#).

This topic describes:

Wired Equivalent Privacy (WEP) was the Original Wireless Encryption

WEP was the original security algorithm for IEEE 802.11 wireless networks, introduced as part of the original 802.11 standard.

WPA Encryption Replaced WEP

WPA addressed the vulnerabilities of WEP, the original, less secure 40 or 104-bit encryption scheme in the IEEE 802.11 standard. WPA also provides user authentication—WEP lacks any means of authentication.

WPA replaced WEP with a stronger encryption technology called Temporal Key Integrity Protocol (TKIP) with Message Integrity Check (MIC). It also provides a scheme of mutual authentication using either IEEE 802.1X/Extensible Authentication Protocol (EAP) authentication or pre-shared key (PSK) technology.

NOTE: You can simultaneously apply both WPA and WPA2 to an SSID. Clients use WPA2 if they have the capability—otherwise the client uses WPA. WPA2 is recommended unless you need to provide access to for legacy devices. All 802.11n devices support WPA2.

WPA2 Is the Strongest Encryption Available

WPA2 is the certified version of the full IEEE 802.11i specification. Like WPA, WPA2 supports either IEEE 802.1X/EAP authentication or PSK technology. It also includes a new advanced encryption mechanism using the Counter-Mode/CBC-MAC Protocol (CCMP) called the Advanced Encryption Standard (AES).

WPA was based on the 802.11i draft, while WPA2 is based on 802.11i final standard. Where WPA encryption was specifically designed to work with some wireless hardware that supported WEP, WPA2 offers stronger security but is not supported by earlier hardware designed for WEP.

NOTE: The Wi-Fi Alliance requires that high-throughput (802.11n) transmissions use WPA2 and CCMP. You can simultaneously apply both WPA and WPA2 to an SSID. Clients use WPA2 if they have the capability—otherwise the client uses WPA.

Security Ciphers for WPA and WPA2

Standard security ciphers are part of both WPA and WPA2 encryption. You choose whether you want to apply either the newer CCMP, or TKIP (an upgrade of original WEP programming), or both for each WLAN Service profile. Both cipher suites dynamically generate unique session keys for each session and periodically change the keys to reduce the likelihood of a network intruder intercepting enough frames to decode a key. The two available ciphers are:

- *Counter Mode with Cipher Block Chaining Message Authentication Code Protocol (CCMP)*—CCMP provides Advanced Encryption Standard (AES) data encryption for WPA and WPA2. To provide message integrity, CCMP also uses the Cipher Block Chaining Message Authentication Code (CBC-MAC).

A radio using WPA/WPA2 with CCMP encrypts traffic for only WPA CCMP clients but not for TKIP clients. The radio disassociates from TKIP clients unless you selected both CCMP and TKIP.

NOTE: The Wi-Fi Alliance requires that high-throughput (802.11n) transmissions use WPA2 and CCMP.

- *Temporal Key Integrity Protocol (TKIP)*—TKIP uses the RC4 encryption algorithm, a 128-bit encryption key, a 48-bit initialization vector (IV), and a message integrity code (MIC). A radio using WPA/WPA2 with TKIP encrypts traffic for only WPA TKIP clients but not for CCMP clients. The radio disassociates from CCMP clients unless you selected both CCMP and TKIP.

TKIP is most useful for upgrading security on devices originally using WEP — it does not address all of the security issues facing WLANs and may not be reliable or efficient enough for sensitive corporate and government data transmission. The 802.11i standard specifies the Advanced Encryption Standard (AES) in addition to TKIP. AES is an additional cipher stream that adds a higher level of security and is approved for government use.

NOTE: TKIP is not permitted for 802.11n-based transmissions. It is only supported for legacy (802.11b, 802.11g and 802.11a) transmissions, which are limited to a maximum of 54 Mbps.

Which Encryption Method Should I Use?

WPA2 is the most secure encryption method available for wireless networks—we recommend using WPA2 with the CCMP cipher whenever possible. WPA2 with CCMP is the only option permitted for high throughput 802.11n transmissions. Eventually, WPA encryption with TKIP will be obsolete as you replace older devices that use only TKIP.

If you need to accommodate legacy devices with an SSID, enable WPA encryption with the TKIP cipher. Keep in mind that this has an effect on performance. The additional AES cipher takes more computing power to run than simple TKIP does, therefore older, smaller devices may not support it.

NOTE: You can create different WLAN Service profiles (SSIDs) for different levels of encryption. This maximizes the use of WPA2 security.

Security always affects performance, so it is really up to you how much bandwidth and processing time you want to devote to it. With newer devices, this is much less of an issue because new devices have plenty of resources for the highest level of security, WPA2 with CCMP.

RELATED DOCUMENTATION

[Creating and Managing a WLAN Service Profile | 1090](#)

[Network Director Documentation home page](#)

Understanding the IEEE 802.11 Standard for Wireless Networks

IN THIS SECTION

- [Differences Between 802.11 Standards | 1076](#)
- [802.11 Divides Each Frequency Band into Channels | 1077](#)
- [What Is 802.11i Security? | 1078](#)
- [What Is 802.11n? | 1078](#)

The IEEE 802.11 standard consists of a series of technological advances that have been developed over many years. Each new advancement is defined by an amendment to the standard that is identified by a one or two letter suffix to "802.11." The original 802.11 standard allowed up to 2 Mbps on only the 2.4-GHz band. 802.11b added new coding schemes to increase throughput to 6 Mbps. 802.11a added support on the 5-GHz band and Orthogonal Frequency Division Multiplexing (OFDM) coding schemes that increase throughput to 54 Mbps. 802.11g brought OFDM from 802.11a to the 2.4-GHz band. 802.11n added an assortment of high throughput advances to increase throughput roughly 10 times, such that high-end enterprise access points achieve signaling throughputs of 450 Mbps. The emerging 802.11ac standard promises to exceed 1 Gbps of throughput. The individual standards in use now are 802.11a, 802.11b, 802.11g, and 802.11n (which uses a more advanced technology than the others). The newest standard, 802.11ac, is the newest and fastest standard.

The segment of the radio frequency spectrum used by 802.11 varies between countries.

This topic describes:

Differences Between 802.11 Standards

The newer the 802.11 standard, the faster it is and the greater its capacity. The new 802.11ac specification will eventually enable multi-station WLAN throughput of 1 gigabit per second. [Table 237](#) lists the differences between current 802.11 standards. The draft 802.11ac estimates are given in the last row of the table for comparison.

Table 237: Differences Between 802.11 Protocols

802.11 Protocol	Frequency Band Used	Bandwidth	Data Rate per Stream
n Up to 4 streams of data	2.4-GHz 5-GHz	20 MHz	7.2 mbps 14.4 mbps, 21.7 mbps 28.9 mbps, 43.3 mbps 57.8 mbps, 65 mbps, 72.2 mbps
		40 MHz	15 mbps, 30 mbps, 45 mbps, 60 mbps, 90 mbps, 120 mbps, 135 mbps, 150 mbps
g 1 stream of data	2.4-GHz	20 MHz	6, 9, 12, 18, 24, 36, 48, 54
b 1 stream of data	2.4-GHz	20 MHz	1 mbps, 2 mbps, 5.5 mbps, 11 mbps
a 1 stream of data	5-GHz 3.7-GHz	20 MHz	6 mbps, 9 mbps, 12 mbps, 18 mbps, 24 mbps, 36 mbps, 48 mbps, 54 mbps
ac (draft) Up to 8 streams of data	5- GHz	20 MHz	up to 87.6 mbps
		40 MHz	up to 200 mbps
		60 MHz	up to 433.3 mbps
		80 MHz	up to 866.7 mbps

NOTE: 802.11ng and 802.11na are Juniper Networks terminology and not part of the 802.11 standard. It is simply Juniper Networks notation for indicating 802.11n use on the 2.4-GHz band (11ng) or 802.11n use on the 5-GHz band (11na).

802.11 Divides Each Frequency Band into Channels

802.11 divides each of the frequency bands listed in [Table 237](#) into channels.

802.11 divides each frequency band into channels in a different way. For example the 2.4000-2.4835-GHz band is divided into 13 channels spaced 5 MHz apart. Channels 1, 6, and 11 were originally the only non-overlapping channels, but with the newer 802.11g standard there are now four non-overlapping channels—1, 5, 9 and 13. (There are now four because the orthogonal frequency-division multiplexing (OFDM) modulated 802.11g channels are 20 MHz wide.)

NOTE: Many countries, including US allow use of channels 1 - 11 only.

The amount of available spectrum for unlicensed use, which varies by country, in the 5-GHz band is much greater, and typical supports many more channels than in the 2.4-GHz band.

Availability of channels is regulated by country and can change. Japan permits the use of 14 channels in the 2.4-GHz band, while other countries such as Spain initially allowed the use of only channels 10 and 11. Europe and Asia now allow channels 1 through 13. North America and some Central and South American countries allow only channels 1 through 11.

What Is 802.11i Security?

Wi-Fi Protected Access (WPA) and Wi-Fi Protected Access II (WPA2) are two security standards and security certification programs developed by the Wi-Fi Alliance to address security issues found in Wired Equivalent Privacy (WEP). Where WPA encryption was specifically designed to work with some wireless devices that support WEP, while WPA2 encryption does not work on any device that supports only WEP. For more information about encryption, see [“Understanding Wireless Encryption and Ciphers” on page 899](#).

What Is 802.1X?

802.1X is an authentication protocol supported by the 802.11 standards that enables mobile devices to be authenticated by username and password or by various types of credentials such as an X.509 certificate, or SIM in cellular phones. 802.1X authentication works in conjunction with an AAA server (typically RADIUS) that provides centralized authentication and user management.

RELATED DOCUMENTATION

[Understanding Wireless Encryption and Ciphers | 899](#)

[Creating and Managing Authentication Profiles | 383](#)

[Understanding Authentication Profiles | 381](#)

[Network Director Documentation home page](#)

Understanding PSK Authentication

IN THIS SECTION

- [What Is PSK? | 1079](#)
- [How Does PSK Work? | 1079](#)
- [When Would I Use PSK Authentication? | 1080](#)
- [Why Would I not Use PSK Authentication? | 1080](#)
- [How Is WPA Encryption Different from WPA-PSK Encryption? | 1081](#)

Pre-Shared Key (PSK) is a client authentication method that uses a string of 64 hexadecimal digits, or as a passphrase of 8 to 63 printable ASCII characters, to generate unique encryption keys for each wireless client. PSK is one of two available authentication methods used for WPA and WPA2 encryption on Juniper Networks wireless networks. PSK is not the default authentication method when creating a WLAN Service profile because the other choice, 802.1X authentication, is the standard and is stronger.

NOTE: 802.1X and PSK authentication types can be applied simultaneously—clients will use the most secure option that they are capable of using. For more information about 802.1X authentication, see [“Understanding the IEEE 802.11 Standard for Wireless Networks”](#) on page 1076.

This topic describes:

What Is PSK?

There are two WPA forms of encryption available with Network Director: Wi-Fi Protected Access (WPA) and the newer WPA2. Pre-shared key (PSK), a shared secret method, can be added to either encryption method:

- WPA/WPA2 Enterprise (requires a RADIUS server) and provides coverage for large entities.
- WPA/WPA2 Personal (also known as WPA-PSK) is appropriate for use in most residential and small business settings.

How Does PSK Work?

With PSK, you configure each WLAN node (access points, wireless routers, client adapters, bridges) not with an encryption key, but rather with a string of 64 hexadecimal digits, or as a passphrase of 8 to 63

printable ASCII characters. Using a technology called TKIP (Temporal Key Integrity Protocol), that passphrase, along with the network SSID, is used to generate unique encryption keys for each wireless client. Those encryption keys are constantly changed. When clients connect, the PSK authentication users provide the password to verify whether to allow them access to a network. As long as the passwords match, a client is granted access to the WLAN.

NOTE: You have the option to encrypt the PSK plain-English passphrase.

When Would I Use PSK Authentication?

PSK was designed for home and small office networks that do not require the complexity of an 802.1X authentication server. Some reasons to use PSK authentication are:

- PSK is simple to implement, as opposed to 802.1X authentication, which requires a RADIUS server.
- Your legacy clients might not support 802.1X or the latest WPA2 standard. You can use both WPA/WPA2 and PSK simultaneously to accommodate all clients.

Why Would I not Use PSK Authentication?

Even if you have a small company, there are drawbacks to using PSK authentication. For example:

- If an administrator leaves the company, you should reset the PSK key. This can become tiresome and be skipped.
- If one user is compromised, then all users can be hacked.
- PSK cannot perform machine authentication the way that IEEE 802.1X authentication can.
- Keys tend to become old because they are not dynamically created for users upon login, nor are the keys rotated frequently. You must remember to change the keys and create keys long enough to be a challenge to hackers. PSK is subject to brute force key space search attacks and to dictionary attacks.
- Because WPA2-Personal uses a more advanced encryption type, additional processing power is required to keep the network functioning at full speed. Wireless networks that use legacy hardware for access points and routers can suffer speed reductions when WPA2-Personal is used instead of WPA, especially when several users are connected or a large amount of data is moving through the network. Because WPA2-Personal is a newer standard, firmware upgrades can also be required for some hardware that previously used WPA exclusively.

How Is WPA Encryption Different from WPA-PSK Encryption?

The primary difference between WPA and WPA2-Personal are the encryption ciphers used to secure the network. WPA can use only the encryption cipher Temporal Key Integrity Protocol (TKIP). WPA2-Personal can use TKIP, but because TKIP security keys are less secure, the WPA2 protocol usually uses the Advanced Encryption Standard. AES uses a much more advanced encryption algorithm that cannot be defeated by the tools that overcome TKIP security, making it a much more secure encryption method.

RELATED DOCUMENTATION

[Understanding the IEEE 802.11 Standard for Wireless Networks | 1076](#)

[Understanding WLAN Service Profiles | 885](#)

[Creating and Managing a WLAN Service Profile | 1090](#)

[Network Director Documentation home page](#)

Understanding WLAN Service Profiles

IN THIS SECTION

- [SSID | 1082](#)
- [Mapping WLAN Service Profiles to Additional Profiles | 1083](#)
- [SSID Encryption | 1083](#)
- [Associated Authentication Profile | 1084](#)
- [Associated Authorization Profile | 1084](#)
- [VLAN Use | 1084](#)
- [Bandwidth Limit for Client Sessions | 1084](#)
- [Load Balancing Between Access Points | 1084](#)
- [Using Proxy ARP | 1084](#)
- [Restricting DHCP | 1085](#)
- [Client Types | 1085](#)
- [Call Admission Control Settings for Voice | 1085](#)
- [Retry Count | 1086](#)
- [Client Timeouts | 1086](#)
- [802.11n Settings | 1086](#)
- [Maximum Bandwidth Used by a WLAN Service Profile's SSID | 1087](#)

- [Maximum Transmission Unit Parameter | 1087](#)
- [Client Probing of Idle Clients | 1087](#)
- [Enable Pre-Shared Key \(PSK\) for WPA or WPA2 | 1087](#)
- [Create a Pre-Shared Key \(PSK\) Phrase for WPA or WPA2 | 1088](#)
- [Create a Pre-Shared Key \(PSK\) Raw Phrase for WPA or WPA2 | 1088](#)
- [Enforce Data Rates | 1088](#)
- [Retry Count for Sending Frames | 1088](#)
- [WPA Encryption Type Used | 1088](#)
- [Shared Key Authentication Values | 1089](#)
- [Radio Transmit Rates Used | 1089](#)
- [WMM Power Save | 1089](#)

A wireless LAN (WLAN) Service profile is a set of configurations, including a unique SSID, that provides clients part of a wireless connection to the wireless network. You must have at least one WLAN Service profile with an SSID on your wireless network for operation. Note that there are no default profiles provided by MSS—you must configure each WLAN Service profile. This topic describes the parameters configured in a WLAN Service profile.

There are many parameters, either optional or mandatory, that are associated with every SSID. The WLAN Service profile provides some of these parameters but not all of them.

This topic describes the parameters configured in a WLAN Service profile:

SSID

The SSID name is the most important configuration in a WLAN Service profile. Only the SSID name and WLAN name are actually required to create a WLAN Service profile—the other parameters have default values.

Wireless networks are identified by unique network names such as Juniper Networks_meetings or employee_patio. The unique name is known as a service set identifier, or SSID and this electronic identifier serves as a password for certain online communications. Most controllers are set to broadcast their SSID using WPA or WPA2 encryption. If an SSID is not broadcast, you must manually configure one or more SSIDs on clients so that the clients automatically find and connects to those SSIDs when it is in range of the controller. On a PC running Microsoft Windows OS, this setting is typically found in the Network Control Panel.

Beaconing the SSID Name

By default, SSIDs are beaconed, which means that an SSID advertises its name on the air. You might want to disable beaconing for security reasons, although doing so can make it more difficult for clients to access the WLAN. When you disable beaconing for an SSID, the radio still sends beacon frames, but the SSID name in the frames is blank. For a non-beaconed SSID, radios respond only to directed 802.11 probe requests that match the non-beaconed SSID string.

NOTE: Disabling beaconing is not particularly effective as a security measure because any sniffer can easily detect the name of the SSID.

Mapping WLAN Service Profiles to Additional Profiles

WLAN Service Profiles rely on other profiles, such as Authentication profiles and Authorization profiles, to provide additional parameters for the SSID. You link these other profiles as part of the WLAN configuration. (For more information, see [“Understanding Network Director SSID Configuration Using Profiles” on page 1064.](#)) The completed WLAN Service profile is then mapped to a Radio profile during configuration of the Radio profile. It is the Radio profile that is actually deployed to controllers, pulling all of the other mapped profiles along with it.

SSID Encryption

Encryption protects information within a wireless session by reading that information in a data stream and altering it to make it unreadable to users outside the network.

Encryption (ssid-type) is either on (Crypto) or off (Clear). If encryption is on, additional configuration indicates the type of encryption (WPA and/or WPA2), any added ciphers such as TKIP (tkip-mc-time) or CCMP. For more information, see [“Understanding Wireless Encryption and Ciphers” on page 899.](#)

Both WPA and WPA2 are enabled by default in Network Director. Clients that use only WPA associate by using WPA and all clients capable of using WPA2 associate by using WPA2.

NOTE: The original wireless encryption method, WEP, is not supported in Network Director.

Authentication Used for Encryption Methods

The standard for wireless LAN authentication is the IEEE 802.1X standard, which is based on the IETF's Extensible Authentication Protocol (EAP). EAP and 802.1X together provide an authentication framework.

The second way to authenticate encrypted traffic is pre-selected key (PSK) authentication. If you use PSK, you must also provide the key or password. For more information, see [“Understanding Wireless Encryption and Ciphers” on page 899](#).

Associated Authentication Profile

An Authentication profile must be associated with each WLAN Service profile—you do this by selecting an Authentication profile while configuring the WLAN. Authentication profiles are described in [“Understanding Authentication Profiles” on page 381](#).

Associated Authorization Profile

An Authorization profile is also mapped to WLAN Service Profiles—you do this by selecting an Authorization profile while configuring the WLAN. Authorization profiles are described in [“Understanding Wireless Authorization Profiles” on page 395](#).

VLAN Use

Clients usually switch VLANs when they switch controllers, usually as a result of roaming, but you can make initially assigned VLANs persist over different controllers. If an 802.1X user is not assigned to a VLAN by AAA, and subsequently roams to a controller where the VLAN he was in does not exist, a tunnel is set up so that he stays in that VLAN. This does not work for Web portal clients, however.

Bandwidth Limit for Client Sessions

You can limit bandwidth for clients to prevent one client from hogging bandwidth.

Load Balancing Between Access Points

RF load-balancing is the ability to reduce network congestion over an area by distributing client sessions across the access point radios with overlapping coverage in the area. Load balancing automatically occurs on the mobility domain to ensure maximum failover capability. For more information, see [“Understanding Load Balancing for Wireless Radios” on page 1070](#).

Using Proxy ARP

Proxy address resolution protocol (ARP) is a technique by which a device on a network answers requests for a different device. Because the proxy knows the location of the traffic’s destination, it offers its own IP address then sends the traffic on to the true destination.

Usually, wireless clients receive an IP address from a router. If you want to, you can have the controller respond to wireless clients’ search for a destination, and then forward the traffic to the router.

Restricting DHCP

The dynamic host configuration protocol (DHCP) is used to configure network devices with IP addresses from a DHCP server.

You can configure a controller to capture but not forward any wireless client traffic except DHCP traffic during authentication and authorization. This is referred to as restricting DHCP and enables a controller to authenticate and authorize new clients more quickly.

Client Types

You decide which client types to support on a WLAN. If you put all clients on one WLAN, they will be reduced to the speed of the slowest client—we do not recommend doing this.

Possible client types are:

- 802.11n clients use newer technology that can produce throughput link rates above 54 MBps, if you select only this client type and enforce the data rate (speed). Typical clients include laptops, PCs, and streaming video.

NOTE: The Wi-Fi Alliance requires that high-throughput (802.11n) transmissions use WPA2 and CCMP.

- 802.11g clients can have throughput link rates as fast as 54 MBps, with a more average rate of 19 MBps. Typical clients include older laptops and PCs.
- 802.11a clients can have throughput link rates as fast as 54 MBps, with a more average rate of 19 MBps.
- 802.11b clients can have throughput link rates as fast as 10 MBps.

Call Admission Control Settings for Voice

Call admission control (CAC) regulates the addition of new real-time media sessions on access point radios, guaranteeing a higher quality of service to a fixed number of clients by limiting either the number of concurrent sessions or the number of concurrent phone calls.

There are two CAC methods, WMM and SVP. WMM is enabled by default and is used by all newer 802.11n devices. SVP includes all the configurations of WMM, but adds Spectralink phones at the top of the priority list. SVP is a legacy technology—Spectralink's new phones use WMM.

As the name indicates, call admission control applies to real-time media traffic as opposed to data traffic. Call admission control mechanisms and quality-of-service settings work together to protect voice traffic from the negative effects of other voice traffic and keep excess voice traffic off the network. For more information, see [“Understanding Call Admission Control” on page 1068](#).

Retry Count

The retry count is the number of times a channel resends a frame without getting a response. You can configure 1 - 15 attempts, with each attempt taking more time and affecting throughput. We recommend configuring five attempts.

You can also specify either a long retry count or a short retry count. The difference between the two methods is the length of the pause between attempts. In general, a short retransmission works best in heavy traffic and is used most often. You might want to try a long transmission if the network is experiencing a lot of interference.

Client Timeouts

Timeouts are used to disconnect clients under certain circumstances. You can change the timeframes for these timeouts:

- User Idle—180 seconds
- Handshake attempt (logon)—20 milliseconds
- Web portal session—0 seconds (which means no timeout)

802.11n Settings

802.11n is the most recent wireless technology that utilizes different mechanisms than previous versions of 802.11. Therefore, there are settings that apply only to 802.11n traffic.

Guard Intervals

A guard interval is the interval is observed before the next bit of traffic is transmitted. This guard interval ensures that bit transmissions do not interfere with one another. As long as the echoes fall within this interval, they do not affect the receiver's ability to safely decode the actual data, because data is interpreted only outside the guard interval—it eliminates intersymbol interference. In normal 802.11 operation, the guard interval is 800 ns. In 802.11n operation, short guard intervals of 400 ns are supported. Shorter guard intervals between symbols increases throughput. Legacy devices might require long guard intervals. By reducing this interval (called *short guard interval*), data bits are transmitted in shorter intervals and provide for increased throughput.

Frame Aggregation

You can enable frame aggregation for certain frame types in 802.11n. Multiple packets of application data can be aggregated into a single packet called an aggregated MAC protocol data Unit (A-MPDU). This improves performance because the number of packets is reduced.

After transmission of every frame, an idle time called Interframe Spacing (IFS) is observed before transmitting the subsequent frame. When frames are aggregated, fewer IFS intervals are used, which in turn reduces the time for data transmission. In addition, when clients operating in 802.11n send acknowledgement for

block of aggregated packets instead of individual packets, overhead involved in frame acknowledgements and increasing overall throughput is reduced.

MAC Service Data Unit (MSDU) Length

With 802.11n, you can change the maximum length for a MAC service data unit (MSDU) to reduce the overhead associated with each transmission. An MSDU is the service data unit received from the logical link control (LLC) sub-layer which lies above the medium access control (MAC) sub-layer in a protocol stack. When 802.11n is an enabled client type for this WLAN, you can configure the maximum aggregated MSDU packet length. This enables joining multiple packets together into a single transmission unit, which reduces the overhead associated with each transmission. MSDU default length is 4K.

MPDU Length

MPDU can have a maximum length for frame aggregation.

Maximum Bandwidth Used by a WLAN Service Profile's SSID

The speed of a computer network is most commonly stated in bandwidth units of Megabits per second (Mbps) or Gigabits (Gbps). This standard measure of communication capacity (data rate) is advertised by all computer networking equipment. When you indicate a maximum bandwidth for a WLAN, you are indicating the highest data rate you support with a given WLAN Service profile. Devices in each category have a maximum bandwidth, so the maximum bandwidth also determines which devices are supported. For example, the 802.11g standard for wireless networking supports a maximum bandwidth of 54 Mbps, including overhead. 802.11n supports a maximum bandwidth of 600 Mbps.

Maximum Transmission Unit Parameter

The maximum transmission unit (MTU) of the communications protocol of a layer is the size (in bytes) of the largest protocol data unit that the layer can forward. MTU parameters usually appear in association with a communications interface such as a NIC card or serial port.

Client Probing of Idle Clients

Idle client probing sends periodic keepalives from a radio to non-transmitting clients. By default, a radio sends idle-client probes every 10 seconds to each client with a session on the radio to verify that the client is still active. The probes are unicast null-data frames. Normally, an active client sends an ACK in reply to an idle-client probe. If a client does not send any data or respond to idle-client probes before the user idle timeout expires, the client session is disassociated.

Enable Pre-Shared Key (PSK) for WPA or WPA2

WPA-PSK is an authentication mechanism in which users provide credentials to verify whether or not to allow them access to a network. This requires that a single password be entered into each WLAN node

(access points, wireless routers, client adapters, bridges). When the passwords match, a client is granted access to a WLAN.

Create a Pre-Shared Key (PSK) Phrase for WPA or WPA2

A pre-shared key (PSK) phrase is the hexadecimal secret phrase used for authenticating WPA or WPA2 clients. Note that either WPA or WPA2 security must be enabled for this to have any effect.

Create a Pre-Shared Key (PSK) Raw Phrase for WPA or WPA2

A pre-shared key (PSK) raw phrase is the raw hexadecimal secret phrase used for authenticating WPA or WPA2 clients. Note that either WPA or WPA2 security must be enabled for this to have any effect.

Enforce Data Rates

By default, a client can associate with and transmit data to an access point by using a slower data rate than the mandatory or standard rate, although the access point does not necessarily transmit data back to the client at the slower rate.

When you enforce data rates, a connecting client must transmit at one of the mandatory or standard rates to associate with the access point. Clients transmitting at slower rates cannot associate with the access point.

Retry Count for Sending Frames

Retry-count settings indicate how many times that the network sends either a long unicast frame or short unicast frame without receiving an acknowledgement. If you set retry count to zero, frames are sent once with no retries.

NOTE: The fragmentation threshold uses the short-retry-count for frames shorter than 2346 bytes and uses the long-retry-count for frames that are 2346 bytes or longer.

WPA Encryption Type Used

Enable and choose WPA encryption—either WPA or WPA2. Either or both TKIP and CCMP cipher algorithms for WPA and WPA2 can be added.

NOTE: The Wi-Fi Alliance requires that high-throughput (802.11n) transmissions use WPA2 and CCMP.

Shared Key Authentication Values

Shared key authentication is a process by which a client gains access to a WLAN by using an encryption key. The key, obtained in advance by the client, must match a key stored at the access point. To begin the connection process, the client sends a request for authentication to the access point. The access point responds by generating a sequence of characters called a challenge text for the computer. The computer encrypts the challenge text with the key and transmits the message back to the access point. The access point decrypts the message and compares the result with the original challenge text. If there are no discrepancies, the access point sends an authentication code to the connecting computer. Finally, the computer accepts the authentication code and becomes part of the network for the duration of the session or for as long as it remains within range of the original access point. If the decrypted message does not precisely agree with the original text, the access point does not allow the computer to become part of the network.

Radio Transmit Rates Used

Radio transmit rates supported by access point radios have defaults, but you can change the transmit rates for the radios. Each type of radio (802.11a, 802.11b, 802.11g, and 802.11n) providing service to an SSID has a set of rates the radio is enabled to use for sending beacons, multicast frames, and unicast data. The rate you set also specifies the rates clients must support to associate with a radio.

WMM Power Save

WMM Power Save is disabled by default, even though it saves client battery life, because clients that use power save must send a separate PSpoll to retrieve each unicast packet buffered by the access point radio. This increases bandwidth and affects performance. For more information, see [“Understanding WMM Power Save and WLAN Client Battery Life” on page 859](#).

RELATED DOCUMENTATION

[Creating and Managing a WLAN Service Profile | 1090](#)

[Understanding WMM Power Save and WLAN Client Battery Life | 859](#)

[Understanding Wireless Authorization Profiles | 395](#)

[Understanding Wireless Encryption and Ciphers | 899](#)

[Understanding the Network Director User Interface | 85](#)

[Understanding Network Director SSID Configuration Using Profiles | 1064](#)

[Network Director Documentation home page](#)

Creating and Managing a WLAN Service Profile

IN THIS SECTION

- [Managing WLAN Service Profiles | 1090](#)
- [Before You Create a WLAN Service Profile | 1092](#)
- [Creating a WLAN Service Profile | 1092](#)
- [Specifying WLAN Service Profile Quick Setup | 1093](#)
- [Specifying WLAN Service Profile Custom Setup Settings | 1102](#)
- [What To Do Next | 1119](#)

WLAN Service profiles create an SSID and provide many of the parameters the SSID needs to operate—you must have at least one WLAN Service profile in your wireless network. There are no default WLAN Profiles provided in Network Director. Therefore, you must create at least one WLAN Service profile. For more information about the individual WLAN Service profile parameters, see [“Understanding WLAN Service Profiles” on page 885](#).

This topic describes:

Managing WLAN Service Profiles

From the Manage WLAN Service Profiles page, you can:

- Create a new WLAN Service profile by clicking **Add**. For directions, see [“Creating a WLAN Service Profile” on page 1092](#).
- Modify an existing profile by selecting it and then clicking **Edit**.
- View information about a WLAN Service profile by either clicking the profile name or by selecting the profile and clicking **Details**.
- Delete profiles by selecting a profile and clicking **Delete**.

TIP: You cannot delete profiles that are in use—that is, assigned to objects or used by other profiles. To see the current assignments for a profile, click the hyperlink provided for the profile name to display the details.

- Clone a profile by selecting a profile and clicking **Clone**.

[Table 238](#) describes the information provided about WLAN Service profiles on the Manage WLAN Service Profiles page. This page lists all WLAN Service profiles defined for your network, regardless of your current selected scope in the network view.

Table 238: Manage WLAN Service Profile Fields

Field	Description
Profile Name	Name given to the profile when the profile was created.
Device Family	Wireless controllers (WLC)
SSID	Name broadcast by access points to radios.
SSID Type	SSID type refers to encryption. Encryption is either on (Crypto) or off (Clear).
Service Profile Type	WLAN Service profiles are tailored for different conditions, such as WEB, Open, Voice, 802.1X, or Custom.
Authorization Profile Name	Associated Authorization profile name. Authorization refers to the levels of information available to each client.
Authentication Profile Name	Associated Authentication profile name. Authentication is the process of identifying yourself to the network, for example logging on.
CoS Profile	Class of Services profile associated with this WLAN Service profile. For more information, see “Understanding Class of Service (CoS) Profiles” on page 609.
Description	Any description provided during creation of the WLAN Service profile.
Creation Time	Date and time when the profile was created.
Last Updated Time	Date and time when the profile was last modified.
User Name	The username of the person or system that created or modified the profile.

TIP: All columns might not be currently displayed. To show or hide fields in the table, click the down arrow on the field header, select **Columns**, and select or clear the check box adjacent to the field that you want to show or hide.

Before You Create a WLAN Service Profile

You need at least one Authentication Profile and one Authorization profile to create a WLAN Service profile—these profiles are mapped to the WLAN Service profile to create an SSID. You can use existing Authentication Profiles and Authorization profiles—see [“Creating and Managing Authentication Profiles” on page 383](#)) or an Authorization profile ([“Creating and Managing Wireless Authorization Profiles” on page 395](#)). You can alternately create the profiles during WLAN Service profile creation.

Creating a WLAN Service Profile


In Network Director, you configure wireless SSIDs by creating WLAN Service profiles. In a WLAN Service profile, at minimum with Quick Setup, you specify:

- A WLAN Profile name
- A Service Profile Type: 802.1X, Voice, Web Portal, Open Access, or Custom
- An SSID name
- Encryption setting (on or off). If encryption is on, you must indicate a method of encryption.
- Authentication for client connection
- Authorization for client connection

To create a WLAN Service profile:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.
3. In the Tasks pane, expand **Wireless**, expand **Profiles**, and then click **WLAN Service**.

The Manage WLAN Service Profiles page opens, displaying the list of currently configured WLAN Service profiles.

4. Click **Add**.

The Create WLAN Profile page is displayed with two tabs, **Quick Setup** and **Custom Setup**. The Quick Setup page is displayed.

5. You can complete only the required settings for a WLAN Service profile as described in both the online help and in *Specifying WLAN Service Profile Quick Setup*.

6. Optionally, click the **Custom Setup** tab and complete any or all advanced settings for a WLAN Service profile as described in both the online help and in *Specifying WLAN Service Profile Custom Setup Settings*.

7. Click **Done**.

The WLAN Service profile is added to the list of profiles that you can include in a Radio profile.

NOTE: Assigned settings from any profile, including this one, have lower priority than settings made directly to a controller or an access point. For more information, see [“Adding and Managing an Individual Access Point” on page 1156](#) and [“Configuring a Controller” on page 1037](#).

SEE ALSO

[Specifying WLAN Service Profile Quick Setup](#)

[Specifying WLAN Service Profile Custom Setup Settings](#)

[Understanding the Network Director User Interface | 85](#)

[Understanding WLAN Service Profiles | 885](#)

Specifying WLAN Service Profile Quick Setup

To configure only the required settings for the WLAN Service profile, enter the settings described in [Table 239](#). Required settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 239: WLAN Service Profile Quick Setup

Field	Action
Profile Name (all Service Profile Types)	<p>Type a unique name that identifies the profile.</p> <p>Use up to 32 characters for wireless profile names. Profile names must not contain special characters or spaces. Note that profiles automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character.</p>
Description (all Service Profile Types)	<p>Type 0 through 256 alphanumeric characters, including spaces and special characters.</p>
Service Profile Type	<p>Indicate one of these Service profiles:</p> <ul style="list-style-type: none"> ● 802.1X Service Profile ● Voice Service Profile ● Web Portal Service Profile ● Open Access Service Profile ● Custom Service Profile <p>The remaining setting options change, depending on which Service profile you selected here.</p>
SSID (all Service profile types)	<p>Type a unique name to be broadcast from access points and selected by clients. Use up to 32 characters and only the special character _. For more information about SSIDs, see “Understanding Network Director SSID Configuration Using Profiles” on page 1064 and “Understanding the Network Terms SSID, BSSID, and ESSID” on page 1061.</p>
SSID Type (Voice, Web-Portal, Open-Access, Custom Service Profile)	<p>SSID type refers to the encryption setting, which is either Encrypted or Unencrypted. For more information, see “Understanding Wireless Encryption and Ciphers” on page 899. If you selected the Service Profile Type 802.1X, the SSID Type is automatically set—you have no option here.</p>

Table 239: WLAN Service Profile Quick Setup (continued)

Field	Action
Vendor (Voice Service Profile Type)	Select one of the supported vendors for voice products: <ul style="list-style-type: none"> • SpectraLink • Vocera • Avaya • Ascom • Aastra • Other
Enable Voice Tracking (Voice Service Profile Type)	You can configure or select a CoS Profile to work with the WLAN Service profile for wireless access to Voice over IP (VoIP) devices.

Security Settings

You can have any or all of the available security types enabled: **RSN (WPA2)**, **WPA**, or **Static WEP**.

Table 239: WLAN Service Profile Quick Setup (continued)

Field	Action
RSN (WPA2) or WPA (802.1X, Voice, and Custom Service Profile Types)	<p>AES (CCMP): WPA2 with CCMP</p> <p>CCMP encryption ciphers are part of WPA2 (RSN) encryption. Counter Cipher Mode with Block Chaining Message Authentication Code Protocol (CCMP) implements the standards of the IEEE 802.11i amendment to the original IEEE 802.11 standard and is based on the Counter Mode with CBC-MAC (CCM) of the AES standard. For more information, see “Understanding Wireless Encryption and Ciphers” on page 899.</p> <p>NOTE: CCMP and TKIP ciphers can be applied simultaneously. In this case, clients will authenticate to the CCMP version if possible, otherwise they will use the TKIP version.</p> <hr/> <p>TKIP: WPA2 with TKIP</p> <p>TKIP encryption ciphers are part of WPA encryption. Temporal Key Integrity Protocol (TKIP) provides link-layer security without requiring replacement of legacy hardware the way CCMP does. For more information, see “Understanding Wireless Encryption and Ciphers” on page 899.</p> <p>NOTE: CCMP and TKIP ciphers can be applied simultaneously. In this case, clients will authenticate to the CCMP version if possible, otherwise they will use the TKIP version.</p>
Static WEP (802.1X, Voice, and Custom Service Profile Types)	<p>To enable static WEP encryption, configure the static WEP keys and assign them to unicast and multicast traffic. Make sure you configure the same static keys on the clients.</p>

Authentication Settings

Authentication is the process of identifying yourself to the network, for example logging on. Here you have the option to use an existing Authentication Profile (**Select Existing Authentication**) or to set authentication settings (**Configure Authentication Settings**). For more information about Authentication Profiles, see [“Understanding Authentication Profiles” on page 381](#).

Table 239: WLAN Service Profile Quick Setup (continued)

Field	Action
<p>Configure Authentication Settings for an 802.1X Service Profile, Voice Service Profile, or Web Portal</p>	<p>Authentication for 802.1X, voice, or a Web portal is done with a RADIUS server. You can either create a RADIUS server configuration here or you can select an existing RADIUS server.</p> <p>The default is to create a RADIUS server. Provide a RADIUS Server Address and Secret.</p> <p>To use an existing RADIUS server, enable Select RADIUS Server, click Select, and then select one of the RADIUS servers listed in the Choose RADIUS Profile window. Click OK.</p>

Table 239: WLAN Service Profile Quick Setup (continued)

Field	Action
Configure Authentication Settings for a Custom Service Profile	<hr/>

Table 239: WLAN Service Profile Quick Setup (*continued*)

Field	Action
	<p>When you enable Configure Existing Authentication for a Custom Service Profile, add one or more Access Rules to the WLAN Profile:</p> <ol style="list-style-type: none"> 1. Click Add. <ul style="list-style-type: none"> The Add access Rule window opens. 2. Select an access type: <ul style="list-style-type: none"> • 802.1X Access—Authenticate the client by using 802.1X authentication. For more information, see “Understanding the IEEE 802.11 Standard for Wireless Networks” on page 1076. • MAC Access—Authenticate the client by using MAC RADIUS authentication. • Web Access—Have clients log into a Web page before granting access to the SSID. • Open Access—Automatically authenticate the client and enable access to the requested SSID without requiring a username and password. <p>TIP: Open Access has no additional authentication settings. You can only indicate that you want to Enable Accounting. You must either enable local accounting or specify an Access Profile.</p> 3. Provide a matching glob, a shorthand method for matching an authentication, authorization, and accounting (AAA) to either a single user or a set of users. A user glob can contain up to 80 characters and cannot include spaces or tabs. The double-asterisk (**) wildcard characters with no delimiter characters match all user names. The single-asterisk (*) wildcard character matches any number of characters up to, but not including, a delimiter character in the glob. Valid user glob delimiter characters are the at (@) sign and the period (.). <p>NOTE: The matching glob value must be unique and cannot be used for any other access rules.</p> 4. Select an EAP type:

Table 239: WLAN Service Profile Quick Setup (continued)

Field	Action
	<ul style="list-style-type: none"> ● External Authentication Server (default)—Use an external server for authentication. Do not Enable Local Authentication if you select this. ● PEAP Offload—Offload all EAP processing from server groups. In this case, the RADIUS server is not required to communicate by using the EAP protocols. ● Local EAP—Use a local database to authenticate clients. Encryption and data integrity checking are provided for the connection. Use only with Local Authentication. <p>5. Either Enable Authentication or Enable Local Authentication, depending on which option you chose in the previous step.</p> <p>6. Optionally, Enable Accounting and/or Enable Local Accounting. Select a Record Type, either Start-Stop or Stop-Only.</p> <p>7. You need to reference an Access Profile.</p> <p>Create an Access Profile by providing a RADIUS Server Address and RADIUS Secret.</p> <p>Select an Access Profile by enabling Select Access Profile and selecting an Access Profile from the list.</p> <p>8. Click OK.</p> <p>The Add access Rule window closes and the access rule is added to the list of Access Rules under Authentication Settings.</p>
<p>Authorization Settings</p> <p>Authorization refers to the levels of information available to each client. For more information, see “Understanding Wireless Authorization Profiles” on page 395.</p>	

Table 239: WLAN Service Profile Quick Setup (continued)

Field	Action
Configure Authorization Settings (all service types)	When you enable Configure Authorization Settings , select either a VLAN Name or a VLAN Pool .
	When you select VLAN Name do the following: <ol style="list-style-type: none"> 1. Click Select. The Choose VLAN Profile window opens. 2. Select one of the existing VLAN Profiles from the list. 3. Click OK. The Choose VLAN Profile window closes and the name of the VLAN Profile you selected appears in the VLAN Name field under Authorization Settings.
	When you select VLAN Pool do the following: <ol style="list-style-type: none"> 1. Click Select. The Choose VLAN Pool window opens. 2. Select one of the existing VLAN pools from the list. 3. Click OK. The Choose VLAN Pool window closes and the name of the VLAN pool you selected appears in the VLAN Pool field under Authorization Settings.

Table 239: WLAN Service Profile Quick Setup (continued)

Field	Action
Select Existing Authorization (all service types)	<p>When you enable Select Existing Authorization, do the following:</p> <ol style="list-style-type: none"> 1. Click Select. The Choose Authorization Profile window opens. 2. Select one of the existing Authorization Profiles from the list. 3. Click OK. The name of the Authorization Profile you selected appears in the Authorization Profile field under Authorization Settings.

SEE ALSO

[Creating and Managing a WLAN Service Profile | 1090](#)

[Understanding the Network Director User Interface | 85](#)

[Understanding WLAN Service Profiles | 885](#)

Specifying WLAN Service Profile Custom Setup Settings**IN THIS SECTION**

- [Specifying Basic Settings for Custom WLAN Profile Setup | 1103](#)
- [Specifying WLAN Settings for Custom WLAN Profile Setup | 1111](#)
- [Specifying Web Portal Settings Under Advanced WLAN Profile Setup | 1112](#)
- [Specifying 802.11n and Client Type Settings Under Advanced WLAN Profile Setup | 1113](#)
- [Specifying Voice Configuration Settings Under Advanced WLAN Profile Setup | 1115](#)
- [Specifying Broadcast Settings Under Advanced WLAN Profile Setup | 1116](#)
- [Specifying Client Timeouts Under Advanced WLAN Profile Setup | 1116](#)
- [Specifying Rate Configuration Under Advanced WLAN Profile Setup | 1117](#)
- [Specifying Device Detection Settings Under Advanced WLAN Profile Setup | 1119](#)

When you select the **Custom Setup** tab, nine different groups of settings are available in a list on the left side of the window: **Basic Settings**, **Web Portal Settings**, **802.11n and Client Type Settings**, **Voice Configuration**, **Broadcast Settings**, **Client Timeouts**, **Rate Configuration**, and **Device Detection**.

Follow these directions to reconfigure any of the nine WLAN Service profile options under the **Custom Setup** tab:

Specifying Basic Settings for Custom WLAN Profile Setup

[Table 240](#) describes the required basic settings for a WLAN Service profile. These are the same settings found under the **Quick Setup** tab.

Table 240: WLAN Service Profile Basic Settings

Field	Action
Profile Name (all Service Profile Types)	Type a unique name that identifies the profile. Use up to 32 characters for wireless profile names. Profile names must not contain special characters or spaces. Note that profiles automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character.
Description (all Service Profile Types)	Type 0 through 256 alphanumeric characters, including spaces and special characters.
Service Profile Type	Indicate one of these Service profiles: <ul style="list-style-type: none"> ● 802.1X Service Profile ● Voice Service Profile ● Web Portal Service Profile ● Open Access Service Profile ● Custom Service Profile <p>The remaining setting options change, depending on which Service profile you selected here.</p>
SSID (all Service profile types)	Type a unique name to be broadcast from access points and selected by clients. Use up to 32 characters and only the special character _. For more information about SSIDs, see “Understanding Network Director SSID Configuration Using Profiles” on page 1064 and “Understanding the Network Terms SSID, BSSID, and ESSID” on page 1061 .

Table 240: WLAN Service Profile Basic Settings (continued)

Field	Action
SSID Type (Voice, Web-Portal, Open-Access, Custom Service Profile)	SSID type refers to the encryption setting, which is either Encrypted or Unencrypted . For more information, see "Understanding Wireless Encryption and Ciphers" on page 899 . If you selected the Service Profile Type 802.1X, the SSID Type is automatically set—you have no option here.
Vendor (Voice Service Profile Type)	Select one of the supported vendors for voice products: <ul style="list-style-type: none"> ● SpectraLink ● Vocera ● Avaya ● Ascom ● Aastra ● Other
Enable Voice Tracking (Voice Service Profile Type)	You can configure or select a CoS Profile to work with the WLAN Service profile for wireless access to Voice over IP (VoIP) devices.

Security Settings

You can have any or all of the available security types enabled: **RSN (WPA2)**, **WPA**, or **Static WEP**.

Table 240: WLAN Service Profile Basic Settings (continued)

Field	Action
RSN (WPA2) or WPA (802.1X, Voice, and Custom Service Profile Types)	<p>AES (CCMP): WPA2 with CCMP</p> <p>CCMP encryption ciphers are part of WPA2 (RSN) encryption. Counter Cipher Mode with Block Chaining Message Authentication Code Protocol (CCMP) implements the standards of the IEEE 802.11i amendment to the original IEEE 802.11 standard and is based on the Counter Mode with CBC-MAC (CCM) of the AES standard. For more information, see “Understanding Wireless Encryption and Ciphers” on page 899.</p> <p>NOTE: CCMP and TKIP ciphers can be applied simultaneously. In this case, clients will authenticate to the CCMP version if possible, otherwise they will use the TKIP version.</p> <hr/> <p>TKIP: WPA2 with TKIP</p> <p>TKIP encryption ciphers are part of WPA encryption. Temporal Key Integrity Protocol (TKIP) provides link-layer security without requiring replacement of legacy hardware the way CCMP does. For more information, see “Understanding Wireless Encryption and Ciphers” on page 899.</p> <p>NOTE: CCMP and TKIP ciphers can be applied simultaneously. In this case, clients will authenticate to the CCMP version if possible, otherwise they will use the TKIP version.</p>
Static WEP (802.1X, Voice, and Custom Service Profile Types)	<p>To enable static WEP encryption, configure the static WEP keys and assign them to unicast and multicast traffic. Make sure you configure the same static keys on the clients.</p>

Authentication Settings

Authentication is the process of identifying yourself to the network, for example logging on. Here you have the option to use an existing Authentication Profile (**Select Existing Authentication**) or to set authentication settings (**Configure Authentication Settings**). For more information about Authentication Profiles, see [“Understanding Authentication Profiles” on page 381](#).

Table 240: WLAN Service Profile Basic Settings (continued)

Field	Action
<p>Configure Authentication Settings for an 802.1X Service Profile, Voice Service Profile, or Web Portal</p>	<p>Authentication for 802.1X, voice, or a Web portal is done with a RADIUS server. You can either create a RADIUS server configuration here or you can select an existing RADIUS server.</p> <p>The default is to create a RADIUS server. Provide a RADIUS Server Address and Secret.</p> <p>To use an existing RADIUS server, enable Select RADIUS Server, click Select, and then select one of the RADIUS servers listed in the Choose RADIUS Profile window. Click OK.</p>

Table 240: WLAN Service Profile Basic Settings (continued)

Field	Action
Configure Authentication Settings for a Custom Service Profile	<hr/>

Table 240: WLAN Service Profile Basic Settings (continued)

Field	Action
	<p>When you enable Configure Existing Authentication for a Custom Service Profile, add one or more Access Rules to the WLAN Profile:</p> <ol style="list-style-type: none"> 1. Click Add. The Add access Rule window opens. 2. Select an access type: <ul style="list-style-type: none"> • 802.1X Access—Authenticate the client by using 802.1X authentication. For more information, see “Understanding the IEEE 802.11 Standard for Wireless Networks” on page 1076. • MAC Access—Authenticate the client by using MAC RADIUS authentication. • Web Access—Have clients log into a Web page before granting access to the SSID. • Open Access—Automatically authenticate the client and enable access to the requested SSID without requiring a username and password. <p>TIP: Open Access has no additional authentication settings. You can only indicate that you want to Enable Accounting. You must either enable local accounting or specify an Access Profile.</p> 3. Provide a matching glob, a shorthand method for matching an authentication, authorization, and accounting (AAA) to either a single user or a set of users. A user glob can contain up to 80 characters and cannot include spaces or tabs. The double-asterisk (**) wildcard characters with no delimiter characters match all user names. The single-asterisk (*) wildcard character matches any number of characters up to, but not including, a delimiter character in the glob. Valid user glob delimiter characters are the at (@) sign and the period (.). <p>NOTE: The matching glob value must be unique and cannot be used for any other access rules.</p> 4. Select an EAP type:

Table 240: WLAN Service Profile Basic Settings (continued)

Field	Action
	<ul style="list-style-type: none"> ● External Authentication Server (default)—Use an external server for authentication. Do not Enable Local Authentication if you select this. ● PEAP Offload—Offload all EAP processing from server groups. In this case, the RADIUS server is not required to communicate by using the EAP protocols. ● Local EAP—Use a local database to authenticate clients. Encryption and data integrity checking are provided for the connection. Use only with Local Authentication. <ol style="list-style-type: none"> 5. Either Enable Authentication or Enable Local Authentication, depending on which option you chose in the previous step. 6. Optionally, Enable Accounting and/or Enable Local Accounting. Select a Record Type, either Start-Stop or Stop-Only. 7. You need to reference an Access Profile. Create an Access Profile by providing a RADIUS Server Address and RADIUS Secret. Select an Access Profile by enabling Select Access Profile and selecting an Access Profile from the list. 8. Click OK. The Add access Rule window closes and the access rule is added to the list of Access Rules under Authentication Settings.
<p>Authorization Settings</p> <p>Authorization refers to the levels of information available to each client. For more information, see “Understanding Wireless Authorization Profiles” on page 395.</p>	

Table 240: WLAN Service Profile Basic Settings (continued)

Field	Action
Configure Authorization Settings (all service types)	When you enable Configure Authorization Settings , select either a VLAN Name or a VLAN Pool .
	When you select VLAN Name do the following: <ol style="list-style-type: none"> 1. Click Select. The Choose VLAN Profile window opens. 2. Select one of the existing VLAN Profiles from the list. 3. Click OK. The Choose VLAN Profile window closes and the name of the VLAN Profile you selected appears in the VLAN Name field under Authorization Settings.
	When you select VLAN Pool do the following: <ol style="list-style-type: none"> 1. Click Select. The Choose VLAN Pool window opens. 2. Select one of the existing VLAN pools from the list. 3. Click OK. The Choose VLAN Pool window closes and the name of the VLAN pool you selected appears in the VLAN Pool field under Authorization Settings.

Table 240: WLAN Service Profile Basic Settings (continued)

Field	Action
Select Existing Authorization (all service types)	<p>When you enable Select Existing Authorization, do the following:</p> <ol style="list-style-type: none"> 1. Click Select. The Choose Authorization Profile window opens. 2. Select one of the existing Authorization Profiles from the list. 3. Click OK. The name of the Authorization Profile you selected appears in the Authorization Profile field under Authorization Settings.

Specifying WLAN Settings for Custom WLAN Profile Setup

Reconfigure any or all of the available advanced WLAN settings listed in [Table 241](#) for any WLAN Service profile.

Table 241: Custom WLAN Settings for WLAN Profiles

Field	Description
Beacon (default is enabled)	Select this check box to indicate that the SSID name of this WLAN will be broadcast. Clear this check box to hide the name of the SSID. See “Understanding Network Director SSID Configuration Using Profiles” on page 1064.
Keep Initial VLAN (default is disabled)	Select this check box to specify that VLANs persist over different controllers. If an 802.1X user is not assigned to a VLAN by AAA, and subsequently roams to a controller where the VLAN he was in does not exist, a tunnel is set up so that the user stays in that VLAN. This, however, does not work for Web portal clients. For more information about configuring VLANs in Network Director, see “Understanding VLAN Profiles” on page 499, and “Creating and Managing VLAN Profiles” on page 502.
Load Balance Exempt (default is disabled)	Select this check box to prevent access points from sharing the data traffic load for this SSID. This only has an effect if the associated Access profile has load-balancing enabled. For more information about load-balancing, see “Understanding Load Balancing for Wireless Radios” on page 1070.

Table 241: Custom WLAN Settings for WLAN Profiles (continued)

Field	Description
Fall Through Access (default is None)	Select the action the system will take when authentication fails. You can indicate that a Web Portal be used login or you can just enable login with open access (Last Resort). The default is None . If you select Web Portal , you must also complete the additional Web Portal Settings under the Authentication Profile tab—see either “Creating and Managing Authentication Profiles” on page 383 or <i>Specifying WLAN Service Profile Quick Setup</i> .
Bandwidth Limit (default is disabled)	Select this option to limit the bandwidth of any client session connected to an access point with this given WLAN Service profile (SSID). You must also indicate a Max Bandwidth—default is 1Kbps. Max Bandwidth: If Bandwidth Limit is selected, select a maximum bandwidth in kilobytes per second to limit any client session connected to an access point with this given WLAN Service profile (SSID). This bandwidth setting in a WLAN Service profile overrides any bandwidth setting configured in a CoS profile as part of an Authorization profile. Default is 1Kbps.
Backup SSID Mode (default is disabled)	You can enable this option and then provide a backup SSID for remote access points on the network. Configure the SSID in “Creating and Managing Remote Site Profiles” on page 1014 . Also see “Assigning Remote Site Profiles to Access Points” on page 1024 and <i>Understanding Remote Access Points</i> Backup SSID Timeout —Check to measure the length of time before the backup SSID starts to broadcast. Backup SSID Timeout —Length of time a remote access point is non-functional before the backup SSID starts to broadcast.
Keep Clients (default is enabled)	Specifies whether clients (sessions) are dropped or not during an outage period. The default is to keep the sessions.
Multicast Conversion (default is disabled)	When checked, this feature enables multicast to unicast conversion on packets.

Specifying Web Portal Settings Under Advanced WLAN Profile Setup

If **Service Profile Type** (under the Basic Settings tab) is set to **Web Portal Service Profile**, complete the Web Portal Settings listed in [Table 242](#).

Table 242: Web Portal Settings for WLAN Profiles

Field	Description
Web Portal ACL (Web Portal Service Profile only)	ACL stands for access control list. To restrict Layer 3 traffic among clients in the same VLAN, use an ACL. You can configure the ACL yourself or use the Restrict L3 Traffic option. The default is portalacl .
Web Portal Login Page (Web Portal Service Profile only)	To add a Web Portal Login page, select Web Portal Logout and indicate the Web Portal Login Page name.
Web Portal Logout (Web Portal Service Profile only)	By default, Web Portal Logout is disabled—you do not need to provide a logout page. If you want to provide a logout page, enable this option and then provide the name of the page under Web Portal Logout .

Specifying 802.11n and Client Type Settings Under Advanced WLAN Profile Setup

If your access points and/or clients are capable of 802.11n transmission, you can change the settings listed in [Table 243](#) for those devices.

Table 243: 802.11n and Client Settings for WLAN Profiles

Field	Description
Client Types (all profile types)	All of the listed client types are enabled by default. Disable any type by removing the check mark.
802.11n	
802.11ng Mode (all profile types) (default is Enabled)	When 802.11n is an enabled client type for this WLAN, enable 802.11ng Mode to accept additional connections from 802.11g clients.
802.11na Mode (all profile types) (default is Enabled)	When 802.11n is an enabled client type for this WLAN, enable 802.11na Mode to accept additional connections from 802.11a clients.
Guard Interval (all profile types) (all profile types) (default is Short)	Select a guard interval value (Long or Short). The guard interval is the space between symbols (characters) being transmitted—it eliminates inter-symbol interference. In normal 802.11 operation, the guard interval is 800 ns (Long). In 802.11n operation, short guard intervals of 400 ns are supported. By reducing this interval (by selecting Short), data bits are transmitted in shorter intervals and provide increased throughput. TIP: Legacy devices might require long guard intervals.

Table 243: 802.11n and Client Settings for WLAN Profiles (*continued*)

Field	Description
Frame Aggregation for 802.11n (all profile types)	<p>When 802.11n is an enabled client type, you can enable frame aggregation for the listed frame types:</p> <ul style="list-style-type: none"> ● All—After transmission of every 802.11n frame, an idle time called Interframe Spacing (IFS) is observed before transmitting the subsequent frame. ● MSDU—Aggregate MAC Service Data Unit aggregation collects Ethernet frames to be transmitted to a single destination and wraps them in a single 802.11n MAC header frame. This is efficient because Ethernet headers are much shorter than 802.11 headers. Only MSDUs with whose destination address and source address map to the same receiver address and transmitter address are aggregated. ● MPDU: Aggregated - MAC Protocol Data Unit aggregation collects multiple 802.11n packets of application data into a single packet called an A-MPDU. This reduces the IFS number, which in turn provides more time for data transmission. In addition, clients operating in 802.11n send acknowledgement for block of packets instead of individual packet acknowledgement, reducing overhead involved in frame acknowledgements and increasing overall throughput. ● Disabled

Table 243: 802.11n and Client Settings for WLAN Profiles (continued)

Field	Description
(default is All)	
A-MSDU Max Length (default is 4K)	Select a maximum length for a MAC service data unit (MSDU)—4K or 8K. An MSDU is the service data unit received from the logical link control (LLC) sub-layer which lies above the medium access control (MAC) sub-layer in a protocol stack. When 802.11n is an enabled client type for this WLAN, you can configure the maximum aggregated MSDU packet length. This enables joining multiple packets together into a single transmission unit, in order to reduce the overhead associated with each transmission. Default is 4K.
A-MPDU Max Length (default is 64K)	Select the MPDU maximum length for 802.11n frame aggregation—8K, 16K, 32K, or 64K. An idle time called Interframe Spacing (IFS) is observed before transmitting a data frame. When 802.11n is an enabled client type for this WLAN, multiple packets of application data are aggregated into a single packet. This is called A-MPDU (Aggregated - MAC Protocol Data Unit). This reduces the number of IFS, which in turn provides more time for data transmission. In addition, clients operating in 802.11n send acknowledgement for block of packets instead of individual packet acknowledgement. This reduces the overhead involved in frame acknowledgements and increases the overall throughput.

Specifying Voice Configuration Settings Under Advanced WLAN Profile Setup

Voice support can be part of any WLAN Service profile and you can reconfigure any of the voice settings listed in [Table 244](#).

Table 244: Voice Settings for WLAN Profiles

Field	Description
CAC Mode (default is None) (all profile types)	Select either None (no call admission control constraint), Session (call admission constrained by total number of sessions on this WLAN), or VoIP Session (call admission constrained by total number of voice over IP sessions on this WLAN). When enabled, CAC limits the number of active sessions to 14 on a radio by default. You can change the maximum number of sessions to a value from 0 to 100. For more information about call admission control, see “Understanding Call Admission Control” on page 1068 .
	Max Associated Sessions —If call admission control (CAC) is set to Session , select a number from 0 through 500 to limit the number of sessions by using this WLAN Service profile. Default is 14.
	Max VoIP Calls —If call admission control (CAC) is set to VoIP Session , select a number from 0 through 100 to limit the voice over IP calls by using this WLAN Service profile. Default is 12.

Table 244: Voice Settings for WLAN Profiles (*continued*)

Field	Description
Short Retry Count (default is 5) (all profile types)	Select the number of times (1 through 15) a channel tries to send a frame without getting a response—the default is 5. By default, the frag-threshold setting uses the short-retry-count for frames less than 2346 bytes.
Long Retry Count (default is 5) (all profile types)	Select number of times (1 through 15) a channel tries to send a frame without getting a response—the default is 5. By default, the frag-threshold setting uses the long-retry-count for frames 2346 bytes or longer.

Specifying Broadcast Settings Under Advanced WLAN Profile Setup

Broadcast Settings are mechanisms to reduce overhead caused by wireless broadcast traffic or traffic from unauthenticated clients. Any of the settings listed in [Table 245](#) can be enabled for any of the profile types:

Table 245: Broadcast Settings for WLAN Profiles

Field	Description
Proxy ARP (disabled by default) (all profile types)	Select to have the controller respond on behalf of wireless clients to ARP requests for IP addresses.
No Broadcast (disabled by default) (all profile types)	Select to send unicasts to clients for ARP requests and DHCP offers. Send ACKs instead of forwarding them as multicasts.
DHCP Restrict (disabled by default) (all profile types)	Select to have controller capture but not forward any traffic except DHCP traffic for a wireless client during authentication and authorization.

Specifying Client Timeouts Under Advanced WLAN Profile Setup

Client Timeout settings determine when clients are dropped by the network. Any of the settings listed in [Table 246](#) can be reconfigured:

Table 246: Client Timeout Settings for WLAN Profiles

Field	Description
User Idle Timeout (default is 180 seconds) (all profile types)	Select the number of seconds (20 through 86400) that a voice call can be idle before it is dropped. Default is 180.

Table 246: Client Timeout Settings for WLAN Profiles (*continued*)

Field	Description
Idle Client Probing (default is enabled) (all profile types)	Select to send keepalives from radios to idle clients on the SSID to check for rogue devices.
Web Portal Session Timeout (default is 5 seconds) (all profile types)	If a Web portal is configured, select the maximum number of seconds (5 through 28800) a user session on a Web portal can last before it is dropped. Default is 5 seconds. NOTE: Web portals are configured in Authentication profiles.
Handshake Timeout (default is 20 milliseconds) (all profile types)	Select the maximum number of milliseconds (20 through 5000) an authentication handshake can last before it is dropped. Default is 20 milliseconds and zero indicates no limit.

Specifying Rate Configuration Under Advanced WLAN Profile Setup

The following rates can be reconfigured for 802.11a, 802.11b, 802.11na, and 802.11ng:

- **Beacon Rate:** Data rate of beacon frames sent by radios. This rate is also used for probe-response frames. The valid rates depend on the radio type.
- **Multicast Rate:** Data rate of multicast frames sent by radios.
- **Transmission Rates:** Data transmission rates supported by each radio type. Select **Mandatory** to indicate that a client must support at least one of these rates to associate. Select **Standard** to indicate that valid rates are neither disabled nor mandatory.

Table 247 lists the radio default settings.

Table 247: Default Rate Settings for Radios

Field	Default
802.11a	Beacon Rate: 6.0
	Multicast Rate: Automatic
	Transmission Rates: <ul style="list-style-type: none"> • 6.0: Mandatory • 9.0: Supported • 12.0: Mandatory • 18.0: Supported

Table 247: Default Rate Settings for Radios (*continued*)

Field	Default
802.11b	Beacon Rate: 2
	Multicast Rate: Automatic
	Transmission Rates: <ul style="list-style-type: none"> • 1.0: Mandatory • 2.0: Mandatory • 5.5: Supported • 11.0: Supported
802.11g	Beacon Rate: 2
	Multicast Rate: Automatic
	Transmission Rates: <ul style="list-style-type: none"> • 1.0: Mandatory • 2.0: Mandatory • 5.5: Mandatory • 6.0: Supported
802.11na	Beacon Rate: 6
	Multicast Rate: Automatic
	Transmission Rates: <ul style="list-style-type: none"> • 6.0: Mandatory • 9.0: Supported • 12.0: Mandatory • 18.0: Supported

Table 247: Default Rate Settings for Radios (continued)

Field	Default
802.11ng	Beacon Rate: 2
	Multicast Rate: Automatic
	Transmission Rates: <ul style="list-style-type: none"> • 1.0: Mandatory • 2.0: Mandatory • 5.5: Mandatory • 6.0: Supported

Specifying Device Detection Settings Under Advanced WLAN Profile Setup

Field	Description
Detection Mode	<p>You can select from the following detection modes:</p> <ul style="list-style-type: none"> • Just Detect enables device detection but does not enforce any rules. • Enforce lets you configure the device detection timeout with a range of 1 to 60 seconds with a default value of 5 seconds. When you select Enforce, the default ACL device ACL is enabled. This ACL prevents access to the network until the device is recognized. • Disable disables the feature which is enabled by default.
Detection Timeout (default is 5 seconds)	When Enforce is selected, indicate the length of time in seconds allowed for device detection on the network.
Pre-detection ACL	When Enforce is selected, configures an ACL for device fingerprinting authorization. The Device Detect ACL is configured automatically when you enable device policy enforcement on a Service profile. This is similar to the way that portalacl is configured when the parameter auth-fallthru is set to web-portal.

What To Do Next

Next, you can either create a new Radio profile and select the WLAN Service profile during creation, or you can edit an existing Radio profile and add the WLAN Service profile to the existing Radio profile. See [“Creating and Managing a Radio Profile” on page 932](#) for directions.

NOTE: Assigned settings from any profile, including this one, have lower priority than settings made directly to a controller or an access point. For more information, see [“Adding and Managing an Individual Access Point”](#) on page 1156 and [“Configuring a Controller”](#) on page 1037.

RELATED DOCUMENTATION

[Understanding the Network Director User Interface | 85](#)

[Understanding WLAN Service Profiles | 885](#)

[Network Director Documentation home page](#)

Understanding Voice Clients and Voice Traffic

IN THIS SECTION

- [What Is Voice Over IP? | 1121](#)
- [How Is Voice Traffic Different From Data Traffic? | 1121](#)
- [What Protocols Are Used for Voice? | 1122](#)
- [What Is Different About Configuring for Voice Traffic? | 1122](#)

Voice, like all other information, travels in packets over IP networks with fixed maximum capacity. The major difference between voice and data traffic is the fact that data packets can be re-sent if they are dropped, and then applied to the empty spots in data, thereby producing complete information. With voice, there is no point in resending packets because voice only makes sense in a stream of contiguous packets. Because resending voice packets is not an option with voice over IP (VoIP), voice traffic must be more carefully configured.

This topic applies to the transmission of voice, multimedia sessions, and video, which are all part of VoIP.

This topic describes:

What Is Voice Over IP?

VoIP refers to the communication protocols, technologies, methodologies, and transmission techniques involved in the delivery of voice communications and multimedia sessions over Internet Protocol (IP) networks. VoIP systems use session control protocols to control the setup and tear-down of calls as well as audio codecs that encode speech, allowing transmission over an IP network as digital audio using an audio stream.

Voice over Wireless Phones

A VoWiFi phone is similar to a cell phone except that the radio is an 802.11 radio instead of a cellular radio. VoWiFi phones are 802.11 client stations that communicate through an access point.

VoIP is also available on smart phones and Internet devices, making it possible for users of portable devices that are not phones to place calls or send text messages over 3G, 4G, or Wi-Fi.

Private Branch Exchange (PBX)

Private Branch Exchanges (PBXs) are telephone exchanges that serve a particular business or office. The PBX makes the connections for internal phone calls and provides a way to dial out to the public network.

How Is Voice Traffic Different From Data Traffic?

Voice transmission has different issues than data transmission because voice traffic cannot be re-sent if it is lost. Three unique problems that affect call quality are latency, jitter, and packet loss.

In addition, voice traffic can have roaming issues because people are more apt to walk around with phones.

Latency, Jitter, and Packet Loss Affect Voice Traffic

Latency, or delay, is the amount of time it takes the sound of your voice to reach the ear of the other person. Maximum acceptable latency limits for VoIP are 150-200 ms, depending on your call quality requirements. Delay levels that exceed 80 ms are indication of delay issues.

Jitter is the variation in delay between packets. Because packet jitter always varies, VoIP phones use jitter buffers to smooth out the variations. A jitter buffer is simply a First-In, First Out (FIFO) memory cache that collects the packets as they arrive, forwarding them evenly spaced and in proper sequence for smooth playback. Increasing jitter buffer size can help with jitter, but only to a point. Significant jitter might cause the jitter buffer to increase to the point that delay reaches unacceptable levels. Jitter levels that exceed 20 ms are indication of jitter issues

Packet loss occurs when the maximum delay specified in the jitter buffer is exceeded. Networks tend to either occasionally drop single packets (called gaps in packet loss), or large numbers of contiguous packets in a burst. Packet loss above 5% is considered unacceptable when using the G.711 codec—sustained bursts of packet loss cause the most problems.

Voice Traffic Is Susceptible to Roaming Issues

A key advantage of voice over WLAN is mobility, but voice roaming is potentially incompatible with 802.11i. In accordance with 802.11i, users traversing a network negotiate new encryption keys with every access point they encounter. If this does not happen quickly enough, voice quality is affected.

The 802.11i fast roaming specification eliminates delays associated with re-authenticating roaming clients. Instead of generating a new encryption key with each access point, the VoIP client can use the same key because the pairwise master keys (PMKs) used are cached at the controller.

What Protocols Are Used for Voice?

Voice solutions require two types of protocol. Signaling protocols are used during call setup, management and teardown. These protocols generally require low bandwidth, might use a connection-oriented model (TCP), and are typically not delay sensitive. Bearer protocols actually carry the stream of voice samples. Bearer protocols are delay sensitive, are connectionless (UDP) and require special treatment to ensure prioritization over other types of traffic. A separate stream is usually required in each direction. Spectralink Radio Priority (SRP) is a legacy bearer protocol that you can configure.

What Is Different About Configuring for Voice Traffic?

IN THIS SECTION

- [Consider Using Call Admission Control \(CAC\) to Limit Clients per Access Point | 1123](#)
- [Ensure That Network Equipment Supports Seamless Roaming | 1123](#)
- [Support Quality of Service on all Hardware Used for Voice | 1123](#)
- [Use Automatic Power Save Delivery to Preserve the Battery Life of Phones | 1123](#)
- [Create a Unique WLAN Service Profile for Voice | 1123](#)
- [Create a Unique Radio Profile for Voice | 1123](#)

Voice traffic needs more bandwidth than data traffic, and that bandwidth needs to be protected so it remains constant and prevents degradation. For these reasons, voice traffic works best when the WLAN Service profile (SSID) and Radio profile used are dedicated to and designed for voice traffic. We recommend that you separate voice and data by SSID and preferably by RF band. Also, be sure to ensure adequate coverage at -60 dBm to -70 dBm level and capacity for the expected voice load.

To optimize voice, also follow these suggestions:

Consider Using Call Admission Control (CAC) to Limit Clients per Access Point

Call admission control, configured in the WLAN Service profile, limits the number of concurrent phone calls. Some VoIP devices use Wi-Fi Multimedia (WMM) to provide call admission control for voice clients and some devices use Spectralink voice protocol SVP to provide call admission control for voice clients. Both methods work—which one you use is dictated by your voice clients. The only difference between WMM and SVP is that SVP gives highest priority to Spectralink traffic and WMM does not. If your network has both WMM (most common) and Spectralink voice devices, we suggest that you provide a dedicated SSID for each type.

Ensure That Network Equipment Supports Seamless Roaming

Voice traffic must use seamless roaming so no call is interrupted. You can achieve seamless roaming with:

- At least 20% overlap between access points
- Between -60 dBm and -70 dBm signal strength wherever voice is required
- Roaming achieved within 50 ms
- Minimized number of router hops between the handsets and the PBX

Support Quality of Service on all Hardware Used for Voice

QoS should be supported end-to-end with voice traffic. Ensure that infrastructure switches and routers support and preserve data prioritization across the full path of any voice streams.

Use Automatic Power Save Delivery to Preserve the Battery Life of Phones

Handset battery life is a major problem. Unscheduled—using Automatic Power Save Delivery (U-APSD) provides dramatic improvements in battery life, for example, from 2 hours talk time to over 10 hours.

NOTE: Handsets that support WMM/802.11e are increasingly available.

Create a Unique WLAN Service Profile for Voice

Create a unique WLAN profile for voice—see [“Creating and Managing a WLAN Service Profile” on page 1090](#).

Create a Unique Radio Profile for Voice

Create a unique Radio profile for voice—see [“Creating and Managing a Radio Profile” on page 932](#).

RELATED DOCUMENTATION

[Creating and Managing a WLAN Service Profile | 1090](#)

[Creating and Managing a Radio Profile | 932](#)

[Network Director Documentation home page](#)

Configuring a Voice SSID with Network Director

IN THIS SECTION

- [Creating a CoS Profile Dedicated to Voice | 1124](#)

Voice over IP has different requirements than data traffic and therefore has unique configuration requirements. For an explanation of voice traffic, see [“Understanding Voice Clients and Voice Traffic” on page 1120](#).

If your network is supporting both wireless data clients and voice clients, we recommend that you have two WLAN Service profiles, one for data and one for voice—this topic describes creating a WLAN Service profile for voice clients.


NOTE: WLAN Service profiles define the SSID name.

This topic describes:

Creating a CoS Profile Dedicated to Voice

To be able to support latency-sensitive traffic such as voice traffic, you must create a dedicated CoS profile by entering voice-specific settings. This topic describes the typical settings that you can use to create a CoS profile dedicated to voice.

To create a CoS profile dedicated to voice:

1. Click the Build Mode icon  in the Network Director banner.
2. Click **CoS** under Profile and Configuration Management in the Tasks pane.
The Manage CoS Profiles page opens, displaying the list of current CoS profiles.
3. Click **Add** on the Manage CoS Profiles page and select **Wireless Controller (WLC)** in the Device Family Chooser page.
The Create CoS Profile for Wireless Controllers (WLC) wizard opens.
4. Configure the following settings:

- Provide a name for the CoS profile.
- In the Voice CoS section, select **VoIP-data**.
- (Optional) In the Voice CoS section, select **Enable Bandwidth Limit** and specify the maximum bandwidth that you want the system to reserve for voice traffic.
- In the Access Categories section, make sure that the action corresponding to the Access Category—Voice is **Permit**.

5. Click **Done** to save the settings.

For an explanation of all the CoS profile parameters, see [“Creating and Managing Wired CoS Profiles” on page 613](#).

After you create the CoS profile, you then associate the CoS profile with an Authorization profile during creation of the Authorization profile. Next, you associate the Authorization profile with a WLAN Service profile to apply the CoS settings to all the users who connect to that SSID. For information on creating a WLAN Service profile (including mapping the Authorization profile) dedicated to voice, see [“Creating and Managing a WLAN Service Profile” on page 1090](#).

RELATED DOCUMENTATION

[Understanding Voice Clients and Voice Traffic | 1120](#)

[Understanding Radio Profiles | 879](#)

[Creating and Managing Wired CoS Profiles | 613](#)

[Creating and Managing a WLAN Service Profile | 1090](#)

[Creating and Managing Wireless Authorization Profiles | 395](#)

[Creating and Managing Authentication Profiles | 383](#)

[Understanding WLAN Service Profiles | 885](#)

[Network Director Documentation home page](#)

Creating and Managing RF Snooping Filter Profiles

IN THIS SECTION

- [Managing Snooping Filter Profiles | 1126](#)
- [Creating an RF Snooping Filter Profile | 1127](#)

- Specifying RF Snooping Settings | 1128
- What To Do Next | 1132

When active scan is enabled in a Radio profile, the radios with the profile actively scan other channels in addition to the data channel that is currently in use. Active scan operates on enabled radios and disabled radios. In fact, using a radio in sentry mode as a dedicated scanner provides better rogue detection because the radio can spend more time scanning on each channel.

When a radio is scanning other channels, active snoop filters on the radio also snoop traffic on the other channels. To prevent monitoring of data from other channels, use the channel option when you configure the filter, to specify the channel on which you want to snoop.

Managing Snooping Filter Profiles

From the Manage RF Snooping page, you can:

- Create a new RF Snooping profile by clicking **Add**. For directions, see [“Creating an RF Snooping Filter Profile” on page 1127](#).
- Modify an existing RF Snooping profile by selecting it and clicking **Edit**.
- Assign a RF Snooping profile to access points by selecting the profile and clicking **Assign**. For directions, see [“Assigning RF Snooping Filter Profiles to Access Points” on page 1132](#).
- Edit an existing RF Snooping profile by selecting it and clicking **Edit Assign**.
- Delete a RF Snooping profile by selecting site name and clicking **Delete**.

TIP: You cannot delete a profile that is in use. To see the current state of a profile, select the site name and click **Details**.

- Clone a RF Snooping profile by selecting a profile and clicking **Clone**.

[Table 248](#) describes the information provided about RF Snooping profiles on the Manage Switching Profiles page. This page lists all RF Snooping profile defined for your network, regardless of the scope you selected in the network view.

Table 248: RF Snooping Profile Information

Field	Description
Snoop Filter Name	Profile name up to 15 alphanumeric characters.

Table 248: RF Snooping Profile Information (continued)


Field	Description
Enabled	A snooping filter can be disabled or enabled.
Owner	Login of user who created the Snoop Filter Profile.
Description	Description for the snoop filter.
Assignment State	Displays the assignment state of the profile. A profile can be: <ul style="list-style-type: none"> • Unassigned—When the profile is not assigned to any object. • Deployed—When the profile is assigned and is deployed from Deploy mode. • Pending Deployment—When the profile is assigned, but not yet deployed in the network.
Creation Time	Date and time when the profile was created.
Last Updated Time	Date and time when the profile was last modified.
User Name	The username of the user who created or modified the profile.

Creating an RF Snooping Filter Profile

To add an RF Snooping Filter Profile, follow these steps:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.
3. In the Tasks pane, expand **Wireless**, expand **Profiles**, and then click **RF Snooping**.
The Manage RF Snooping Profile page appears, displaying the list of currently configured RF Snooping profiles.
4. Click **Add**.

The Create RF Snooping Profile page opens.

5. Provide the RF Snooping settings listed in [“Specifying RF Snooping Settings”](#) on page 1128.
6. Click **Done**.

The new RF Snooping profile is added to the list on the Manage RF Snooping Profile page.

Specifying RF Snooping Settings

Specify the RF Snooping settings described in [Table 249](#).

Table 249: RF Snooping Settings

Field	Description
Snoop Filter Name	Type a snooping filter name up to 15 characters long.
Description	Provide a description of the snooping filter profile.
Enable	Turn the Snooping profile on and off by adding a check mark to enable it or removing the check mark to disable it.

Snoop Observer

You can either select an existing snooping observer or you can create a new snooping observer.

If the snooping filter meets these conditions, the observer must also be in the same subnet.

- Filter is running on a distributed access point.
- Access point used a DHCP server in a local subnet to configure the IP information, and the access point did not receive a default router (gateway) address as a result, Without a default router, the access point cannot find the observer.

TIP: Do not specify an observer associated with the access point with an assigned snooping filter. This configuration causes an endless cycle of snooping traffic.

Table 249: RF Snooping Settings (continued)

Field	Description
Task: Select an existing snooping observer	<ol style="list-style-type: none"> 1. Click Select. The Select Snooping Observer window opens. 2. Place a check mark next to one of the snoop observers listed in the Select Snooping Observer window. 3. Click Select in the Select Snooping Observer window. The selection window closes and the selected observer now appears in the Observer Name field.
Task: Create a new snooping observer	<ol style="list-style-type: none"> 1. Click Create. The Create Snooping Observer window opens. 2. Provide the following settings: <ul style="list-style-type: none"> • Name for the snooping observer 1-15 characters long. • Target IP Address of the snoop observer. • Optionally enable Snap Length Limit and provide a snap length in bytes. Snap length specifies the maximum number of bytes to capture. If you do not specify a length, the entire packet is copied and sent to the observer. Juniper Networks recommends specifying a snap length of 100 bytes or less. 3. Optionally, enable Frame Gap Limit and provide a frame gap value in milliseconds. Frame gap refers to the minimum time period allowed between transmission of packets. 4. Select a Mode for the snoop observer: <ul style="list-style-type: none"> • tzsp: When an 802.11 packet matches all conditions in a filter, the access point encapsulates the packet in a Tazmen Sniffer Protocol (TZSP) packet and sends the packet to the observer host IP addresses specified by the filter. TZSP uses UDP port 37008 for transport. • batched tzsp: Every captured packet matching a configured filter will be encapsulated in a TZSP header and snoop record header and marshaled in a buffer maintained on a per snoop observer basis, for later transmission. The marshaled snoop packets will be copied to a UDP datagram and sent to the observer as soon as the buffer size exceeds the UDP datagram data field size. 5. Click Done in the Create Snoop Observer window. The window closes and the created observer now appears in the Observer Name field.

Table 249: RF Snooping Settings (continued)

Field	Description
-------	-------------

Snooping Conditions

The snooping conditions specify the match criteria for packets. Conditions in the list are appended. Therefore, to be copied and sent to an observer, a packet must match all snooping conditions. You can specify up to eight of the following conditions in a filter, in any order or combination:

- Frame Type
- Channel
- BSSID
- Transmitter Type
- Source MAC Address
- Destination MAC Address
- MAC Host
- MAC Pair
- Direction

Task: Add a Snooping Condition

A snooping condition consists of three parts, a Type, an Operation, and a Direction. The end result resembles an equation. To create a condition, Click **Add** under Snooping Conditions. The Create Snooping Condition window opens with three conditions, a Type, an Operation, and a third attribute. You can create eight of the following combinations:

Type	Operation	Third Attribute
Frame Type	Equals or Not Equals	Frame Types: Management, Control, Data, Beacon, or Probe
Direction	Equals or Not Equals	Direction: Receive or Transmit
Channel	Equals or Not Equals	Channel: 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14
BSSID	Equals or Not Equals	When the operation type is Glob the fields OUID and Vendor Name are available. <ul style="list-style-type: none"> • Glob—Indicate a host address. An asterisk like this * (the default) is a wildcard meaning all hosts. • OUID—Indicate a vendor's OUID code.
Transmitter Type	Equals or Not Equals	Transmitter Type: Member AP

Table 249: RF Snooping Settings (continued)

Field	Description	
Source MAC,	<ul style="list-style-type: none"> • Equals—EQ • Not Equals—NEQ • Less than—LE • Greater than—GE • Wild card pattern— Glob 	<p>When the operation type is Glob the fields OUID and Vendor Name are available.</p> <ul style="list-style-type: none"> • Glob—Indicate a host address. An asterisk like this * (the default) is a wildcard meaning all hosts. • OUID—Indicate a vendor's OUID code.
Destination MAC	<ul style="list-style-type: none"> • Equals—EQ • Not Equals—NEQ • Less than—LE • Greater than—GE • Wild card pattern— GLOB 	<p>When the operation type is Glob the fields OUID and Vendor Name are available.</p> <ul style="list-style-type: none"> • Glob—Indicate a host address. An asterisk like this * (the default) is a wildcard meaning all hosts. • OUID—Indicate a vendor's OUID code.
MAC Host	<ul style="list-style-type: none"> • Equals—EQ • Not Equals—NEQ • Less than—LE • Greater than—GE • Wild card pattern— GLOB 	<p>When the operation type is Glob the fields OUID and Vendor Name are available.</p> <ul style="list-style-type: none"> • Glob—Indicate a host address. An asterisk like this * (the default) is a wildcard meaning all hosts. • OUID—Indicate a vendor's OUID code.
MAC Pair	Type two MAC addresses	

Task: Edit a Snooping Condition

Task: Edit a Snooping Condition	<ol style="list-style-type: none"> 1. Select an existing snooping condition from the list of Snooping Conditions on the Create RF Snooping Filter Profile page. 2. Click Edit. The Edit Snooping Condition window opens with the Type, Operation, and Direction of the condition displayed. 3. Make any needed changes to the condition and then click Done. The Edit Snooping Condition window closes and the condition is updated in the Snooping Condition list.
---------------------------------	---

Task: Delete a Snooping Condition

Table 249: RF Snooping Settings (continued)

Field	Description
Task: Delete a Snooping Condition	<ol style="list-style-type: none"> 1. Select an existing snooping condition from the list of Snoop Conditions on the Create RF Snooping Filter Profile page. 2. Click Delete. <p>The condition disappears from the Snoop Condition list.</p>

NOTE: The AP running a snooping filter forwards snooped packets directly to the observer. This is a one-way communication, from the AP to the observer. If the observer is not present, the access point still sends the snooped packets, which uses bandwidth. If the observer is present but is not listening to TZSP traffic, the observer continuously sends ICMP error indications back to the access point. These ICMP messages can affect network and access point performance.

What To Do Next

Assign the Snooping Filter Profile to an access point following the directions in [“Assigning RF Snooping Filter Profiles to Access Points”](#) on page 1132. You can also map a Radio profile to a snooping profile—see [“Creating and Managing a Radio Profile”](#) on page 932.

RELATED DOCUMENTATION

[Assigning RF Snooping Filter Profiles to Access Points | 1132](#)

[Network Director Documentation home page](#)

Assigning RF Snooping Filter Profiles to Access Points

IN THIS SECTION

- [Assign an RF Snooping Profile to Access Points | 1133](#)
- [What To Do Next | 1134](#)

RF snooping refers to access points listening for any kind of RF communication on the network—this topic describes assigning a snooping configuration. You must have an existing RF Snooping profile to assign to devices—for directions, see [“Creating and Managing RF Snooping Filter Profiles” on page 1125](#).

TIP: You can also map a Radio profile to a snooping profile—see [“Creating and Managing a Radio Profile” on page 932](#).


You can map up to eight snoop filters to a radio. A filter does not become active until you enable it. Filters and their mappings are persistent and remain in the configuration following a restart. The filter state is also persistent across restarts. Once a filter is enabled, if the controller or the access point is subsequently restarted, the filter remains enabled after the restart. To stop using the filter, you must manually disable it.

Assign an RF Snooping Profile to Access Points

To assign an RF Snooping profile to access points, follow these steps:

1. Under Views, select one of these options: **Logical View**, **Location View**, **Device View** or **Custom Group View**.

TIP: Do not select **Dashboard View**, **Datacenter View**, or **Topology View**.

2. Click  in the Network Director banner.
3. In the Tasks pane, expand **Wireless**, expand **Profiles**, and then click **RF Snooping**.
The Manage RF Snooping Profile page appears, displaying the list of currently configured Snooping Profiles.
4. Select one of the listed profiles and then click **Assign**.
The Assign RF Snooping Profile wizard opens with three sections, **Device Selection** (displayed), **Profile Assignment**, and **Review**.
5. Under **Device Selection**, select one or more devices. Snooping profiles are assigned to access point radios. If you select a controller, all associated access points receive the Snooping profile.
6. Click either **Profile Assignment** or **Next**.

The Assign Snoop Profile page opens.

7. On the Assign Snoop Profile page, select any listed devices for Snooping profile assignment.
8. Click either **Review** or **Next**.
9. Make any needed changes and then click **Finish**.

The assignment is added to the Manage Snooping Profiles page.

What To Do Next

Next, deploy the devices with the new Snooping profile assignment—for directions, see “[Deploying Configuration to Devices](#)” on page 1180.

RELATED DOCUMENTATION

[Creating and Managing RF Snooping Filter Profiles | 1125](#)

[Deploying Configuration to Devices | 1180](#)

[Network Director Documentation home page](#)

Managing Network Devices

IN THIS CHAPTER

- Viewing the Device Inventory Page | 1136
- Viewing Device Connectivity | 1139
- Viewing Profiles Assigned to a Device | 1144
- Viewing the Physical Inventory of Devices | 1146
- Viewing Licenses With Network Director | 1147
- Viewing a Device's Current Configuration from Network Director | 1150
- Assigning Devices to Logical Category | 1150
- Accessing a Device's CLI from Network Director | 1151
- Accessing a Device's Web-Based Interface from Network Director | 1153
- Deleting Devices | 1154
- Rebooting Devices | 1155
- Viewing Virtual Machines | 1155
- Adding and Managing an Individual Access Point | 1156

Viewing the Device Inventory Page

The Device Inventory page lists devices managed by Network Director and provides basic information about the devices, such as IP address and current operating status. The Device Inventory page is available in Build and Deploy mode and is the default landing page for Build mode.

The scope you have selected in the View pane and the network view that you have selected from the View selector determines which devices are listed in the Device Inventory page. For example:

- If you are in the Logical View and select My Network, all devices managed by Network Director are listed.
- If you are in the Datacenter View and select Network Devices under a data center, only the network devices that are part of that data center are displayed.
- If you select a building in Location view, only those devices assigned to that building (including the floors and closets in the building) are listed.
- If you select a wireless LAN controller, only that wireless LAN controller and any access points it manages are listed.

NOTE: If you have configured access points using the Manage Access Point task but have not yet deployed the configuration on the wireless LAN controller, the Device Inventory page does not list those access points.

The Device Inventory page provides three pie charts that summarize the status of the devices in your selected scope:

- Devices by Family—Indicates the proportion of devices in each device family.
- Connection State—Shows the proportion of devices that are up or down. In this chart, Virtual Chassis count as one device.
- Configuration State—Shows the proportion of devices in each configuration state. See the Config State entry in [Table 250](#) for definitions of the configuration states.

Mouse over a pie segment to view the actual number of devices and the percentage represented by that pie segment.

[Table 250](#) describes the fields in the Device Inventory table.

Table 250: Fields in the Device Inventory Table

Field	Description
Hostname	Configured name of the device or IP address if no hostname is configured.

Table 250: Fields in the Device Inventory Table (continued)

Field	Description
IP Address	IP Address of the device.
Serial Number	Serial number of device chassis.
Platform	Model number of the device.
OS Version	Operating system version running on the device.
Device Family	<p>Device family of the device:</p> <ul style="list-style-type: none"> • JUNOS-EX for EX Series switches • JUNOS for Campus Switching ELS • JUNOS-QFX for QFX Series switches and QFabric members • JUNOS-QF for QFabric systems • MSSOS for WLC Series wireless LAN controllers • WLC-AP for access points managed by WLC Series wireless LAN controllers
Device Type	<p>Type of the device:</p> <ul style="list-style-type: none"> • AP—Wireless LAN access point • Fabric Member—QFabric member switch • QFabric—QFabric system • Switch—Standalone switch • VC—Virtual Chassis master • VC Member—Virtual Chassis member switch • WLC—Wireless LAN controller • XRE—External Routing Engine for EX8200 Virtual Chassis
Connection State	<p>Connection status of the device in Network Director:</p> <ul style="list-style-type: none"> • UP—Device is connected to Network Director. • DOWN—Device is not connected to Network Director. • N/A—Access point state is unavailable to Network Director.

Table 250: Fields in the Device Inventory Table (continued)

Field	Description
Config State	<p>Displays the configuration status of the device:</p> <ul style="list-style-type: none"> • In Sync—The configuration on the device is in sync with the Network Director configuration for the device. • Out Of Sync—The configuration on the device does not match the Network Director configuration for the device. This state is usually the result of the device configuration being altered outside of Network Director. You cannot deploy configuration on a device from Network Director when the device is Out Of Sync. To resolve this state, use the Resynchronize Device Configuration task in Deploy mode. • Sync failed—An attempt to resynchronize an Out Of Sync device failed. • Synchronizing—The device configuration is in the process of being resynchronized. • N/A—The device is down or is an access point.
Manageability State	<p>Displays if the device is directly manageable or not.</p> <p>This is a hidden field. To display the Manageability State field, click any column, click the down arrow to expand the list, select Columns from the list, and then enable Manageability State.</p>

NOTE: Juniper Networks devices require a license to activate the feature. To understand more about Network Director Licenses, see [“Viewing Licenses With Network Director” on page 1147](#).

RELATED DOCUMENTATION

[Understanding the Network Director User Interface | 85](#)

[Understanding Resynchronization of Device Configuration | 1214](#)

[Device Inventory Report | 1511](#)

[Viewing Licenses With Network Director | 1147](#)

[Network Director Documentation home page](#)

Viewing Device Connectivity

At the device level, you can view the connectivity details of a device and the details of all the devices that are connected to the specified device by using the Device Connectivity task in Network Director. The Device Connectivity page displays various details about a selected device and its immediate neighbors. The level of detail that Network Director displays in the Device Connectivity page depends on the type of device that you select.

To view the connectivity details of devices:

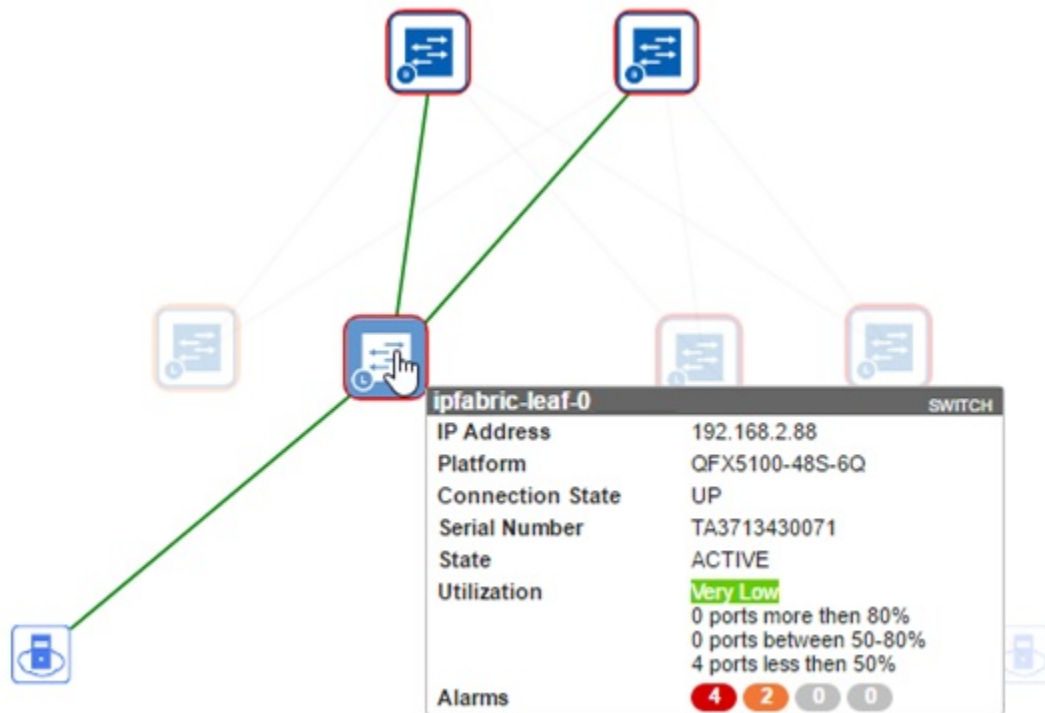
1. While in the Logical, Location, Device, Custom, or Topology View, under Build Mode, select the device for which you want to view connectivity from the View pane or the network topology (in case of Topology View) and click **Connectivity > View Device Connectivity** from the Tasks pane.

The Device Connectivity page opens. You can view the device connectivity details either in graph view or in grid view. The default view is the graph view.

In the graph view, each device and its network connectivity to all the connected devices are displayed as shown in [Figure 48](#). Mouse over a device to select a device and view details of the device.

NOTE: If the selected device is connected to a device that is not managed by Network Director, the latter appears dimmed in the Device Connectivity page.

Figure 48: Displaying Connection Details in Graph View



If the selected device is connected to more than sixty devices, then all the connected devices are highlighted in a circular form or a grid form. If the selected device is connected to less than 60 devices, then the links between the interconnected devices are displayed.

The device details displayed include name, IP address, and the alarm state information in colored labels that provide health and reachability information. You can also view the details of the hosts or virtual machines that are connected to the devices.


You can view the following details when you mouse over a device in the Device Connectivity—Graph view:

NOTE: The level of detail displayed depends on the type of device selected.

- Name—The name of the device provided while configuring the device. The device name and the device type are displayed in a label.
- IP Address—The IP address of the device.
- Platform—The device family and platform information. For example, EX4300-24P, EX4200-24T, QFX10002-36Q, and so on.
- Connection State—Connection state of the device. Connection state can be UP, DOWN, or N/A. Network Director updates and displays the connection state changes in real time.

- Config State—Device state. The device state can be online/provisioned or down/provisioned.
- Serial Number—Serial number of the device.
- Link status—Indicates whether the link between two devices is UP or DOWN. Network Director updates and displays the link status changes in real time. It might take up to 2 minutes for the updates to reflect.
- Utilization—Overall color-coded bandwidth utilization level and the breakup of bandwidth utilization by each port on the device.
- Alarms—Alarm details displaying the number of critical, major, minor alarms, or info for the device. Alarms details are color coded to indicate their severity level. Network Director updates and displays the alarm status changes in real time.
- Slot Number—(Applicable to Junos Fusion satellite devices) The FPC identifier of the satellite device in Junos Fusion. Slot number ranges from 65 through 255, and functions as the FPC identifier in the interface name when satellite device interfaces are being configured.
- LAG—Identifies connections that are configured as LAGs as shown in the following figure.

Figure 49: LAGs in the Device Connectivity view



You can view the following details of virtual machines (VMs) that are connected to hosts:

- Virtual Machine—Name of the virtual machine.
- Host Name—Name of the host to which the virtual machine is connected to.
- VNetwork—Name of the virtual network.
- OS—Name of the operating system on which the virtual machine is running.
- Connection State—Connection status of the virtual machine. Connection state can be UP, DOWN, or N/A.
- Power State—State of the power supply: Powered On or Powered Off.

You can view the following details of the Desktop machine:

- Host Name—Name of the host to which Desktop machine connected to.
- OS—Name of the operating system on which the Desktop machine is running.
- Connection State—Connection status of the Desktop machine. Connection state can be UP, DOWN, or N/A.

In the graph view, each networking device has a unique icon for easy identification. [Table 251](#) describes the networking device that each of these icons indicate.

Table 251: Icons on the Device Connectivity Page










Icon	Description
	Juniper Networks switch
	Desktop computer
	Desktop IP phone
	Wireless LAN controller
	WLAN access point
	Printer
	Satellite device cluster in a Junos Fusion system. Double-click this icon to view the connectivity of the devices that are part of the cluster.

Table 251: Icons on the Device Connectivity Page (continued)

Icon	Description
	Hypervisor server (KVM or ESXi host).
	Bare metal server

- You can perform the following tasks from the graph view of the Device Connectivity page:
 - Click the number adjacent to a device icon to view details about the device alarms. The Alarm Details by Severity page opens. For more details, see [“Alarm Detail Monitor” on page 1459](#).
For example, in [Figure 50](#) click the number 8 to view details of the alarms on the corresponding device.

Figure 50: Alarm Count in the Device Connectivity Page



- Double-click the satellite device cluster icon to expand and view the member satellite devices and their connectivity.
 - Click **Links** and select **Color Code Port Utilization** to view the color-coded port utilization level in the graph view. Network Director displays a port utilization legend in the upper right corner of the graph view, which you can use to identify links that are optimally used, overutilized, or underutilized and take necessary corrective actions.
 - Select **Stop Updates**, to freeze the link status changes in real time in the Device Connectivity page that might be required while the user is performing some tasks in this page.
 - Enter the device name, IP address, or the tag name of the device in the search field to quickly locate a device in the graphical view.
- Click **Show Grid View** to view the device connectivity details in a tabular format as displayed in [Figure 51](#). This view has two tabs: the External Links tab and the Fabric Links tab. Clicking the **External Links** tab displays the external device interface details that are connected to the fabric devices. Clicking the **Fabric Links** tab displays fabric link interface details.

Figure 51: Displaying the Connection Details in Grid View

Device Connectivity : nd-72q1-ellit						
External Links Fabric Links						
Show Graph View						
Source Device	Source Port	Source Port Bandw...	Destination Device	Destination Port	Destination Port Bandw...	Link Status
nd-72q1-ellit	[LAG] ae0	NA	nd-36q1-ellit	[LAG] ae0	NA	↑
nd-72q1-ellit	[LAG] ae1	NA	nd-36q1-ellit	[LAG] ae1	NA	↓
nd-72q1-ellit	et-0/0/0 (ae0)	0	nd-36q1-ellit	et-0/0/0 (ae0)	0	↑
nd-72q1-ellit	et-0/0/1 (ae1)	0	nd-36q1-ellit	et-0/0/1 (ae1)	0	↓
nd-72q1-ellit	et-0/0/2 (ae2)	0	nd-opus-48s4	et-0/0/48	0	↑

The following details are displayed in the grid view:

- Source Device—Name of the device specified while configuring the device.
- Source Port—Source port of the device.
- Source Port Bandwidth %—Real-time percentage of bandwidth utilized at the source port.
- Destination Device—Name of the destination device or devices the source device is connected to.
- Destination Port—The port number on the destination device to which the source device is connected to.
- Destination Port Bandwidth %—Real-time percentage of bandwidth utilized at the destination port.
- Link Status—Indicates whether the link to the device is up or down.

You can sort the details in the table in the ascending order or descending order for each column. You can also use filters to display device connectivity details for specific devices. If you type a text string and click **Go**, entries that do not contain the text string (filter criterion) are removed from the table

RELATED DOCUMENTATION

| [Setting Up the Topology View](#) | 250

Viewing Profiles Assigned to a Device

View Assigned Profile page list all the profiles associated with a selected device or with an object such as ports, access points, or radios within that device. To view the profiles assigned to a device, you must have the profiles already assigned to the devices, ports, access points, or radios within that device. Only those profiles that are assigned to a specified object will be displayed in the Profiles Assigned to the Device page. In addition to displaying profiles assigned to objects, the Profiles Assigned to Device page also shows link aggregation groups (LAGs) assigned to devices.

The View Assigned Profiles task is available in the Logical, Location, and Device panes for EX Series switches. For wireless LAN controllers, the View Assigned Profiles task is available only in the Logical pane as cluster functionality is available only in Logical pane.

You can access the View Assigned Profiles page by selecting the object (device, port, or radio) and clicking the View Assigned profiles menu. However, view assigned profiles task is unavailable for auto access point (Auto AP) profiles.

You can view the profiles assigned to an EX Series switch or a wireless LAN controller. To view the assigned profiles to a particular device:

While in Build mode, select an EX Series switch from the Switching Network cabinet or a controller from the Wireless Network cabinet under the View pane and select **View Assigned profiles** from the Tasks pane.

The Profiles Assigned to the Device page displays a list of profiles that are already assigned to the selected device. The details displayed are described in [Table 252](#).

Table 252: Details of Assigned Profiles to a Device

Field	Description
Profile Type	The type of the profile. The profiles are grouped based on the type of the profile. The profiles that are directly deployed on the device are displayed in the list. For example, Device, VLAN, Port, Radio and so on.
Profile Name	Name of the profile that was specified at the time of creating the profile.
Object	The name of the device (EX Series switch or a wireless LAN controller), port, access point, or radio.
Object Type	Specifies whether the object is a device (EX Series switch or a controller), a port, an access point, or a radio. TIP: For an EX Series switch, the object type is the device or a port For a wireless LAN controller, the object type is a controller, an access point, a radio, or a port.
Assignment State	The status of the profile whether it is deployed or in progress. <ul style="list-style-type: none"> • Deployed—the profile is provisioned to the device • Pending deployment—the profile is assigned to the device, but pending provisioning.

RELATED DOCUMENTATION

[Assigning Device Common Settings to Devices | 331](#)

[Assigning a VLAN Profile to Devices or Ports | 531](#)

[Network Director Documentation home page](#)

Viewing the Physical Inventory of Devices

You can view the physical inventory of all the devices in your network in the Device Physical Inventory page. The Device Physical Inventory page displays information about the slots that are available for a device and provides information about power supplies, chassis cards, fans, part numbers, and so on. Network Director displays hardware inventory by device name, based on data retrieved both from the device during discovery and resynchronizing operations, and from the data stored in the hardware catalog. For each managed device, the physical inventory page provides descriptions for field replaceable units (FRUs), part numbers, model numbers, and the pluggable locations from which empty slots are determined.

To view the Device Physical Inventory page, while in the Build mode select a standalone EX Series switch, a Virtual Chassis, or a wireless LAN controller (WLC) from the View pane and select **Device Management > Physical Inventory** from the Tasks pane.

The physical inventory page displays the model number, part number, serial number, and description for the following, depending on the device that you selected:

- For standalone EX Series switches and Virtual Chassis, the page displays details of the switch, the chassis, the Flexible PIC Concentrator (FPC), the PIC slot, the PIC installed in the PIC slot, the power supply, the fan tray, and the routing engine.
- For a controller, the page displays details about the fan and optics used by the controller.

You can view the following details from the Device Physical Inventory page as described in [Table 253](#).

Table 253: Fields in the Device Physical Inventory Table

Field	Description
Item	Name of the device and the components that are part of the device. By default, Network Director displays the device and components in an expanded tree structure. You can click a device or component to collapse or expand the sub-components.
Model Number	<ul style="list-style-type: none"> • For standalone EX Series switches and Virtual Chassis, the full Junos EX Series model number of the device. • For a controller, the model number of the controller.
Part Number	Part number of the EX Series switch chassis component.
Serial Number	The hardware serial number of the device.

Table 253: Fields in the Device Physical Inventory Table (continued)

Field	Description
Description	The description about the component.

NOTE: Juniper Networks devices require a license to activate the feature. To understand more about Network Director Licenses, see [“Viewing Licenses With Network Director” on page 1147](#). Please refer to the product Data Sheets for further details, or contact your Juniper Account Team or Juniper Partner.

RELATED DOCUMENTATION

[Understanding the Network Director User Interface | 85](#)

[Viewing Licenses With Network Director | 1147](#)

[Network Director Documentation home page](#)

Viewing Licenses With Network Director

Juniper Networks devices require a license to operate some features. You can view the licenses for devices connected to Network Director.

To view the license for a Juniper Networks device on your network:

1. Select the **Build** icon in the Network Director banner.
2. In the View pane, select a wireless or wired device.
3. In the Tasks pane, select **View License Information**.

The Licenses page for that object is displayed with the fields listed in [Table 254](#).

Table 254: Viewing Licenses with Network Director

Field	Description
Feature Name	Name of the licensed SKU or feature. It can be used to look up the license with Juniper Networks. Not all devices support this.

Table 254: Viewing Licenses with Network Director (continued)

Field	Description
License Count	Number of times an item has been licensed. This value can have contributions from more than one licensed SKU or feature. Alternatively, it can be 1, no matter how many times it has been licensed.
Used Count	Number of times the feature is used. For some types of licenses, the license count will be 1, no matter how many times it is used. For capacity-based licensable items, if infringement is supported, the license count can exceed the given count, which has a corresponding effect on the need count.
Need Count	Number of times the feature is used without a license. Not all devices can provide this information.
Given Count	Number of instances of the feature that are provided by default.

NOTE: If a device does not have a license, a blank page is displayed with the message, **No license is installed on this device.** If you are sure the device has a license, try resynchronizing the device before displaying the license again.

- Optionally, expand the license information by feature name to view the feature SKU information. [Table 255](#) describes the additional fields that are displayed.

Table 255: Additional Licensing Information

Field	Description
Validity Type	Validity type can be Databased (license expires on end date), Permanent, Countdown (license expires when time remaining is zero), or Trial. If the validity type is either Databased or Countdown, more information is displayed—License Name, License Version, License State, and Time Remaining. Additional information can be added in the details grid based on the SKU type (SKU or Feature)—Start Date, End Date, or Original Time Allowed.
License Name	If the validity type is either Databased or Countdown, the identifier associated with a license key is displayed.
License Version	If the validity type is either Databased or Countdown, the version of a license is displayed. The version indicates how the license is validated, the type of signature, and the signer of the license key.

Table 255: Additional Licensing Information (*continued*)

Field	Description
License State	If the validity type is either Databased or Countdown, the state of the license is displayed—Valid, Invalid, or Expired.
Time Remaining	If the validity type is either Databased or Countdown, the remaining time left on the license is displayed. For a trial license, the number of days remaining after you installed the device is displayed. For a commercial license, the time remaining is unlimited.
Start Date	Based on the SKU type, the start date of the license can be displayed in the details grid.
End Date	Based on the SKU type, the end date of the license can be displayed in the details grid.
Original Time Allowed	Based on the SKU type, the original license timeframe can be displayed here.

NOTE: If you apply a new license to an existing wireless LAN controller, you must resynchronize the device before the new license is seen in Network Director. For directions, see [“Resynchronizing Device Configuration” on page 1220](#).

RELATED DOCUMENTATION

[Resynchronizing Device Configuration | 1220](#)

[Network Director Documentation home page](#)

Viewing a Device's Current Configuration from Network Director

You can view a device's current configuration from Network Director. This is a convenient way to view device configurations without leaving Network Director.

To view a device's current configuration:

1. Click **Build** or **Deploy** in the Network Director banner.
2. Select the device in the View pane.
3. Select **Device Management > Show Current Configuration** in the Tasks pane.
4. The device's current configuration displays in the main window.

RELATED DOCUMENTATION

[Understanding the Network Director User Interface | 85](#)

[Understanding the Build Mode Tasks Pane | 189](#)

[Network Director Documentation home page](#)

Assigning Devices to Logical Category

Network Director enables you to tag the available EX Series switches to different categories—Core, Aggregation, Access, or Unassigned categories. Once you tag a switch, it appears under the tagged category after refreshing the screen.

To assign a group of EX Series switches under Core, Aggregation, Access, or Unassigned category, select Logical View in Build mode (available only in Build mode) and select one of the cabinets under the Switching Network cabinet:

1. Select **Assign Device to Logical Category** under Device Management from the Tasks pane.
The Assign Device to Logical category page is displayed.
2. Click **Add** from the Selected Devices table. The Please select devices dialog box is displayed.
3. Select the device or devices that you want to assign by selecting the check box next to the hostname. Click **OK**. The selected device or devices with the details appears in the Selected Devices table.

4. Select a role from the New Role list.

The available roles are: Access, Aggregation, Core, and Unassigned.

5. Click **Done** to change the role or click **Cancel** if you do not want to change the role. The message: **Device Role successfully changed** appears if you have selected Done.

To assign an EX Series switch under Core, Aggregation, Access, or Unassigned category, select Logical View in Build mode (available only in Build mode) and select any EX Series switch from the one of the cabinets under the Switching Network cabinet:

1. Select **Assign Device to Logical Category** under Device Management from the Tasks pane.

The Assign Device to Logical category page is displayed. The page displays the name of the selected device and the current role of the device.

2. Select a role from the New Role list to change the current role.

The available roles are: Access, Aggregation, Core, and Unassigned.

3. Click **Done** to change the role or click **Cancel** if you do not want to change the role. The message: **Device Role successfully changed** appears if you have selected Done.

RELATED DOCUMENTATION

[Viewing Profiles Assigned to a Device | 1144](#)

[Network Director Documentation home page](#)

Accessing a Device's CLI from Network Director

Network Director enables you to connect to the CLI for switches and wireless LAN controllers in your network, using SSH.

This topic describes the steps to connect to a switch or a controller by using SSH (Secure Shell). SSH is a cryptographic network protocol used for remote shell services or command execution. SSH is one of the many access services that are supported on the Juniper Networks devices. All Juniper Network devices have SSH enabled by default.

To connect to a device by using SSH:

1. Do one of the following:

- In the View pane, select the device to which you want to connect.
 - In the Topology View, locate the device to which you want to connect.
2. Do one of the following:
 - With the device selected in the View pane, select **Build** mode and select **Tasks > Device Management > SSH to Device**.
 - While in the Topology View, select the device to which you want to launch the SSH connection and click **Device Management > SSH To Device**.

The SSH to Device dialog box appears.

3. Enter the username and password to connect to the selected device and click **Connect**.

NOTE: Ensure that you have removed Pop-Up blockers, if any, before you click Connect.

The SSH console to the switch or controller opens in a separate browser tab or window depending on your browser settings. Refer to the [EX Series documentation](#) or the [Wireless LAN Services \(WLS\) Product Documentation](#) for more information about using the CLI for EX Series switches and WLC wireless LAN controllers respectively.

NOTE: Any configuration changes you make to a device, using the CLI qualify as out-of-band changes in Network Director. Out-of-band configuration changes can cause the configuration state of a managed device to become out of sync, which indicates that the device configuration no longer matches the Build mode configuration for the device. Use the Resynchronize Device Configuration task in Deploy mode to resynchronize the device configuration.

RELATED DOCUMENTATION

[Understanding Resynchronization of Device Configuration | 1214](#)

[Accessing a Device's Web-Based Interface from Network Director | 1153](#)

[Network Director Documentation home page](#)

Accessing a Device's Web-Based Interface from Network Director

Network Director enables you to connect to the switches and wireless LAN controllers in your network, using the device Web-based interface.

This topic describes the steps to connect to a switch by using the J-Web interface or to a controller by using Web View. The J-Web interface is a graphical user interface, using which you can monitor, configure, troubleshoot, and manage switches. Web View is a web-based management application that enables you to perform common configuration and management tasks on wireless devices.

You can connect and configure a device by using the J-Web interface or Web View only if the device is configured to accept HTTP or HTTPS as a management service. You can configure HTTP or HTTPS as a management service using the Device Common Settings profile. For more information, see [“Creating and Managing Device Common Settings” on page 291](#).

To connect to a device using the J-Web interface or Web View:

1. Do one of the following:
 - In the View pane, select the device to which you want to connect.
 - In the Topology view, locate the device to which you want to connect.
2. Do one of the following:
 - While selecting the device in the View pane, select Build mode and select **Tasks** pane > **Device Management** > **Launch Web View**.
 - While in the Topology View, select the device for which you want to launch the Web connection and click **Device Management** > **Launch Web View**.

The Web View or J-Web Login page appears.

3. Enter the username and password to connect to the selected switch and click **Login**.

If the credentials that you entered are valid, the system displays the J-Web or Web View home page for the selected device.

NOTE: Any configuration changes you make to a device using the Web interface qualify as out-of-band changes in Network Director. Out-of-band configuration changes can cause the configuration state of a managed device to become out of sync, which indicates that the device configuration no longer matches the Build mode configuration for the device. Use the Resynchronize Device Configuration task in Deploy mode to resynchronize the device configuration.

RELATED DOCUMENTATION

[Understanding Resynchronization of Device Configuration | 1214](#)

[Accessing a Device's CLI from Network Director | 1151](#)

[Network Director Documentation home page](#)

Deleting Devices

You can delete devices that are no longer used from Network Director. Deleting a device removes all device configuration and device inventory information from the Junos Space database. Once a device is deleted from the database, all the profiles associations, device configurations, and inventory information of the deleted device are also deleted. However, the system maintains the audit logs and monitoring data for the device even after the device is deleted.

Use the Delete Devices page to delete devices from Network Director. While in Build mode, click **Delete Devices** from the **Tasks > Device Management** menu. The Delete Devices page appears.

The Delete Devices page displays the devices contextually depending on your selection in the View pane. For example, if you select a site in Location view and click Delete Devices, Network Director displays all the devices that are assigned to the buildings or floors in the selected site in the Delete Devices page. If you select a particular switch family in Device View and click Delete Devices, only switches that belong to that switch family are displayed.

To delete devices, complete the following tasks:

1. Select the check box adjacent to the switch, controller, or a group of access points that you want to delete.
2. Click **Done**.

Network Director prompts you to confirm the deletion. Click **Yes** to confirm the deletion or **No** to go back and make changes to the selection.

RELATED DOCUMENTATION

[Understanding the Network Director User Interface | 85](#)

[Discovering Devices in a Physical Network | 204](#)

[Viewing the Device Inventory Page | 1136](#)

[Network Director Documentation home page](#)

Rebooting Devices

Use the Reboot Devices task to immediately reboot the selected device. This task is available in all scopes when in Build mode. To reboot one or more devices immediately:

1. Select the scope in the View pane that contains the devices you want to reboot.
2. Select Reboot Devices from the Tasks pane.
3. Expand the tree on the page as needed to locate the available devices.
4. Select the check box for one or more devices.
5. Click **Done** to start the reboot or click **Cancel** to return to the Device Inventory page.

The rebooting process triggers a Cold Start Alarm that can be seen in Fault mode.

RELATED DOCUMENTATION

[Understanding the Build Mode Tasks Pane | 189](#)

[Understanding the Network Director User Interface | 85](#)

[Network Director Documentation home page](#)

Viewing Virtual Machines

EX Series switches (standalone and Virtual Chassis), QFX Series switches, and QFabric systems in your network can be connected to one or more ESX/ESXi hosts. Each host can have one or more virtual machines running on them.

You can use the View Virtual Machine task to view details about virtual machines that are connected to a switch or a QFabric system.

To view the virtual machines

1. While in the Logical View with Build mode selected, select the standalone switch, virtual chassis or the QFabric system for which you want to view the connected hosts.
2. Click **Connectivity > View Virtual Machines** from the Tasks pane.

3. The View Virtual Machines table displays the details of the virtual machines that are connected to the selected switch or QFabric system. [Table 256](#) describes the fields in this table:

Table 256: Manage Virtual Machines Page Field Descriptions

Field	Description
Switch Port	The switch port on the physical switch or the QFabric system that is connected to the host.
Host	Name of the host on which the virtual machine is running.
Host NIC	The network adapter on the host that connects the physical switch to the host.
VLANs	The VLANs configured on the physical switch port.
Virtual Machines	The name of the virtual machines that are running on the given host. Mouse over this field to view the number of virtual machines that are running on the host.

RELATED DOCUMENTATION

[Viewing the Virtual Machine Inventory in a Cloud Infrastructure | 803](#)

[Network Director Documentation home page](#)

Adding and Managing an Individual Access Point

IN THIS SECTION

- [Managing Access Points | 1157](#)
- [Adding an Access Point to a Wireless Network | 1159](#)
- [Specifying Basic Access Point Settings | 1161](#)
- [Enabling Local Switching on an Access Point | 1163](#)
- [Configuring an Access Point as Remote | 1164](#)
- [Configuring Link Layer Discovery Protocol \(LLDP\) on an Access Point | 1165](#)
- [Configuring Bonjour on an Access Point | 1167](#)
- [Specifying Access Point Radio Settings | 1168](#)

You can explicitly configure access points on a controller with unique information to identify them. If you use this method to add all of your access points, it prevents rogue or neighbor access points from being added accidentally. Configured access points are persistent—the configuration is not lost if an access point loses contact with the controller.

You can also use this feature to tailor an access point that is using a Radio profile and WLAN Service profile. Any configurations you make here take precedence over the profile settings.

NOTE: The option to manage access points is only available when a controller or controller cluster is selected in the leftmost pane of Network Director and **Logical View** is the selected view.

TIP: The feature Auto AP adds access points to controllers temporarily—see [“Understanding Auto AP Profiles” on page 883](#) and [“Creating and Managing Wireless Auto AP Profiles” on page 980](#).

This topic describes how to use Network Director to add or modify specific access points.

Managing Access Points

The option to manage access points is only available when a controller or controller cluster is selected in the leftmost pane of Network Director and **Logical View** is the selected view. From the **Manage Access Points** page, you can:

- Manually add an access point to a controller by clicking **Add**. For detailed steps, see [“Adding an Access Point to a Wireless Network” on page 1159](#).
- Modify an existing access point configuration by selecting it and clicking **Edit**.
- Delete manually configured access points by selecting access points and then clicking **Delete**.

TIP: You can delete access points that are in use—the clients from the access point re-associate to another access point.

[Table 257](#) describes the information provided about access points on the Manage Access Points page, which is available only in Logical View. The access points displayed are those managed by the controller or the cluster selected in the View pane.

Table 257: Manage Access Point Fields

Field	Description
AP Name	Name given to the access point when the access point was created.
AP ID	Number from 1 through 9999 that identifies an access point.
Model	Juniper Networks model number of the access point
Serial Number	Serial number on the back of the access point
Fingerprint	<p>Access points are configured with an encryption key pair at the factory. The fingerprint for the public key is displayed on a label on the back of the access point, in the format: RSA a:aaaa:aaaa:aaaa: aaaa:aaaa:aaaa:aaaa</p> <p>TIP: This field might be blank because a fingerprint is optional during access point configuration.</p>
Connection	Distributed access points are connected to the network. Direct access points are connected to a controller port with no switch or router in between.

Adding an Access Point to a Wireless Network

You can explicitly configure access points so that the controller finds only those access points.

To add individual access points to Network Director:

1. Gather the following information:

- Access point model—The model is listed on the back of the access point.
- Access point IDs that have already been assigned, for example, 1, 2, and so on. You cannot repeat these numbers for new assignments.

TIP: Some controllers keep track of the IDs that are in use, and suggest an unused ID for assignment.

- Serial number located on the back of the access point.

TIP: Serial numbers never contain spaces.

- Optionally, for increased security, you can configure the access point fingerprint, which is an encryption key pair generated at the factory. The fingerprint for the public key is printed on a label located on the back of the access point in the following format: RSA a:aaa:aaa:aaa:aaa:aaa:aaa:aaa
- Optionally, if you added an external antenna or plan to add an external antenna to the access point, you need the model number of the antenna.

TIP: If you do not specify an antenna, the access point's internal antenna is used.

2. Under Views, select **Logical View**.

3. Click  in the Network Director banner.

4. Under the Wireless Network in the View pane, select either a controller or a cluster of controllers.

TIP: Do not select the whole wireless network, a custom group, or an individual access point—these selections are not viable for adding an access point, because access points cannot be added to an entire network nor to individual access points. You do not see **Manage Access Point** under Key Tasks unless you select a controller or a cluster of controllers.

5. Click **Manage Access Point** under Key Tasks in the Tasks pane.

The Manage Access Points page is displayed with a list of all access points configured for the selected controller or controller cluster.

6. Click **Add**.

The Add AP window opens with the **Access Point** tab selected.

7. On the **Access Point** tab, complete the basic access point settings as described in both the online help and in [“Specifying Basic Access Point Settings” on page 1161](#).

8. Optionally, assign the access point to a local switching profile by clicking the tab **Local Switching** and completing the settings described in both online help and [“Enabling Local Switching on an Access Point” on page 1163](#).

9. Optionally, configure the access point as a remote access point by clicking the tab **Remote WLA** and completing the settings described in both the online help and in [“Configuring an Access Point as Remote” on page 1164](#).

10. Optionally, reconfigure the default LLDP or LLDP-MED settings for the access point by clicking the tab **LLDP** and completing the settings described in both the online help and in [“Configuring Link Layer Discovery Protocol \(LLDP\) on an Access Point” on page 1035](#).

11. Click the tab **Radio 1** and reconfigure any radio information as described in both the online help and in [“Specifying Access Point Radio Settings” on page 1168](#).

12. If it is displayed, click the **Radio 2** tab and reconfigure that radio information as described in both the online help and in [“Specifying Access Point Radio Settings” on page 1168](#).

13. Click **OK**.

The access point is added to the Manage Access Points list on the screen.

14. Add as many access points as needed, and then click **Done**.

The access points you added now appear on the inventory list, but have not yet been added to the network—the access points become part of the network when you deploy the controller.



CAUTION: Be sure to click **Done** to add each access point—otherwise, the configuration is lost.

15. Deploy the controller or controller cluster associated with the access points, following the directions in [“Deploying Configuration to Devices” on page 1180](#).

The access points become part of the network when you deploy the associated controller or cluster.

Specifying Basic Access Point Settings

To configure an access point, provide the basic access point information listed in [Table 258](#). Required settings are indicated by a red asterisk (*) that appears next to the field label in the user interface.

Table 258: Basic Access Point Configuration

Field	Description
AP Name	Type a unique name that identifies the access point, using up to 32 characters. Names must not contain special characters or spaces. Note that access point names automatically created by Network Director as part of device discovery or out-of-band changes might contain the underscore (_) character.
AP ID	An unassigned ID number appears by default. You can replace it with any ID number from 1 through 99999 that has not yet been assigned.
Remote Site	If the access point location is in a remote site (connected by a WAN link to the central network), select that Remote Site Profile from the list. TIP: Remote Sites are created by following the directions in “Creating and Managing Remote Site Profiles” on page 1014 .
Model	An access point model appears by default. Replace it with any access point model from the list—the list is based on the country code of the remote site.
Connection	Distributed access points (default) are connected to the network, usually through a switch. Direct access points are connected to a controller port with no switch or router in between. Ports depend on the controller model. Also, if the port is already connected to another direct access point, that port is not listed.
Serial Number	Type the serial number found on the back of the access point.

Table 258: Basic Access Point Configuration (continued)

Field	Description
Fingerprint	You can type the fingerprint for the access point public key, which is displayed on a label on the back of the access point, in the following format: RSA a:aaaa:aaaa:aaaa: aaaa:aaaa:aaaa:aaaa
Bias (default is high)	<p>If the access point is directly connected to the selected controller, bias does not matter. An access point always attempts to boot on AP port 1 first, and if a controller is directly attached on AP port 1, the access point boots from there regardless of the bias settings.</p> <p>Bias matters for access points indirectly connected to the controller through an intermediate Layer 2 or Layer 3 network. If the access point is indirectly connected to the selected controller, indicate that the access point has a high bias for the controller (default), has a low bias, or is sticky which means it will continue to use the current controller for the active data link even if another controller configured with high bias becomes available.</p> <p>For more information, see “Understanding Access Point Bias for Controllers” on page 852.</p>
Enable Firmware Update (enabled by default)	When the controller receives a later version of access point firmware, the access point will be updated unless you remove the check mark to disable updates.
Force Image Download (disabled by default)	Select for automatic image updates. When the controller receives a later version of the access point image, the access point will be updated.
Enable Blink (disabled by default)	Blink mode makes an access point blink so that it is easy to identify. When blink mode is enabled, the health and radio LEDs alternately blink green and amber. By default, LED blink mode is disabled. Once enabled, blink mode continues until you disable it. Changing the LED blink mode does not alter operation of the access point. Only the behavior of the LEDs is affected.
LED Mode (automatic by default)	Set LED mode either to auto (default) to have LEDs operate normally, or to static to have the LEDs operate with normal flashing patterns converted to a static On pattern. You can also turn the LEDs off to disable flashing.
Description	Enter up to 256 characters to describe the access point.
Location	Enter up to 256 characters to describe the location of the access point.
Contact	Enter up to 256 characters of contact information for a network administrator.

Table 258: Basic Access Point Configuration (*continued*)

Field	Description
WLA Communication Timeout (default is 25 seconds)	Number of seconds (default is 25) that the access point times out after the last communication.
Power Mode	On some access point models, power mode can be set to auto or high .
Antenna Mode	Antenna mode is derived from the access point model. Options are displayed only for access points with external antennas. Access points with internal antennas only have antenna mode set to Internal .
Enable Data Security (default is disabled)	Check to configure an access point for data path encryption. In cluster mode, the primary seed (PS) ensures that an access point with data security enabled is not assigned a primary access manager (PAM) or secondary access manager (SAM) that does not support this feature. If such a controller is not located on the network, the access point remains unassigned.
High Latency Mode (default is disabled)	Check to enable high latency mode. Bandwidth and latency are two elements that affect network speed. Latency refers to delays typically encountered when processing network data. A low latency network connection is one that has small delay times, while a high latency network typically encounters long delays. High latency mode configures attributes that can mitigate the association problems on a high latency network.

Next, optionally click the **Local Switching** tab to enable local switching for the access point. For directions, see [“Enabling Local Switching on an Access Point” on page 1163](#).

Enabling Local Switching on an Access Point

You can enable Local Switching on the access point. Local Switching means that packets switch directly from access points to the wired network instead of passing through a controller. For information about local switching, see [“Understanding Local Switching on Access Points” on page 907](#).

To configure Local Switching settings on the access point from the **Local Switching** tab of the Add AP process, complete the settings in [Table 259](#).

Table 259: Assign Local Switching Profile to Access Point

Field	Description
Enable Local Switching (disabled by default)	Select this option to have packets switch directly from access points to the wired network.
	Enable WLA Tunneling: WLA tunneling extends the WLC-WLC tunnel feature to allow access points that are using local switching to create and terminate client VLAN tunnels. This eliminates the need for a VLAN on every access point.
	Tunnel Affinity: Indicate a tunnel affinity from 0 through 10. When tunneling is enabled, an access point can create a tunnel to a controller or to another access point, depending on the configured tunnel affinity. By default, a VLAN on the controller has an affinity of 5 and a VLAN on an access point has an affinity of 4. If the tunnel affinity is the same, then the node with the lowest load is selected as a tunnel endpoint.

Next, optionally click the **Remote WLA** tab to make the access point a remote access point. For directions, see [“Configuring an Access Point as Remote” on page 1164](#).

Configuring an Access Point as Remote

You can mark the access point as remote, which means that the access point is connected by a WAN link to the central network. For information about remote access points, see *Understanding Remote Access Points*.

To configure an access point as a remote access point from the **Remote WLA** tab of the Add AP process, configure the settings shown in [Table 260](#).

Table 260: Configuring a Remote Access Point

Field	Description
Enable Remote WLA (disabled by default)	Select Enable Remote WLA to mark the access point as remote. A check mark appears in the corresponding box and the remaining options become available.
	Outage Duration: Number of hours between periodic checks of the state of the controller connection.
	Connection Evaluation Period: Maximum number of seconds that the access point remains in outage mode before rebooting.
Path MTU (default is zero)	Indicate the size of the maximum transmission unit (MTU), which is the largest packet that a network protocol can transmit. Default is zero.

Next, optionally click the **LLDP** tab to assign Link Layer Discovery Protocol to the access point. For directions, see [“Configuring Link Layer Discovery Protocol \(LLDP\) on an Access Point” on page 1035](#).

Configuring Link Layer Discovery Protocol (LLDP) on an Access Point

Link Layer Discovery Protocol (LLDP) is a link layer protocol used by network devices to advertise identity, capabilities, and neighbors. It also provides additional TLVs for capabilities discovery, network policy, Power over Ethernet (PoE), and inventory management.

Link Layer Discovery Protocol-Media Endpoint Discovery (LLDP-MED) is an extension to LLDP that operates between endpoint devices such as IP phones to provide support for voice over IP (VoIP) applications. LLDP-MED endpoints determine the capabilities of a connected device and whether those capabilities are enabled.

TIP: LLDP and LLDP-MED cannot operate simultaneously on a network. By default, access points send only LLDP packets until LLDP-MED packets are received from an endpoint device. The access point then sends out LLDP-MED packets until it receives LLDP packets.

For information about LLDP, see [“Understanding LLDP and LLDP-MED” on page 930](#).

To configure LLDP on the access point from the **LLDP** tab of the Add AP process, complete the configuration described in [Table 227](#).

Table 261: LLDP Settings for an Access Point

Field	Description
LLDP Mode (enabled by default)	<p>Link Layer Discovery Protocol (LLDP) on an access point is enabled for transmission by default, which means that access point transmits its identity, capabilities, and neighbors with LLDP. If you do not want the access point to use LLDP, disable LLDP by selecting Disable instead of transmit (TX).</p> <p>NOTE: You can enable both LLDP Mode and LLDP-MED Mode, but only one can operate at a time. By default, network devices send only LLDP packets until LLDP-MED packets are received from an endpoint device. The network device then sends out LLDP-MED packets until it receives LLDP packets.</p>

Table 261: LLDP Settings for an Access Point (continued)

Field	Description
LLDP-MED Mode (enabled by default)	<p>Media Endpoint Discovery is an enhancement of LLDP that is disabled by default. To have the access point transmit its identity, capabilities, and neighbors with the LLDP-MED protocol, enable LLDP-MED Mode and complete the two configuration options for LLDP-MED.</p> <p>NOTE: You can enable both LLDP Mode and LLDP-MED Mode, but only one can operate at a time. By default, network devices send only LLDP packets until LLDP-MED packets are received from an endpoint device. The network device then sends out LLDP-MED packets until it receives LLDP packets.</p> <hr/> <p>Power via MDI: Enable Power via Media Dependent Interface (MDI) to have the access point also convey power information, such as the type of power, power priority, and the amount of power required by the device. Information is collected on the Ethernet interface.</p> <hr/> <p>Inventory: Enable Inventory to have the access point also transmit detailed inventory information to a controller. Inventory information includes hardware revision, firmware version, software version, serial number, manufacturer name, model name, and asset ID.</p>

TIP: You can also configure LLDP on a controller—see [“Configuring a Controller”](#) on page 1037.

Next, click the **Radio 1** tab to configure the access point’s radio information. For directions, see [“Specifying Access Point Radio Settings”](#) on page 1168.

SEE ALSO

[Understanding LLDP and LLDP-MED | 930](#)

[Configuring a Controller | 1037](#)

Configuring Bonjour on an Access Point

Bonjour Zero configuration IP networking enables users to find printers, network resources, or music sharing on a network. If you are running Bonjour on your network, users can instantly find printers, or a friend's network game or music device, and share those files with someone else. These Bonjour services are available in Network Director—Apple TV, Internet printer, or Digital Auto Access Protocol (iTunes).

For more information, see [“Understanding Bonjour” on page 995](#).

To configure Bonjour on the access point from the **Bonjour** tab of the Add AP process, complete the configuration described in [Table 262](#).

Table 262: Bonjour Settings for an Access Point

<p>Enable Bonjour (disabled by default)</p>	<p>When you enable Bonjour, you must also provide a Location where Bonjour (mDNS) is located.</p>
<p>Task: Add services to this Bonjour configuration</p>	<p>To add a Bonjour service to access point:</p> <p>TIP: You can add multiple services.</p> <ol style="list-style-type: none"> Click Add under Service Name. The phrase <i>Enter service here</i> is displayed. Enter one of the following: <ul style="list-style-type: none"> Host Name Glob—Indicate a host for the Bonjour Profile. An asterisk like this * (the default) is a wildcard meaning all hosts. Service Name—Select one or more services for Bonjour, either Apple TV, Internet printer, or Digital Auto Access Protocol (iTunes): <ul style="list-style-type: none"> _airplay._tcp—Apple TV _ipp._tcp—Internet printer _daap._tcp—Digital Auto Access Protocol (iTunes) Service Type—Select either Discover (default) or Advertise. VLAN Scope—Select either Global or Local. <ul style="list-style-type: none"> Local—Service information is only relevant on the local VLAN. Global—Service information is relevant beyond the local VLAN. Click OK. The Bonjour service is added to the list of Service Names. Click OK.

Specifying Access Point Radio Settings

Access points can have one radio or two radios, depending on the model. Provide the settings for each radio, as described in [Table 263](#).

Table 263: Access Point Radio Settings

Field	Action
Radio Type	Select one of the available radio types on the access point: 11a, 11b, 11g, 11na, or 11ng. For a description of these radio types, see “Understanding the IEEE 802.11 Standard for Wireless Networks” on page 1076.
Radio Mode	Radio Mode is Enabled by default. Turn off the radio by selecting Disabled . Put the radio into sentry mode, which means the radio scans for interference but does not transmit user traffic, by selecting Sentry .
Antenna Type	If you did not add any external antenna, leave the default Internal selected to use the built-in antenna for this radio. If you added an antenna to the radio, select one of the supported add-on antenna models from the list.
Antenna Location	Leave the location set to Indoor if you are using an indoor antenna for this radio—this includes the built-in antenna. If you added an outdoor antenna, select Outdoor from the list.
Auto Channel	Auto channel is enabled by default, which means that the access point will switch channels depending on the optimum available choices. You can disable auto channel by removing the check mark and then assign a channel for the access point.
Channel Number	If Auto Channel is disabled, select a channel number for the radio. The channel numbers listed here are based on the country code. For more information about radio channels, see “Understanding Wireless Radio Channels” on page 856. NOTE: If auto-channel is enabled, the radio might change channels after it is deployed. For more information about radio channels, see “Understanding Adaptive Channel Planner” on page 861.
Auto Power	Auto Power is enabled by default, which means that the access point will modify its power depending on the circumstances. You can disable Auto Power by removing the check mark and then assign a transmit power for the access point.

Table 263: Access Point Radio Settings (continued)

Field	Action
Transmit Power	<p>If Auto Power is disabled, select the transmit power for the radio. The range available is 1 milliwatt (dBm) through 20 milliwatts (dBm). Transmit power is limited by some country codes. Unless you have a reason to do otherwise, we recommend that you set the power as high as possible.</p> <p>NOTE: If auto-power is enabled, the radio might change transmit power after it is deployed. For more information about automatic power tuning, see “Understanding Auto Tune Power Policy for Wireless Radios” on page 866.</p>
Max Transmit Power	<p>Each radio has a maximum transmit power in decibels, which is set to Default unless you change it. If the default option is chosen, the maximum power setting that RF Auto-Tuning can set on a radio is the highest setting allowed for the country of operation or highest setting supported on the hardware, whichever is lower. You can reset maximum transmit power by selecting a number from 1 through 20 decibels to indicate the maximum power setting RF Auto-Tuning can assign to the radio.</p>

Load Balancing

Load balancing distributes a workload across multiple entities, in this case wireless radios, to achieve optimal utilization, maximize throughput, minimize response time, and avoid overload. For more information, see [“Understanding Load Balancing for Wireless Radios”](#) on page 1070.

Enable Load Balancing	<p>Load balancing for radios is enabled by default, which means that the radio shares traffic equally with other radios in a load-balancing group. To turn off load-balancing on this radio, remove the check mark.</p> <hr/> <p>Load Balance Group Name: When load-balancing is enabled, you can create a load-balancing group and assign this radio to it by selecting this option and providing a group name. You can also add this radio to an existing load-balance group by selecting this option and providing an existing group name.</p> <hr/> <p>Enable Rebalance Clients: Optionally, when load-balancing is enabled, you can enable client rebalancing for this radio. This means that the access point radio can disassociate client sessions and rebalance them whenever a new access point radio is added to the same load-balancing group.</p>
------------------------------	--

If the access point has two radios, click the **Radio 2** tab to add the other radio's information. Then, click **OK**.

RELATED DOCUMENTATION

[Understanding the Network Director User Interface | 85](#)

[Understanding Auto AP Profiles | 883](#)

[Understanding Access Point Bias for Controllers | 852](#)

[Understanding Auto Tune Power Policy for Wireless Radios | 866](#)

[Understanding Wireless Radio Channels | 856](#)

[Understanding the IEEE 802.11 Standard for Wireless Networks | 1076](#)

[Network Director Documentation home page](#)

4

PART

Working in Deploy Mode

[About Deploy Mode | 1172](#)

[Deploying and Managing Device Configurations | 1180](#)

[Deploying and Managing Software Images | 1235](#)

[Managing Devices | 1248](#)

[Setting Up Zero Touch Provisioning for Devices | 1260](#)

About Deploy Mode

IN THIS CHAPTER

- [Understanding Deploy Mode in Network Director | 1172](#)
- [Understanding the Deploy Mode Tasks Pane | 1177](#)

Understanding Deploy Mode in Network Director

IN THIS SECTION

- [Deploying Configuration Changes | 1173](#)
- [Managing Software Images | 1174](#)
- [Zero Touch Provisioning | 1175](#)
- [Managing Devices | 1175](#)
- [Managing Device Configuration Files | 1175](#)
- [Managing Baseline Configuration | 1176](#)

The Deploy mode enables you to deploy configuration changes and software upgrades to devices and perform several device management and configuration file management tasks.

NOTE: Deploy mode is disabled for devices in your Datacenter View. This is because you can only discover, manage, and monitor devices in your virtual network. None of the deploy mode tasks are applicable to these devices.

This topic describes:

Deploying Configuration Changes

When you make configuration changes in Build mode, the changes are not deployed to devices automatically. You must manually deploy the changes to devices in Deploy mode. Every time you make configuration changes in Build mode that affect a device, the device is automatically added to the list of devices with pending changes.

NOTE: The device is added to the list of devices with pending changes only when you make the device configuration changes. If you make changes to the configurations (associated with the device) that are specific only to Network Director the device is not listed with pending changes. For example, when you make changes to the profile name associated with the device, the device is not added to the list of devices with pending changes.

Configuration changes are deployed to devices at the device level. When you deploy configuration changes to a device, all pending configuration changes for that device are deployed.

You can deploy the device configurations in the following two ways:

- **Auto Approval**—In this mode, the device configuration changes are approved automatically by the system and do not require explicit (manual) approval by a configuration approver before they can be deployed. This is the default approval mode.
- **Manual Approval**—In this mode, the device configuration changes are required to be explicitly approved by a configuration approver before the changes can be deployed to the device.

For more information about enabling these modes, see [“Setting Up User and System Preferences” on page 108](#).

For manual approval, the Network Director - Configuration Approver role is available in Junos Space, which is specific to the Network Director. A user with this role reviews device configurations and proposed changes to device configurations and can either approve or reject them.

An operator performs device configurations and creates a change request for that configuration and submits it for approval to an approver. The approvers are notified by e-mail whenever a change request is created. If a configuration or a change to it is approved by an approver, then the operator is able to deploy it. If a configuration is rejected then the operator must make the necessary changes, resubmit the change request, and procure an approval before the configuration can be deployed. For more information, see [“Approving Change Requests” on page 1206](#)

NOTE: You can specify any number of approvers. If you specify more than one approver while configuring the Manual Approval mode, once an approver accepts or rejects the proposed change, the change request is not listed for the other approvers and they cannot approve or reject the same change request.

You can do the following configuration deployment tasks on devices that have pending changes:

- Run configuration deployment jobs immediately or schedule them for future times.
- Approve the change requests for pending configurations, if you have selected the Manual Approval mode.
- Preview pending configuration changes before deploying the changes.
- Validate that the pending changes are compatible with the device's configuration.
- Manage configuration deployment jobs.

Configuration changes are validated for each device both in Network Director and on the device. If any part of a configuration change for a device fails validation, no configuration changes are deployed to the device. You can see the results of each validation phase separately.

Network Director will not deploy configuration to a device with a configuration that is out of sync (meaning that the device's configuration differs from Network Director's version of that device's configuration), or to a device that has uncommitted changes to its candidate configuration. Deployment to such devices will fail.

When you schedule a deployment job, that job and any profiles and devices assigned to that job are locked within Network Director. You cannot edit the job or any of its assigned profiles until the job runs or gets cancelled. This locking feature prevents you from deploying unintended configuration changes that could result from editing profiles and devices that are already scheduled to deploy. To change any properties of a scheduled job, cancel the job and create a new scheduled job with the desired properties. You cannot edit the profile assignments of a device that has scheduled pending configuration changes.

Managing Software Images

Network Director can manage software images on the nodes it manages. You can do the following software image management tasks:

- Deploy a software image stored in an image repository on the Network Director server to multiple devices with a single job.
- Track the status of software image management jobs.
- Stage and install software images as separate tasks.

- Schedule staging and installation to happen at independent future times.
- Perform several software image upgrade options, such as rebooting devices automatically after the upgrade finishes.

NOTE: Using nonstop software upgrade (NSSU) to upgrade EX Series switches is supported in Network Director.

Zero Touch Provisioning

Zero touch provisioning enables you to provision new Juniper Networks switches in your network automatically—without manual intervention. When you physically connect a switch to a network and boot it with the factory-default configuration, the switch attempts to upgrade the Junos OS software automatically and autoinstall a configuration file from the network.

Managing Devices

In Deploy mode you can perform several device management tasks, including:

- View the device inventory.
- Show a device's current configuration.
- Resynchronize the device configuration maintained in Build mode with the configuration on the device. For more information about resynchronization of device configuration, see [“Understanding Resynchronization of Device Configuration” on page 1214](#)
- Convert access points that were added to a controller using an Auto AP profile configuration to a persistent access point configuration on the controller.
- Enable or disable switch network ports.
- Manage QFabric node groups.
- Convert QSFP+ port configuration.

Managing Device Configuration Files

You can back up device configuration files to the Network Director server. You can perform several actions on backed up configuration files, such as restoring configuration files to devices, and viewing and comparing configuration files.

Managing Baseline Configuration

You can create a baseline of the Network Director device configuration and the OS version on the Network Director server. You can perform several actions on the baseline configuration files, such as restoring configuration files to devices, and viewing and comparing configuration files.

RELATED DOCUMENTATION

[Deploying Configuration to Devices | 1180](#)

[Deploying Software Images | 1238](#)

[Viewing the Device Inventory Page | 1136](#)

[Viewing a Device's Current Configuration from Network Director | 1150](#)

[Resynchronizing Device Configuration | 1220](#)

[Enabling or Disabling Network Ports on Switches | 1250](#)

[Converting Automatically Discovered Access Points to Manually Configured Access Points | 1248](#)

[Managing Device Configuration Files | 1225](#)

[Network Director Documentation home page](#)

Understanding the Deploy Mode Tasks Pane

The Tasks pane in Deploy mode lists the available tasks. All Deploy mode tasks are always available, regardless of the scope selected in the View pane.

Deploy mode tasks are divided into the following categories:

- **Configuration Deployment**—These tasks enable you to deploy configuration changes to devices and manage configuration deployment jobs. [Table 264](#) describes the configuration deployment tasks.
- **Image Management**—These tasks enable you to manage software images on devices. [Table 265](#) describes the image management tasks.
- **Device Management**—These tasks enable you to view the device inventory, resynchronize the configuration of out-of-sync devices, manage the administrative state of ports, manage QFabric node groups, and convert QSFP+ port configuration. [Table 266](#) describes the device management tasks.
- **Device Configuration File Management**—These tasks enable you manage configuration files on managed devices. [Table 267](#) describes the device configuration file management tasks.
- **Baseline Management**—These tasks enable you to manage baseline configuration of devices. [Table 268](#) describes the baseline management tasks.
- **Zero Touch Provisioning**—These tasks enable you to provision new Juniper Networks switches in your network automatically—without manual intervention. [Table 269](#) describes the zero touch provisioning tasks.
- **Key Tasks**—Network Director enables you to group the tasks that you perform frequently and create a list of key tasks. You can add any task from the Tasks pane to the Key Tasks list by selecting a task and clicking the plus (+) sign that appears adjacent to the task. For some modes, you can see that Network Director has predefined some key tasks for you. You can modify this set of tasks to suit your requirements. This feature is available in Task pane irrespective of your current mode, scope, or view.

[Table 264](#) through [Table 267](#) describe the tasks in each task category.

Table 264: Configuration Deployment Tasks

Task	Description
Deploy Configuration Changes	Deploys pending configuration changes to devices.
Approve Change Requests	Enables a configuration approver to approve or reject a change request, which has been submitted for approval by an operator.
Set SNMP Trap Configuration	Enables SNMP traps on network devices so that Network Director can collect and manage event and error information from these devices.
View Deployment Jobs	Manages configuration deployment jobs.

Table 265: Image Management Tasks

Task	Description
Manage Image Repository	Manages the software images repository on the server.
Deploy Images to Devices	Deploys software images from the repository to devices.
View Image Deployment Jobs	Manages software image deployment jobs.

Table 266: Device Management Tasks

Task	Description
Convert Auto AP	Converts access points that were added to a controller using an Auto AP profile configuration to a persistent access point configuration on the controller.
Convert Ports	Converts QSFP+ port configuration.
Manage Node Groups	Manages QFabric node groups.
Manage Port Admin State	Enables or disables switch network ports.
Resynchronize Device Configuration	Resynchronizes the device configuration maintained in Build mode with the running configuration on the devices.
Show Current Configuration	Shows the selected device's current configuration.
View Inventory	Displays the device inventory of the selected node.

Table 267: Device Configuration File Management Tasks

Task	Description
Manage Device Configuration Files	Manages backup device configuration files.
View Configuration File Mgmt Jobs	Manages device configuration file management jobs.

Table 268: Baseline Management Tasks

Task	Description
Manage Baseline	Manages baseline configuration files.
View Baseline Mgmt Jobs	Manages baseline configuration file management jobs.

Table 269: Zero Touch Provisioning Tasks

Task	Description
Setup	Set up the zero touch provisioning profile to configure the DHCP server and to upload the software image and configuration file to a file server.
Monitor	View details of the devices that are provisioned using a given zero touch provisioning profile.

RELATED DOCUMENTATION

[Understanding Deploy Mode in Network Director | 1172](#)

[Understanding the Network Director User Interface | 85](#)

[Network Director Documentation home page](#)

Deploying and Managing Device Configurations

IN THIS CHAPTER

- [Deploying Configuration to Devices | 1180](#)
- [Managing Configuration Deployment Jobs | 1194](#)
- [Deploy Configuration Window | 1196](#)
- [Importing Configuration Data from Junos OS Configuration Groups | 1198](#)
- [Enabling High-Frequency Traffic Statistics Monitoring on Devices | 1202](#)
- [Configuring Network Traffic Analysis | 1204](#)
- [Approving Change Requests | 1206](#)
- [Enabling SNMP Categories and Setting Trap Destinations | 1208](#)
- [Understanding Resynchronization of Device Configuration | 1214](#)
- [Resynchronizing Device Configuration | 1220](#)
- [Managing Device Configuration Files | 1225](#)
- [Creating and Managing Baseline of Device Configuration Files | 1230](#)

Deploying Configuration to Devices

IN THIS SECTION

- [Selecting Configuration Deployment Options | 1181](#)
- [Using the Change Request Details Page | 1185](#)
- [Creating a Change Request | 1186](#)
- [Validating Configuration | 1186](#)
- [Discarding the Pending Configurations | 1187](#)
- [Viewing Pending Configuration Changes | 1187](#)
- [Using the Pending Changes Window | 1187](#)
- [Using the Configuration or Pending Configuration Window | 1188](#)
- [Using the Deploy Configuration Errors/Warnings Window | 1189](#)

- Using the Configuration Validation Window | 1189
- Deploying Configuration Changes to Devices Immediately | 1189
- Scheduling Configuration Deployment | 1190
- Specifying Configuration Deployment Scheduling Options | 1190
- Editing Change Requests | 1191
- Deleting Change Request | 1192
- Resubmitting a Change Request | 1192
- Performing a Rollback | 1193

When you make configuration changes in Build mode, the changes are not deployed to devices automatically. You must manually deploy the changes to devices in Deploy mode.

To start deploying configuration changes:

1. Click **Deploy** in the Network Director banner.
2. Select a node in the View pane that contains the devices to which you want to deploy.
3. In the Tasks pane, select **Configuration Deployment > Deploy Configuration Changes**.

Depending upon the type of approval mode you select different windows are displayed.

If you select the Auto Approval mode, the Devices with Pending Changes page opens in the main window, listing the devices within the selected node that have pending configuration changes.

If you select the Manual Approval mode, the following two sections open in the main window:

- Devices with recent configuration changes—This section lists the devices with pending changes (along with the details of the change) performed by the user currently logged into the system.
- Change Requests—This section lists the change requests created by the user currently logged into the system.

This topic describes:

Selecting Configuration Deployment Options

Based on the approval mode, you can choose to deploy the device configuration changes in the following ways:

- When you select the auto approval mode, the page Devices with Pending Changes open. From the Devices with Pending Changes page, you can:

- Deploy configuration changes immediately by selecting one or more devices and clicking Deploy Now. For more information, see [“Deploying Configuration Changes to Devices Immediately” on page 1189](#).
- Schedule configuration deployment by selecting one or more devices and clicking Schedule Deploy. For more information, see [“Scheduling Configuration Deployment” on page 1190](#).
- View configuration changes that are pending on a device by clicking View in the Configuration Changes column. For more information, see [“Viewing Pending Configuration Changes” on page 1187](#).
- Validate that the pending changes for a device are compatible with the device’s configuration by selecting up to ten devices and clicking Validate Pending Configuration Changes. For more information, see [“Validating Configuration” on page 1186](#).
- Discard the pending configuration changes. For more information, see [“Discarding the Pending Configurations” on page 1187](#).

[Table 270](#) describes the information provided in the table on the Devices with Pending Changes page. Only the subset of devices within the selected object that have pending configuration changes are listed in the table.

Table 270: Devices with Pending Changes Page

Table Column	Description
Check box	Select to perform an action on the device in that row
Name	Device name
IP Address	Device IP address
Model	Device Model
OS Version	Operating system version running on device
Connection State	State of the connection to the device: <ul style="list-style-type: none"> • Up—Network Director can communicate with the device. • Down—Network Director cannot communicate with the device. You cannot deploy configuration to devices that are down.

Table 270: Devices with Pending Changes Page (continued)

Table Column	Description
Configuration State	<p>Indicates whether the device's configuration is in sync with Network Director's version:</p> <ul style="list-style-type: none"> • In Sync—The configuration on the device is in sync with the Network Director configuration for the device. • Out Of Sync—The configuration on the device does not match the Network Director configuration for the device. This state is usually the result of the device configuration being altered outside of Network Director. <p>You cannot deploy configuration on a device when the device is Out Of Sync. To resolve this state, use the Resynchronize Device Configuration task in Deploy mode.</p> <ul style="list-style-type: none"> • Synchronizing—The device configuration is in the process of being resynchronized. • Sync failed—An attempt to resynchronize an Out Of Sync device failed.
Configuration Changes	Click to view pending configuration changes for a device. The Pending Changes window opens.

If you select the Manual Approval mode, the windows Devices with recent configuration changes and Change Requests opens.

From the Devices with recent configuration changes window, you can:

- Create a device configuration change request approval and submit it for approval. Upon submission, all device changes made by an operator are validated and all the approvers are notified of the details of the proposed change request by e-mail. For more information, see [“Creating a Change Request” on page 1186](#).
- View configuration changes that are pending on a device by clicking View in the Configuration Changes column. For more information, see [“Viewing Pending Configuration Changes” on page 1187](#).
- Validate that the pending changes for a device are compatible with the device's configuration . For more information, see [“Validating Configuration” on page 1186](#).
- Discard the pending configuration changes. For more information, see [“Discarding the Pending Configurations” on page 1187](#).

NOTE: You cannot delete a device from the Devices with Pending Changes list. To remove a device from the list, you must undo the Build mode configuration changes that placed the device on the list.

[Table 271](#) describes the information provided in the table on the Devices with recent configuration changes page.

Table 271: Devices with recent configuration changes

Table Column	Description
Name	Indicates the name of the device and profile node. Below each device node, a profile node is listed.
Change Type	Indicates the type of the configuration change done to the device.
Associations Added	Lists the ports that are added to that profile.
Associations Deleted	Lists the ports that are deleted from that profile.
Configuration	Click to view pending configuration changes for a device. The Pending Changes window opens.
Deployment State	Indicates the deployment state of a change request.

From the Change Requests window, you can:

- Deploy configuration changes immediately by selecting one or more devices and clicking Deploy Now. For more information, see [“Deploying Configuration Changes to Devices Immediately” on page 1189](#).
- Schedule configuration deployment by selecting one or more devices and clicking Schedule Deploy. For more information, see [“Scheduling Configuration Deployment” on page 1190](#).
- Resubmit for the change request for approval after making the necessary modifications. For more information, see [“Resubmitting a Change Request” on page 1192](#).
- Edit or delete the change requests by selecting one or more change requests and clicking Edit or Delete respectively. For more information, see [“Editing Change Requests” on page 1191](#) and [“Deleting Change Request” on page 1192](#).
- Roll back the device configuration that is already deployed. For more information, see [“Performing a Rollback” on page 1193](#).
- View the details of the change request created. For more information, see [“Using the Change Request Details Page” on page 1185](#)

[Table 272](#) describes the information provided in the table on the change requests submitted for the devices for which configuration changes are sought.

Table 272: Change Requests

Table Column	Description
Check Box	Select to perform an action on the device in that row.

Table 272: Change Requests (continued)

Table Column	Description
Change Request No	Indicates the change request number of the change request that is waiting to be deployed.
Title	Indicates the title name of the change request.
Created On	Indicates the change request creation date.
Approver	Indicates the username of the configuration approver.
Last Action On	Indicates the date on which the change request status is changed.
Approval Status	Indicates whether a change request is approved or rejected by the approver.
Deployment Status	Indicates whether a change request is deployed after the approval.
History Icon	Records the audit trail details of a change request, such as operation performed on a change request during a given period of time, username of the approver or operator, and so on.

Using the Change Request Details Page

Use the Change Request Details window to view the details of the change request before you either approve or reject a change request. This window provides you the details such as change request number, title, username of the user who created the change request, change request creation date and so on. A Devices table is also displayed showing the deployment status. [Table 273](#) describes the fields in this table.

Table 273: Change Request Details

Column	Description
Name	Indicates the name of the device and profile node. Below each device node, a profile node is listed.
Change Type	Indicates the type of the configuration change done to the device.
Associations Added	Lists the ports that are added to that profile.
Associations Deleted	Lists the ports that are deleted from that profile.
Configuration	Click to view pending configuration changes for a device. The Pending Changes window opens.

Table 273: Change Request Details (continued)

Deployment Status	Indicates the deployment state of a change request.
-------------------	---

Creating a Change Request

To create a change request for device configurations approval:

1. Click **Create Change Request** in the Devices with recent configuration changes page.

The Create Change Request page opens.

2. Enter the change request number.

You can either enter a number or retain the autogenerated number in this field.

3. Enter an appropriate title name for the change request.

4. Optionally, you can enter comments for the device configuration changes.

5. Click **Submit**.

The Create Change Request page opens, listing the change request details such as change request number, title, and comment, along with the change request submission job details. A Devices table is also displayed showing the validation status of the device and configuration generated for that device.

6. Click **Close**.

A new change request entry with the status Pending Approval is added to the Change Request section.

Validating Configuration

When you deploy configuration changes to a device, validation checks are performed to validate that the pending changes are compatible with the device. You can also perform this validation without deploying.

NOTE: You can also verify the configuration from the Build mode by clicking **Tasks > Domain Management > Validate Pending Configuration**.

To validate that the pending changes for devices are compatible with the device configuration:

1. For Auto Approval mode, select up to ten devices in the Devices with Pending Changes page.

NOTE: For Manual Approval mode, you cannot choose the devices for which validation needs to be done. All the configuration changes for all the devices are validated.

2. Click **Validate Pending Configuration Changes**.

The Configuration Validation window opens. See [“Using the Configuration Validation Window” on page 1189](#) for a description of the window.

Discarding the Pending Configurations

Use the Discard Local Configuration Changes Results window to discard all the pending configurations that were made on a device. Once you discard the local configuration changes on a device, the configuration state of the device changes to In Sync or Out of Sync based on the system of record (SOR) mode set for the Junos Space Network Management Platform. If the SOR mode is set to Network as system of record (NSOR), then the configuration state changes to In Sync and if the SOR mode is set to Junos Space as system of record (SSOR), then the configuration state changes to Out of Sync.

To discard the configuration changes:

1. For Auto Approval mode, select the devices for which you want to discard the pending configuration and click **Discard Pending Configuration**.

The Discard Local Configuration Changes Results window opens displaying the status of the discard pending configuration job.

2. Click **Close** to close the Discard Local Configuration Changes Results window.

Viewing Pending Configuration Changes

To view pending configuration changes for a device, click **View** in the Pending Changes column.

The Pending Changes window opens. See [“Using the Pending Changes Window” on page 1187](#) for a description of the window.

Using the Pending Changes Window

Use the Pending Changes window to view the pending Network Director changes for a device. [Table 274](#) describes the fields in this window.

Table 274: Pending Changes Window

Field	Description
Name	Lists each selected device. Expand a device by clicking its plus sign to see its pending changes. Each pending change to a profile or other configuration object for the device is listed.
State	Describes the nature of the pending change to the configuration object. These are the possible states: <ul style="list-style-type: none"> • Added—The profile or configuration object was added to this device. • Removed—The profile or configuration object was removed from the device • Updated—The profile or configuration object was updated.
Configuration	Click View to view the pending configuration changes for a device. The Pending Configuration window opens. See “Using the Configuration or Pending Configuration Window” on page 1188 for information about the window. NOTE: The device configuration state must be In Sync for you to view the pending configuration changes.
Close	Click to close the window.

Using the Configuration or Pending Configuration Window

Use the Pending Configuration window to view the configuration changes that will be deployed to a device when a job runs. Use the Configuration window to see changes that were deployed to a device when a completed job ran. The configuration changes are shown in these formats:

- Select the **XML View** tab to view the configuration changes in XML format. This view shows the XML-formatted configuration that will be deployed to the device’s Device Management Interface (DMI), which is used to remotely manage devices.
- Select the **CLI View** tab to view the configuration changes in CLI format. This view shows the Junos configuration statements that will be deployed to the device.

In both views, the content is color-coded for easier reading:

- Black text indicates configuration that is already active on the device, and will not be changed if you deploy.
- Green text indicates configuration that will be added if you deploy.
- Red text indicates configuration that will be removed if you deploy.

Using the Deploy Configuration Errors/Warnings Window

Use the Deploy Configuration Errors/Warnings window to view the results of deploying configuration to a device. The Errors/Warnings in validating the device configuration pane shows the results of configuration validation by Network Director. The Errors/Warnings in Updating Device configuration pane shows the results of configuration validation on the device.

Using the Configuration Validation Window

Use the Configuration Validation window to validate that the pending changes for a device are compatible with the device's configuration. [Table 275](#) describes this window.

Table 275: Configuration Validation Window

Table Column	Description
Object name	Lists the devices you selected for validation. Click the arrow next to a device to expand it. If there are no errors or warnings, one item labeled No Validation warnings appears. If the device has errors or warnings, they appear under the device. The device contains a list of the profiles that caused errors or warnings. Expand a profile name to see the of errors and warnings it caused.
Errors/Warnings	Describes the error or warning.

Deploying Configuration Changes to Devices Immediately

To deploy configuration changes to devices immediately:

1. Select the device or devices in the Devices with Pending Changes page.

2. Click **Deploy Now**.

The Deploy Options window opens.

3. In the Deploy Options window, enter a job name in the Deployment Job Name field, then click **OK**.

The configuration deployment job runs. The Deploy Configuration window opens and shows the results of the deployment job. For a description of fields in this window, see [“Deploy Configuration Window” on page 1196](#).

Scheduling Configuration Deployment

To schedule configuration deployment to devices:

1. Select the device or devices in the Devices with Pending Changes page.

2. Click **Schedule Deploy**.

The Deploy Options window opens.

3. Use the Deploy Options window to schedule the configuration deployment. See [“Specifying Configuration Deployment Scheduling Options”](#) on page 1190 for a description of the window.

Specifying Configuration Deployment Scheduling Options

Use the Deploy Options window to schedule configuration deployment jobs. [Table 276](#) describes the actions for the fields in this window.

Table 276: Deploy Options Window

Field	Action
Deployment Job Name	Enter a job name.
Date and Time	Enter the job’s start date and time.
OK	Click to accept changes and exit the window.
Cancel	Click to cancel changes and exit the window.

Editing Change Requests

You can edit a change request to change the profile that was added to a device or delete some of the profile associations. After editing a change request, you can resubmit the change request for approval. While editing a change request, if you try to delete all the profile associations in a given change request, the system prompts a message that a change request should have at least one valid association. Deleting all the associations in a change request makes it invalid. Hence, you cannot delete all the associations in a given change request. However, you can delete a change request itself to delete all the associations for that change request.

NOTE: You are unable to delete a change request or an association of a change request if an association is in pending removal state.

You are unable to edit a change request that is in Cancelled, Deployed, Rollback Success, or Rollback Failed state.

To change a profile or delete the profile associations of a change request:

1. Select the change request in the Change Requests pending action page to edit.

2. Click **Edit**.

The Edit Change Request window opens.

3. Click the call out symbol to change the profile and choose the new profile that you want to assign the change request.

4. To delete a profile association, click **Delete**

5. Click **Save**.

The Edit Change Request window opens, listing the change request details such as change request number, title, and comment, along with the change request submission job details. A Devices table is also displayed showing the validation status of the device and configuration generated for that device.

6. Click **Close**.

Deleting Change Request

Sometimes you might need to delete a change request from the change request list. A change request is assigned with profile associations. If you delete a change request, all the associations of that change request are also deleted.

To delete a change request:

1. Select the change request or change requests in the Change Requests pending action window.
2. Click **Delete**.

The Delete Change Request window opens, displaying the message: **Are you sure you want to delete Change Request?**

3. Click **Yes** to delete the change request; else click **No**.

If you clicked **Yes**, the message: **Change Request deleted successfully** appears.

4. Click **OK**.

Resubmitting a Change Request

You can resubmit only those change requests that are in Pending Approval, Pending Deployment, Deploy Failed, and Create Failed state. You are unable to resubmit change requests in Deployed, Cancelled, Rollback Success, or Rollback Failure state.

In certain situations, a device can go out of sync while a user is creating a change request for that device. The change request is created, but the configuration changes for that change request are not generated. You can select the change request and resubmit it after the device is in sync again, which generates the configuration for this change request. You can resubmit change requests only for devices that have pending configuration changes.

To resubmit a change request:

1. Select the change request in the Change Requests pending action window to edit.
2. Click **Resubmit**.

A warning message pops up indicating if you want to resubmit the change request.

3. Click **Yes**.

The Resubmit Change Request window opens, listing the change request details such as change request number, title, and comment, along with the change request submission job details. A Devices table is also displayed showing the validation status of the device and configuration generated for that device.

4. Click **Close**.

Performing a Rollback

In case of any misconfigurations, you can choose to roll back a configuration that has already been deployed to the device. The following conditions apply for a rollback operation:

- The maximum number of change requests that you can roll back is the rollback limit specified in Preferences.
- Change requests are rolled back in reverse chronological order; the later change requests are rolled back first. If there are any conflicting change requests, roll back is not supported. For example, assume that a user assigns port profile P1 to ge-0/0/1 and creates a change request CR1 and deploys the profile. After this, if the user edits P1, creates another change request CR2 and deploys and removes P1 from the port by assigning some other port profile and deploys device changes or configurations as part of CR3. If the user now tries to roll back CR1, an error message about the conflicting change requests CR2 and CR3 is shown. To roll back CR1, the user must roll back CR3, then CR2, and then CR1.

To roll back a device:

1. Select the device in the left navigation pane for which you want to perform the roll back operation.
2. Select **Rollback Configuration Changes** task under Configuration Deployment.
3. All the devices with previously stored configuration of the device are listed in the working area of the right pane.

You can view the stored configuration also user can view the difference of the current device configuration and stored configuration.

4. Choose the configuration for which you want to perform the roll back.

NOTE: You can choose only one rollback configuration out of the available configurations per device however you can choose multiple devices.

5. Click **Rollback**.

A rollback job is started with all the selected devices and all the devices are resynchronized after the configuration is pushed.

RELATED DOCUMENTATION

- [Deploying Configuration Changes | 1173](#)

- [Managing Configuration Deployment Jobs | 1194](#)

- [Approving Change Requests | 1206](#)

- [Setting Up User and System Preferences | 108](#)

- [Network Director Documentation home page](#)

Managing Configuration Deployment Jobs

IN THIS SECTION

- [Selecting Configuration Deployment Job Options | 1194](#)
- [Viewing Configuration Deployment Job Details | 1195](#)
- [Canceling Configuration Deployment Jobs | 1196](#)

When you deploy configuration changes or schedule a configuration deployment, a configuration deployment job is created.

To start managing configuration deployment jobs:

1. Click **Deploy** in the Network Director banner.
2. In the Tasks pane, select **Device Management > View Deployment Jobs**.

The Deploy Configuration page opens in the main window. The table on that page lists configuration deployment jobs.

This topic describes:

Selecting Configuration Deployment Job Options

From the Deploy Configuration page, you can:

- View the details of a configuration deployment job by clicking Show Details. See [“Viewing Configuration Deployment Job Details” on page 1195](#) for more information.

- Cancel a scheduled configuration deployment job by clicking Cancel Job. See [“Canceling Configuration Deployment Jobs”](#) on page 1196 for more information.

Table 277 describes the information provided on the Deploy Configuration page

Table 277: Deploy Configuration Table Description

Table Column	Description
Check box	Select to perform an action on the job in that row.
Job Id	An identifier assigned to the job.
Job Name	Job name (user-created).
Percent	Percentage of the job that is complete.
Status	Job status. The possible states are: <ul style="list-style-type: none"> • CANCELLED—The job was cancelled by a user. • FAILURE—The job failed. This state is applied if any of the devices in the job failed. But some of the devices might have completed successfully. View the job details for the status of each device. • INPROGRESS—The job is running. • SCHEDULED—The job is scheduled but has not run yet. • SUCCESS—The job completed successfully. This state is applied if all of the devices in the job completed successfully.
Summary	Job summary.
Scheduled Start Time	Job’s scheduled start time
Actual Start Time	Time when the job started.
End Time	Time when the job ended
User	User who created the job
Recurrence	This field is not used for configuration deployment jobs.

Viewing Configuration Deployment Job Details

To view the details of a configuration deployment job:

1. Select the job in the table.

2. Click **Show Details**. The Deploy Configuration window opens. See [“Deploy Configuration Window” on page 1196](#) for a description of the window.

Canceling Configuration Deployment Jobs

To cancel a configuration deployment job:

1. Select the job in the table.
2. Click **Cancel Job**.
3. Click **Yes** in the confirmation window that opens.

RELATED DOCUMENTATION

[Deploying Configuration Changes | 1173](#)

[Deploying Configuration to Devices | 1180](#)

[Network Director Documentation home page](#)

Deploy Configuration Window

The Deploy Configuration window shows the results of a completed deployment job or information about a scheduled job. See [Table 278](#) for a description of the fields in this window.

Table 278: Deploy Configuration Window

Field	Description
Job Name	Job name.
Job Start Time	Time when job started or will start.
Job End Time	Time when job ended.
Percentage Completed	Percentage of the job that is complete.
Number of Devices	Number of devices in the deployment job.

Deployed Devices table

Table 278: Deploy Configuration Window (continued)

Field	Description
Name	Device name.
IP Address	Device IP address.
Deployment Status	<p>Status of configuration deployment on device:</p> <ul style="list-style-type: none"> • Scheduled—Job is scheduled for future deployment. • In Progress—Deployment is in progress. • Success—Deployment completed successfully. • Failed—Deployment failed.
Configuration	<p>Click View to see the configuration changes that were deployed to the device. See “Using the Configuration or Pending Configuration Window” on page 1188 for more information.</p> <p>For a scheduled job, this column does not contain a link. See “Deploying Configuration to Devices” on page 1180 for information about viewing pending configuration changes for a device.</p>
Result Details	<p>Click View to see the results of configuration deployment for the device. See “Using the Deploy Configuration Errors/Warnings Window” on page 1189 for more information.</p>
Close	Click to close the window.

RELATED DOCUMENTATION

[Deploying Configuration Changes | 1173](#)

[Deploying Configuration to Devices | 1180](#)

[Managing Configuration Deployment Jobs | 1194](#)

[Network Director Documentation home page](#)

Importing Configuration Data from Junos OS Configuration Groups

IN THIS SECTION

- [Enabling Import of Configuration Group Data | 1198](#)
- [Viewing Configuration Group Data | 1199](#)
- [Using the Configuration or Pending Configuration Window | 1200](#)
- [Deploying Configuration Group Changes to Devices Immediately | 1201](#)
- [Scheduling Configuration Group Change Deployment | 1201](#)
- [Specifying Configuration Deployment Scheduling Options | 1201](#)
- [Using the Deploy Configuration Errors/Warnings Window | 1202](#)

The configuration groups feature in Junos OS enables you to create a group containing configuration statements and to direct the inheritance of that group's statements in the rest of the configuration on a device. The same group can be applied to different sections of the configuration, and different sections of one group's configuration statements can be inherited in different places in the configuration.

Configuration groups enable you to create smaller, more logically constructed configuration files, making it easier to configure and maintain Junos OS. For example, you can group statements that are repeated in many places in the configuration, such as when configuring interfaces, and thereby limit updates to just the group.


You can use the Migrate Config Groups task in Network Director to import and deploy supported configurations from these configuration groups on devices.

This topic describes:

Enabling Import of Configuration Group Data

For Network Director to be able to import configuration group data.

To enable the import of configuration group data:

1. Click  in the Network Director banner and select **Preferences**.
2. In the Preferences window, select the **Config & Deploy** tab.

3. Select the **Enable migration from Ethernet Design** check box to enable import of configuration group data.
4. Click **Save** to save and close the preferences.

Viewing Configuration Group Data

After you enable the import of configuration group data, Network Director adds a new task—Migrate from Ethernet Design—in the Deploy mode. You can use this task to flatten the active group configurations. All the devices which are discovered and managed by Network Director is listed in the Migrate from Ethernet Design page. Network Director displays the devices in this page based on you selection in the Tree view. You can view the default configuration, active configuration, and the configurations that are to be deployed on the devices from this page.

To view the configuration group data:

1. Click **Deploy** in the Network Director banner.
2. In the Tasks pane, select **Configuration Deployment > Migrate From Ethernet Design**. The Migrate Ethernet Design Configuration Groups page opens displaying all the devices that has configurations from configuration groups that are not deployed on the device yet.

[Table 279](#) describes the information provided in the table on the Migrate Config Groups page.

NOTE: This table also lists the wireless controllers in your network. Ignore these devices as you cannot view or deploy changes to controllers using this task.

Table 279: Migrate Config Groups Page

Table Column	Description
Check box	Select to perform an action on the device in that row
Name	Device name
IP Address	Device IP address
Model	Device Model
Device Family	Device family can be junos, junos-ex, or junos-qfx
OS Version	Operating system version running on device

Table 279: Migrate Config Groups Page (continued)

Table Column	Description
Connection State	<p>State of the connection to the device:</p> <ul style="list-style-type: none"> • Up—Network Director can communicate with the device. • Down—Network Director cannot communicate with the device. You cannot deploy configuration to devices that are down.
Configuration State	<p>Indicates whether the device's configuration is in sync with Network Director's version:</p> <ul style="list-style-type: none"> • In Sync—The configuration on the device is in sync with the Network Director configuration for the device. • Out Of Sync—The configuration on the device does not match the Network Director configuration for the device. This state is usually the result of the device configuration being altered outside of Network Director. You cannot deploy configuration on a device when the device is Out Of Sync. To resolve this state, use the Resynchronize Device Configuration task in Deploy mode. • Synchronizing—The device configuration is in the process of being resynchronized. • Sync failed—An attempt to resynchronize an Out Of Sync device failed.
Configuration Changes	<p>Click to view the configuration changes for a device. The Pending Changes window opens. For more details on using the Pending Changes window, see “Using the Configuration or Pending Configuration Window” on page 1200.</p>

Using the Configuration or Pending Configuration Window

Use the Pending Configuration window to view the configuration changes that will be deployed to a device when a job runs. Use the Configuration window to see changes that were deployed to a device when a completed job ran. The configuration changes are shown in these formats:

- Select the **XML View** tab to view the configuration changes in XML format. This view shows the XML-formatted configuration that will be deployed to the device's Device Management Interface (DMI), which is used to remotely manage devices.
- Select the **CLI View** tab to view the configuration changes in CLI format. This view shows the Junos configuration statements that will be deployed to the device.

In both views, the content is color-coded for easier reading:

- Black text indicates configuration that is already active on the device, and will not be changed if you deploy.
- Green text indicates configuration that will be added if you deploy.
- Red text indicates configuration that will be removed if you deploy.

Deploying Configuration Group Changes to Devices Immediately

To deploy configuration group changes to devices immediately:

1. Select the device or devices in the Devices with Pending Changes page.

2. Click **Deploy Now**.

The Deploy Options window opens.

3. In the Deploy Options window, enter a job name in the Deployment Job Name field, then click **OK**.

The configuration deployment job runs. The Deploy Configuration window opens and shows the results of the deployment job. For a description of fields in this window, see [“Deploy Configuration Window” on page 1196](#).

Scheduling Configuration Group Change Deployment

To schedule configuration group change deployment to devices:

1. Select the device or devices in the Devices with Pending Changes page.

2. Click **Schedule Deploy**.

The Deploy Options window opens.

3. Use the Deploy Options window to schedule the configuration deployment. See [“Specifying Configuration Deployment Scheduling Options” on page 1201](#) for a description of the window.

Specifying Configuration Deployment Scheduling Options

Use the Deploy Options window to schedule configuration deployment jobs. [Table 280](#) describes the actions for the fields in this window.

Table 280: Deploy Options Window

Field	Action
Deployment Job Name	Enter a job name.
Date and Time	Enter the job's start date and time.
OK	Click to accept changes and exit the window.
Cancel	Click to cancel changes and exit the window.

Using the Deploy Configuration Errors/Warnings Window

Use the Deploy Configuration Errors/Warnings window to view the results of deploying configuration to a device. The Errors/Warnings in validating the device configuration pane shows the results of configuration validation by Network Director. The Errors/Warnings in Updating Device configuration pane shows the results of configuration validation on the device.

RELATED DOCUMENTATION

[Deploying Configuration Changes | 1173](#)

[Managing Configuration Deployment Jobs | 1194](#)

[Network Director Documentation home page](#)

Enabling High-Frequency Traffic Statistics Monitoring on Devices

To use Network Director monitoring analytics features such as latency heat maps and congestion monitoring, you must enable high-frequency traffic statistics monitoring on the QFX devices to be monitored. You can also configure the high-frequency traffic statistics sampling rate and event thresholds on devices.

NOTE: High-frequency traffic statistics monitoring requires Cloud Analytics Engine support. For more information, see [“Understanding Cloud Analytics Engine and Network Director” on page 83](#).

NOTE: You must specify the IP address of DLE server under **Preferences > Monitoring > Data Learning Engine Settings** before you can enable high-frequency traffic statistics monitoring as described in this topic.

To enable high-frequency traffic statistics monitoring on devices:

1. Log in to Network Director.
2. Under Views, select **Logical View**, **Location View**, **Device View**, or **Custom Group View**.
3. Click **Deploy** in the Network Director banner to open Deploy mode.
4. Select the task **Configuration Deployment > Enable High Frequency Stats** in the Tasks pane.

The Enable High Frequency Stats page opens. It contains a table listing the QFX devices that support high-frequency statistics monitoring.

5. To enable or disable high-frequency traffic statistics monitoring on a device, use the check box in the Enable column.
6. In the Device / Port column, specify whether you want to enable high-frequency traffic statistics monitoring on a device as a whole or on selected ports.

If you select **Device**, high-frequency traffic statistics monitoring is enabled on all ports on the device with the settings you specify. If you select **Port**, you can selectively enable/disable high-frequency traffic statistics monitoring on individual ports on the device and specify different settings for each port.
7. If you have selected **Port** in the previous step, click the arrow next to the device name to expose the list of ports and to enable high-frequency traffic statistics monitoring on selected ports.
8. To change high-frequency traffic statistics monitoring settings, double-click a current setting in the table to edit it. [Table 281](#) describes these settings.
9. To deploy the settings currently configured on the page, click **Deploy**.
10. To reset all settings on the page to the default values and deploy those settings, click **Restore all values to default and Deploy**.

Table 281: High-Frequency Traffic Statistics Monitoring Settings

Setting	Description
Traffic Stats Sampling Interval (seconds)	Sets the interval for traffic statistics sampling, in seconds.
Latency Stats Sampling Interval (milli seconds)	Sets the interval for latency statistics sampling, in milliseconds.
Latency Threshold (nano seconds)	Sets the latency threshold, in nanoseconds. Monitored latency values higher than this threshold are considered congestion events.

RELATED DOCUMENTATION

[Understanding Cloud Analytics Engine and Network Director | 83.](#)

[Understanding Deploy Mode in Network Director | 1172](#)

[Understanding Monitor Mode in Network Director | 1269](#)

Configuring Network Traffic Analysis

Network Traffic Analysis (NTA) is a monitoring technology for high-speed switched or routed networks. Once enabled, Network Director randomly samples network packets and sends the samples to a data learning engine (DLE) for analysis. Network traffic analysis uses packet-based sampling. Network Director samples one packet out of a specified number of packets from an interface enabled for network traffic analysis and sends the packet to the DLE. DLE uses this sampling information to create a picture of the network traffic, which includes the applications that contribute to the traffic, traffic statistics, and the top applications.

You can enable network traffic analysis on all devices, except the wireless devices, that are managed by Network Director.

Before you configure network traffic analysis, ensure that:

- You have configured one or more data learning engines to analyze the network traffic.
- You have specified the IP address and port number of these DLEs in the System Preferences window. For detailed steps, see the [Specifying the Data Learning Engine \(DLE\) Settings in the “Setting Up User and System Preferences” on page 108](#) topic.

NOTE: On devices for which you want to enable NTA, sFlow must not be configured. Before you enable NTA, make sure that sFlow is not enabled on the devices to be monitored. To verify that the sFlow is not enabled, log in to the device CLI and execute the following command and verify the output:

```
[root@user ~]# show protocols sflow | display set
```

No output is displayed when sFlow is not configured on the device.

NOTE: Network Director supports the sFlow on physical port only and not in aggregation links or logical ports.

To configure network traffic analysis:

1. Log in to Network Director.
2. Under Views, select **Logical View**, **Location View**, **Device View**, or **Custom Group View**.
3. Click **Deploy** in the Network Director banner.

4. In the Tasks pane, select **Configuration Deployment > Enable Network Traffic Analysis**.

The Enable Network Traffic Analysis page opens.

5. Select the check box adjacent to **Enable Traffic Analysis on Devices when Port Utilization exceeds certain percentage** to enable network traffic analysis. The default port utilization percentage value is 90%. You can change the default value to a value that is appropriate for your network.

6. Specify the number of packets from which a packet must be sampled in the **Sample rate** field.

Network Director samples one packet out of a specified number of packets from an interface enabled for network traffic analysis and sends the packet to DLE.

For example, if you configure a sampling rate of 10, one packet is sampled from every 10 packets.

7. Click **Add Devices** to add new devices for network traffic analysis.

The Add Devices window opens.

8. In the Add Devices window, select the devices for which you want to enable network traffic analysis.

9. Click **Add**.

Network Director adds the selected devices to the list in the Enable Network Traffic Analysis page.

To remove a device, select a device from the list and click **Remove**.

10. Click **Save** to save the network traffic analysis configuration details.

Network Director initiates traffic analysis when traffic utilization on any interface of the devices added to the Enable Network Traffic Analysis page exceeds the port utilization that you specified. You can view the traffic analysis details from the **Monitor mode > Traffic Analysis** or the Device & Port Utilization dashboard widget.

RELATED DOCUMENTATION

[Device & Port Utilization Widget | 151](#)

[Monitoring Port Traffic Statistics | 1284](#)

Approving Change Requests

NOTE: This option is available only for the users who are assigned a Configuration Approver role.

When you select the Approve Change Request option, the page Change request(s) pending approval and the page approved/declined change request(s) open in the top and bottom panels respectively.

The [Table 282](#) shows details of the change requests that are pending for approval by the approver.

Table 282: Change request(s) pending approval

Table Column	Description
Change Request No	Indicates the change request number that was either approved or rejected by the approver.
Title	Indicates the title of the change request.
Created By	Indicates the operator name who created the change request and submitted it for approval.
Created On	Indicates the date on which the change request was created.
Age	Indicates the age of the change request, time since the change request was created.
Deployment Status	Indicates the deployment state of the change request.
History Icon	Records the audit trail details of a change request, such as operation performed on a change request during a given period of time, username of the operator, and so on.

The [Table 283](#) shows the change requests that were approved or rejected by the currently logged in approver. The approver can also provide comments

Table 283: approved/declined change request(s)

Table Column	Description
Change Request No	Indicates the change request number that was either approved or rejected by the approver.
Title	Indicates the title of the change request.

Table 283: approved/declined change request(s) (continued)

Table Column	Description
Created By	Indicates the operator name who created the change request and submitted it for approval.
Created On	Indicates the date on which the change request was created.
Age	Indicates the age of the change request, time since the change request was created.
Approval Status	Indicates the approval state of the change request.
Deployment Status	Indicates the deployment state of the change request.
History Icon	Records the audit trail details of a change request, such as operation performed on a change request during a given period of time, username of the approver, and so on.

To approve or reject the change requests submitted by an operator:

1. Select **Approve Change Requests** under Configuration Deployment.
2. Select the check box against the change request and click on a change request in the change request(s) pending approval page.
The Change Request Details page opens.
3. Review details of the profile and its associations.
4. Click on the **View** link.
The Pending Configuration device name page opens.
5. Click **Close**.
6. Click **Approve** or **Reject** to approve or reject the device configuration changes respectively.
The Change Request Details page opens.
7. Type your comments and click **Approve** to approve; else click **Reject**.
After the successful approval, you can deploy the device configurations immediately or schedule the deployment for a later period.

RELATED DOCUMENTATION

[Setting Up User and System Preferences | 108](#)

[Deploying Configuration to Devices | 1180](#)

[Network Director Documentation home page](#)

Enabling SNMP Categories and Setting Trap Destinations

IN THIS SECTION

- [Viewing Eligible Devices for Trap Forwarding | 1208](#)
- [Enabling Trap Forwarding | 1209](#)
- [Deploying SNMP Trap Configurations | 1210](#)

SNMP traps must be enabled on network devices for Network Director to collect and manage event and error information from these devices.

Network Director organizes switch and controller traps by categories. These categories must be enabled and deployed in order to forward trap information to Network Director.

NOTE: Network Director uses protocol port 10162 for receiving traps from devices. This port must be open on the devices.

NOTE: Network Director only supports SNMP V1 and V2C traps.

This topic describes:

Viewing Eligible Devices for Trap Forwarding

Traps are enabled on the Devices page in Deploy mode. To locate this page:

1. Select **Deploy** in the Network Director banner.

2. Select **Set SNMP Trap Configuration** in the Tasks pane. The Devices page opens. For a description of fields in the Devices page, view [Table 284](#).

Table 284: Device Page Fields

Field	Description
Name	Either the hostname or the IP address of the device.
IP Address	Device IP address.
Model	Device model number.
OS Version	Version and release level of the operating system running on the device.
Connection State	State of connection to the device. Valid states are: <ul style="list-style-type: none"> • Up—Network Director is in communication with the device. • Down—Network Director cannot communicate with the device. You cannot enable traps on devices that are in this state.
Configuration State	Either the device's configuration is in sync or out-of-sync with Network Director's version: <ul style="list-style-type: none"> • IN_SYNC—The configuration is in-sync with the database. • OUT_OF_SYNC—The configuration is out-of-sync with the database.

Enabling Trap Forwarding

Select **Set SNMP Trap Configuration** in Deploy mode to enable your network devices to pass SNMP traps and events to Network Director. Network Director creates a target group called *networkdirector_trap_group* using target port 10162. The Community name is *public* and the access is *read-write-notify*.

Before enabling trap forwarding, complete device discovery for all the devices and ensure they are in the up state. Down devices cannot be enabled for trap forwarding.

Selecting Set SNMP Trap Configuration displays the Devices page which contains a table of all discovered switches and controllers in the network. To enable SNMP traps on switches and controllers:

1. Either select individual check boxes for devices, or select the check box next to the Name heading to select all devices. These devices must be up and in the same device family. So if you have both wireless devices and switches, you need to deploy the trap configurations in separate passes.
2. Click **Deploy Trap Configuration**. The Deploy Options window opens.

3. Fill in a new deployment job name or accept the default name of Deploy SNMP Targets.
4. Either select check boxes for individual traps, or select the check box next to the Trap Name heading to select all traps. These traps are discussed further in [“Deploying SNMP Trap Configurations” on page 1210](#).

TIP: To clear an existing configuration, do not select any of the check boxes.

5. Click **Ok**. The Deploy Configuration window opens, which shows the status of deploying the configuration change.
6. Review the outcome of the deployment.

After enabling the traps, enable the alarms and establish the alarm retention period. These tasks are located in Preferences in the Network Director banner.

Deploying SNMP Trap Configurations

The Deploy Options for trap forwarding enable you to select individual traps or all traps for the selected device family.

The device family determines which traps are displayed in the Deploy Options window. The following tables map the trap to one or more MIBs being used.

- EX Series switches traps and related MIBs are shown in [Table 285](#).
- Controllers traps and related MIBs are shown in [Table 286](#).

Table 285: EX Series Switches Traps

Trap	MIB
Chassis	jnxExMibRoot.mib
Link	snmpTraps.mib
Configuration	jnxCfgMgmt.mib
Authentication	jnxJsAuth.mib
Remote operations	jnxPing.mib
Routing	jnx-ipv6.mib

Table 285: EX Series Switches Traps (continued)

Trap	MIB
Startup	snmpTraps.mib
Rmon-alarm	jnxRmon.mib
Vrrp-events	rfc2787a.mib
Services	jnxServices.mib
Sonet-alarms	jnx-sonetaps.mib
Otn-alarms	jnxMlbs.mib
PoE-alarms	mib-rfc3621.mib

Table 286: Controllers Traps

Trap	MIB
LinkDown	snmpTraps.mib
LlinkUp	snmpTraps.mib
Authentication	snmpTraps.mib
DeviceFail	trpzTrapsV2.mib
DeviceOkay	trpzTrapsV2.mib
PoEFail	trpzTrapsV2.mib
MobilityDomainJoin	trpzTrapsV2.mib
MobilityDomainTimeout	trpzTrapsV2.mib
RFDetectAdhocUser	trpzTrapsV2.mib
ClientAuthenticationFailure	trpzTrapsV2.mib
ClientAuthorizationFailure	trpzTrapsV2.mib
ClientAssociationFailure	trpzTrapsV2.mib

Table 286: Controllers Traps (continued)

Trap	MIB
ClientDeAssociation	trpzTrapsV2.mib
ClientRoaming	trpzTrapsV2.mib
AutoTuneRadioPowerChange	trpzTrapsV2.mib
AutoTuneRadioChannelChange	trpzTrapsV2.mib
CounterMeasureStart	trpzTrapsV2.mib
CounterMeasureStop	trpzTrapsV2.mib
ClientDot1xFailure	trpzTrapsV2.mib
RFDetectDoS	trpzTrapsV2.mib
RFDetectDoSPort	trpzTrapsV2.mib
ClientIpAddrChange	trpzTrapsV2.mib
ClientAssociationSuccess	trpzTrapsV2.mib
ClientAuthenticationSuccess	trpzTrapsV2.mib
ClientDeAuthentication	trpzTrapsV2.mib
RFDetectBlacklisted	trpzTrapsV2.mib
RFDetectAdhocUserDisappear	trpzTrapsV2.mib
ApRejectLicenseExceeded	trpzTrapsV2.mib
ClientDynAuthorChangeSuccess	trpzTrapsV2.mib
ClientDynAuthorChangeFailure	trpzTrapsV2.mib
ClientDisconnect	trpzTrapsV2.mib
MobilityDomainFailOver	trpzTrapsV2.mib
MobilityDomainFailBack	trpzTrapsV2.mib

Table 286: Controllers Traps (continued)

Trap	MIB
RFDetectRogueDeviceDisappear	trpzTrapsV2.mib
RFDetectSuspectDeviceDisappear	trpzTrapsV2.mib
RFDetectedClientViaRogueWiredAP	trpzTrapsV2.mib
RFDetectedClassificationChange	trpzTrapsV2.mib
ConfigurationSaved	trpzTrapsV2.mib
APNonOperStatus	trpzTrapsV2.mib
MichaelMICFailure	trpzTrapsV2.mib
ApManagerChange	trpzTrapsV2.mib
ClientCleared	trpzTrapsV2.mib
MobilityDomainResiliencyStatus	trpzTrapsV2.mib
ApOperRadioStatus	trpzTrapsV2.mib
ClientAuthorizationSuccess	trpzTrapsV2.mib
RFDetectRogueDevice	trpzTrapsV2.mib
RFDetectSuspectDevice	trpzTrapsV2.mib
ClusterFailure	trpzTrapsV2.mib
MultimediaCallFailure	trpzTrapsV2.mib
ApTunnelLimitExceeded	trpzTrapsV2.mib
WsTunnelLimitExceeded	trpzTrapsV2.mib
RFNoiseSource	trpzTrapsV2.mib
M2UConvNotPossibleTrap	trpzTrapsV2.mib
M2UConvAvailabilityRestored	trpzTrapsV2.mib

RELATED DOCUMENTATION

[Setting Up User and System Preferences | 108](#)

[Understanding Fault Mode in Network Director | 1448](#)

[Network Director Documentation home page](#)

Understanding Resynchronization of Device Configuration

IN THIS SECTION

- [The Resynchronize Device Configuration Task | 1215](#)
- [How Resynchronization Works in NSOR Mode | 1216](#)
- [How Resynchronization Works in SSOR Mode | 1218](#)
- [How Network Director Resynchronizes the Build Mode Configuration | 1220](#)

In a network managed by Network Director, three separate repositories about device configuration are maintained:

- The configuration information on the devices themselves. Each switch and wireless LAN controller maintains its own configuration record.
- The configuration information maintained by the Junos Space Network Management Platform. When a device is discovered, either by Junos Space or Network Director, Junos Space stores a record of the configuration on that device.

Network Director uses the configuration record maintained by Junos Space to determine what configuration commands need to be sent to the device when you deploy configuration on the device in Deploy mode.

- The configuration information maintained by Network Director in Build mode. This information takes the form of the profiles assigned to the device, plus the additional configuration, such as LAG and access point configuration, that you can do under device management.

In Network Director, the configuration state of a device is shown as In Sync when the configuration information in all three repositories match. If there is a conflict between the configuration information in one or more of the repositories, Network Director shows the device configuration state as Out of Sync.

An Out of Sync state is usually the result of out-of-band configuration changes—that is, configuration changes made to a device using a management tool other than Network Director. Examples of such changes include changes made by:

- Using the device CLI.
- Using the device Web-based management interface (the J-Web interface or Web View).
- Using the Junos Space Network Management Platform configuration editor.

NOTE: You cannot use the Junos Space configuration editor to configure wireless LAN controllers.

- Using RingMaster software.
- Restoring or replacing device configuration files.

You cannot deploy configuration on a device when the device configuration state is Out of Sync.

This topic describes how Network Director enables you to resynchronize the device configuration state. It covers:

The Resynchronize Device Configuration Task

Network Director provides a task in Deploy mode that enables you to resynchronize the repositories of

configuration information. When an out-of-band configuration change is made, you can use this task to resynchronize both the Junos Space configuration record and the Build mode configuration with the configuration on the device.

How Network Director performs resynchronization depends on the system of record (SOR) mode set for the Junos Space Network Management Platform. There are two possible modes:

- Network as system of record (NSOR). This is the default mode.
- Junos Space as system of record (SSOR).

You set the mode in Junos Space under Administration > Applications > Network Management Platform > Modify Application Settings.

How Resynchronization Works in NSOR Mode

In NSOR mode, the network device is considered the system of record for device configuration, which means the configuration maintained by the device takes precedence over the configuration maintained by Junos Space and Network Director. Thus when you perform a resynchronization, the Junos Space configuration record and the Network Director Build mode configuration are updated to match the device configuration.

When an out-of-band change is made on a managed device when Junos Space is in NSOR mode:

1. Junos Space detects that a configuration change has occurred on the device and informs Network Director about the change.
2. Both Junos Space and Network Director set the device configuration state to Out of Sync.
3. Junos Space and Network Director automatically resynchronizes its configuration record to match the device configuration and set the device configuration state to In Sync when the synchronization completes. Network Director performs auto-synchronization when it is operating in the Network as System Of Record (NSOR) mode. The auto-resynchronization parameters are defined in the Preferences page. These parameters enables auto-resynchronization after the interval specified in the Preferences page. For more information see, [“Setting Up User and System Preferences” on page 108](#).

NOTE: Auto-synchronization is not supported for wireless devices.

4. When the device out-of-band changes does not conflict with Network Director, Network Director automatically resynchronizes the network changes and retains the local changes in Network Directory. The configuration state of the device and the profile associated with that device remain unaffected. For example, if you modify the MTU value of the port ge-0/0/1 in Network Director and another user modifies the MTU value of port ge-0/0/2 on the same device, Network Director automatically

resynchronizes the changes on ge-0/0/2 into Network Director and retains the local changes on ge-0/0/1. The profile corresponding to ge-0/0/1 continues to remain in Pending Deployment state and the profile corresponding to port ge-0/0/2 is in Deployed state.

5. When the device out-of-band changes conflict with the changes made in Network Director, Network Director does not automatically resynchronize the device changes into Network Director. The device is marked as Conflict. You must manually resynchronize the changes by using the Resynchronize Configuration task. After this, the local changes are discarded and are replaced by the latest network configuration. For example, if you modify MTU of ge-0/0/1 from Network Director and another user modifies MTU of the same port on the device, Network Director does not automatically synchronize and marks this device as Out Of Sync.
6. When a profile associated with a device is either added or removed from that device while another user tries to change the attributes corresponding to that profile, Network Director does not automatically synchronize the device and marks the device as conflict, and you must manually resynchronize the changes by using the Resynchronize Configuration task.
7. When you make local changes to profiles, the changes are merged with the new profiles if there is no conflicting configuration. If there are conflicting changes, Network Director receives an Out Of Sync message from Junos Space and you need to manually choose the appropriate profile value.

When you do not make any local changes on a profile, the device association with the profile is deleted and a new device association is created. However, when a profile has local changes, the device association of the profile is not deleted.

NOTE: Automatic resynchronization, as described in Step 3 above, is a default setting for the Junos Space Network Management Platform. If automatic resynchronization is disabled, you must manually resynchronize the Junos Space configuration with the device configuration. You can do so in two ways:

- Use the Resynchronize with Network action in Junos Space. The Junos Space configuration is synchronized with the device configuration. However, the Build mode configuration is not synchronized, so the device state in Network Director remains Out of Sync. You must use the Resynchronize Device Configuration task in Deploy mode to resynchronize the Build mode configuration.
- Use the Resynchronize Device Configuration task in Deploy mode. In this case, Network Director resynchronizes both the Junos Space configuration and the Build mode configuration with the device configuration.

How Resynchronization Works in SSOR Mode

When Junos Space is in SSOR mode, Junos Space is considered the system of record for device configuration. In this mode, when an out-of-band configuration change occurs on a device, you can choose whether to accept the change or to overwrite the change with the configuration maintained by Junos Space.

When an out-of-band change is made on a managed device when Junos Space is in SSOR mode:

1. Junos Space detects that a configuration change has occurred on the device and informs Network Director about the change.
2. Junos Space sets the device configuration state as Device Changed, and Network Director sets the device configuration state to Out of Sync.

Network Director sets the device configuration state to Out of Sync even if the configuration change does not affect configuration you can perform in Build mode. This allows you to resolve the Device Changed configuration state for Junos Space from Network Director.

3. In Network Director, use the Resynchronize Device Configuration task to accept or reject the out-of-band changes:
 - If you accept the out-of-band changes, both the Junos Space configuration record and the Network Director Build mode configuration are resynchronized to reflect the out-of-band configuration changes.
 - If you reject the out-of-band changes, the configuration on the device is overwritten by the configuration record maintained by Junos Space. The Network Director Build mode configuration remains unchanged.
4. Both Junos Space and Network Director set the device configuration state to In Sync.

The above process differs somewhat when out-of-band configuration changes are made through the Junos Space configuration editor. In this case:

1. Junos Space sets the device configuration state as Space Changed after the configuration change is saved.

At this point, the changes have been made only in the Junos Space configuration record and the changes have not yet been deployed to the device. Network Director shows the device configuration state as In Sync.

NOTE: Because the device configuration state is In Sync in Network Director, you can deploy configuration on the device from Network Director at this point. If you do so, the Network Director changes are deployed on the device, but the Junos Space changes are not. The device state in Junos Space remains Space Changed.

2. When the changes are deployed to the device from Junos Space, Junos Space changes the device state to In Sync, while Network Director changes the device state to Out of Sync.
3. In Network Director, use the Resynchronize Device Configuration task to resolve the Out of Sync state. In this case, because the Junos Space configuration record and the device configuration are in sync, you cannot reject the changes. When you resynchronize the device in Network Director, the Build mode configuration is updated to reflect the configuration changes.
4. Network Director sets the device configuration state to In Sync.

If you use Junos Space instead of Network Director to resolve out-of-band configuration changes in SSOR mode, note the following:

- If you reject an out-of-band change, the device state becomes In Sync in both Network Director and Junos Space.
- If you accept an out-of-band change that does not affect the Build mode configuration, the device state becomes In Sync in both Network Director and Junos Space.
- If you accept an out-of-band change that affects the Build mode configuration, the device state becomes In Sync in Junos Space but remains Out Of Sync in Network Director. You must use the Resynchronize Device Configuration task to resolve the Out of Sync state.

NOTE: When Junos Space is in SSOR mode, we recommend that you do not make out-of-band changes to the cluster configuration on the secondary seeds and member controllers of a mobility domain, such as disabling the cluster on these devices. Use Network Director to modify the cluster configuration on these devices.

How Network Director Resynchronizes the Build Mode Configuration

When you use the Resynchronize Device Configuration task to resynchronize the Build mode configuration to the device configuration, Network Director launches a resynchronization job. This job deletes all profile assignments configured for the device. The profiles themselves are not deleted—just the assignments of the profiles to the device are deleted. It then reimports the device configuration, as if the device were a newly discovered device. It reassigns existing profiles and creates new profiles as necessary. Profiles that were originally assigned to the device will be reassigned to the device if the profiles were unaffected by the out-of-band changes. All profiles assigned to the device are in a deployed state at the end of the process. Any profile that is not reassigned to the device and is not assigned to any other device will be in a unassigned state.

RELATED DOCUMENTATION

[Resynchronizing Device Configuration | 1220](#)

[Network Director Documentation home page](#)

Resynchronizing Device Configuration

IN THIS SECTION

- [The Resynchronize Device Configuration List of Devices | 1221](#)
- [Resynchronizing Devices When Junos Space Is in NSOR Mode | 1222](#)
- [Resynchronizing Devices When Junos Space Is in SSOR Mode | 1223](#)
- [Resynchronizing Devices in Manual Approval Mode | 1224](#)
- [Viewing the Network Changes | 1224](#)
- [Viewing Resynchronization Job Status | 1225](#)

A network managed by Network Director has three repositories of information about the configuration of a network device—the configuration stored on the device itself, the device configuration record maintained by Junos Space, and the Build mode configuration maintained by Network Director.

When the configuration contained in all three repositories match, the device configuration state is shown as In Sync in Network Director. When the repositories do not match, the configuration state is shown as

Out of Sync. A common cause for this state is out-of-band configuration changes—that is, configuration changes made to a device outside of Network Director.

When a device state is Out of Sync, you cannot deploy configuration changes on the device in Deploy mode. Use the Resynchronize Device Configuration task to resynchronize the three configuration repositories and change the device configuration state back to In Sync.

How the Resynchronize Device Configuration task performs the resynchronization depends on the system of record (SOR) mode setting for the Junos Space Network Management Platform:

- When Junos Space is in network as system of record (NSOR) mode, the device is considered the system of record for configuration. When you resynchronize a device when Junos Space is in NSOR mode, both the Junos Space configuration record and the Network Director Build mode configuration are updated to reflect the device configuration—in other words, the out-of-band configuration changes are incorporated into both the Junos Space and the Network Director configuration repositories.
- When Junos Space is in Junos Space as system of record (SSOR) mode, you can choose whether accept or reject the out-of-band changes reflected in the device configuration. If you accept the changes, both the Junos Space configuration record and the Network Director Build mode configuration are updated to reflect the device configuration. If you reject the changes, the out-of-band changes are rolled back on the device so that the device configuration matches the Junos Space configuration record and the Network Director Build mode configuration.

For more information about out-of-band configuration changes, Junos Space SOR modes, and how Network Director resynchronizes device configuration, see [“Understanding Resynchronization of Device Configuration” on page 1214](#).

This topic covers:

The Resynchronize Device Configuration List of Devices

The Resynchronize Device Configuration page displays a list of all devices in the selected scope whose configuration was successfully imported during device discovery and whose configuration state is now Out Of Sync. You can select devices from this list and resynchronize them.

[Table 287](#) describes the fields in the list of devices.

Table 287: Resynchronize Device Configuration Fields

Field	Description
Name	Device hostname or device IP address.
IP address	IP address of device.
Model	Model number of the device.

Table 287: Resynchronize Device Configuration Fields (continued)

Field	Description
OS Version	Operating system version currently running on the device.
Connection State	<p>Connection state:</p> <ul style="list-style-type: none"> • UP—Network Director is connected to the device • DOWN—Network Director cannot connect to the device
Configuration State	<p>Shows the configuration state of the device:</p> <ul style="list-style-type: none"> • Out Of Sync—The device configuration is out of sync with either the Network Director Build mode configuration or the Junos Space configuration record or both. • Resynchronizing—The device configuration is in the process of being resynchronized. • Sync Failed—The resynchronization attempt failed. <p>If the resynchronization is successful, the device is removed from the table.</p>
Local Changes	<p>Specifies whether configuration changes have been made in Build mode and are pending deployment on the device.</p> <ul style="list-style-type: none"> • None—There are no configuration changes pending deployment. • View—There are configuration changes that are pending deployment. Click View to view the changes. These changes will be lost if you resynchronize the Build mode configuration to match the device configuration. <p>NOTE: The Pending Changes window that appears when you click View allows you to see what profiles have been added, modified, or changed. However, because the device is not in sync, you cannot view the specific changes in CLI or XML format.</p>
Network Changes	<p>Indicates whether you can view the out-of-band changes:</p> <ul style="list-style-type: none"> • None—The out-of-band changes are not available for viewing. You cannot view out-of-band changes in NSOR mode. In SSOR mode, you cannot view the out-of-band changes if they are already resolved in Junos Space—that is, the device configuration state in Junos Space is In Sync. • View—You can view the out-of-band changes made on the device. Click View to view the changes presented in XML format.

Resynchronizing Devices When Junos Space Is in NSOR Mode

To resynchronize devices when the Junos Space Network Application Platform is in NSOR mode:

1. On the Resynchronization Device Configuration page, select the device or devices that you want to resynchronize.

2. (Optional) View any pending changes to a device's configuration in Network Director by clicking **View** in the Local Changes column. These pending changes are deleted when you resynchronize the device.

3. Click **Resynchronize Configuration**.

The Resynchronize Device Configuration Results window appears. This window will be updated with status of the resynchronization when the resynchronization completes.

Resynchronizing Devices When Junos Space Is in SSOR Mode

To resynchronize devices when the Junos Space Network Management Platform is in SSOR mode:

1. On the Resynchronization Device Configuration page, select the device or devices that you want to resynchronize.
2. (Optional) View any pending changes to a device's configuration in Network Director by clicking **View** in the Local Changes column. These pending changes are deleted if you accept the out-of-band changes when you resynchronize the device.
3. (Optional) View the out-of-band configuration changes by selecting **View** in the Network Changes column. If you accept the out-of-band changes when you resynchronize the device, these changes will be reflected in the Build mode configuration. If you reject the out-of-band changes when you resynchronize the devices, these changes will be deleted from the device. For more information about viewing the out-of-band changes, see ["Viewing the Network Changes" on page 1224](#).

NOTE: Out-of-band changes that were made with the Junos Space configuration editor or that were already accepted in Junos Space are not shown. Such changes also cannot be rejected.

4. Click **Resynchronize Configuration**.

5. In the Confirm dialog box:

- Click **Accept device changes** if you want to accept the out-of-band changes.
- Click **Reject device changes** if you want to reject the out-of-band changes and have the configuration that existed on the device before the out-of-band changes were made be reinstated.

click **Submit**.

The Resynchronize Device Configuration Results window appears. This window will be updated with status of the resynchronization when the resynchronization completes.

NOTE: Device changes made by the Junos Space configuration editor or device changes that have been accepted in Junos Space cannot be rejected. Even if you select Reject device changes, these changes will not be rejected and instead will be incorporated into the Build mode configuration.

Resynchronizing Devices in Manual Approval Mode

When out-of-band changes exist, device resynchronization merges the changes done by using the CLI with the local changes provided that there are no conflicts. If there are conflicting changes, the changes made using the CLI take precedence over the local changes. Therefore, configuration changes that are part of a change request might be lost. The configuration change requests that are lost are marked as Cancelled against the corresponding device. When device resynchronization is initiated for a device, a message is displayed that lists the change requests that will be lost because of conflicting CLI and local changes. All other changes remain unaffected.

Viewing the Network Changes

The Network Changes window shows the out-of-band configuration changes made to a device when Junos Space is in SSOR mode.

Not all out-of-band configuration changes are shown in this window. Configuration changes are shown only when the device configuration differs from the Junos Space configuration record—that is, when the device configuration state in Junos Space is not In Sync. For example, if the out-of-band changes were deployed from the Junos Space configuration editor or if the out-of-band changes were already accepted in Junos Space, the configuration changes will not appear in this window.

The configuration changes are shown in XML format. If there have been multiple out-of-band changes—that is, there has been more than one configuration commit, or save, on the device—the changes are grouped by each commit.

The following information is provided for each configuration commit:

- `junos:commit-seconds`—Specifies the time when the configuration was committed as the number of seconds since midnight on 1 January 1970.
- `junos:commit-localtime`—Specifies the time when the configuration was committed as the date and time in the device's local time zone.
- `xmlns:junos`—Specifies the URL for the DTD that defines the XML namespace for the tag elements.
- `junos:commit-user`—Specifies the username of the user who requested the commit operation.

Viewing Resynchronization Job Status

The Resynchronize Device Configuration Results window appears after you start a resynchronization job. This window is automatically updated with the resynchronization status for each device when the job completes.

You can also view the status of the resynchronization jobs using the Manage Jobs task in System mode. The following jobs are associated with resynchronization:

- Resynch Network Elements—This job runs in NSOR mode and resynchronizes the Junos Space configuration record with the device configuration.
- Resolve OOB Changes—This job runs in SSOR mode and resolves the out-of-band changes for Junos Space—either accepting the changes and updating the Junos Space configuration or rejecting the changes and rolling back the changes on the device.
- Resynchronize devices—This job runs in both NSOR and SSOR mode and resynchronizes the Build mode configuration with the device configuration.

RELATED DOCUMENTATION

[Understanding Resynchronization of Device Configuration | 1214](#)

[Managing Jobs | 105](#)

[Setting Up User and System Preferences | 108](#)

[Network Director Documentation home page](#)

Managing Device Configuration Files

You can back up device configuration files to the Network Director server. You can perform several actions on backed up configuration files, such as restoring configuration files to devices, and viewing and comparing configuration files.

To start managing device configuration files:

1. Click **Deploy** in the Network Director banner.
2. In the Tasks pane, select **Device Configuration Files > Manage Device Configuration Files**.

The Manage Device Configuration page opens in the main window. The table lists the devices that have configuration files backed up.

This topic describes:

- [Selecting Device Configuration File Management Options | 1226](#)
- [Backing Up Device Configuration Files | 1227](#)
- [Restoring Device Configuration Files | 1227](#)
- [Viewing Device Configuration Files | 1228](#)
- [Comparing Device Configuration Files | 1228](#)
- [Deleting Device Configuration Files | 1229](#)
- [Managing Device Configuration File Management Jobs | 1229](#)

Selecting Device Configuration File Management Options

From the Manage Device Configuration page, you can:

- Back up device configuration files by clicking Backup. See [“Backing Up Device Configuration Files” on page 1227](#) for more information.
- Restore backup device configuration files to devices by selecting devices and clicking Restore. See [“Restoring Device Configuration Files” on page 1227](#) for more information.
- View backed up configuration files by selecting a device and clicking View Configuration File. See [“Viewing Device Configuration Files” on page 1228](#) for more information.
- Compare backed up device configuration files by selecting devices and clicking Compare Config Files. See [“Comparing Device Configuration Files” on page 1228](#) for more information.
- Delete backup device configuration files by selecting devices and clicking Delete. See [“Deleting Device Configuration Files” on page 1229](#) for more information.

[Table 288](#) describes the information provided in the Manage Device Configuration table.

Table 288: Manage Device Configuration Table

Table Column	Description
Device Name	Device name.
Config File Version	Version number of the backup configuration file.
First Backup on	Date when the oldest version of the backup configuration file was created.
Most Recent Backup on	Date when the configuration file was backed up most recently.

Backing Up Device Configuration Files

To back up device configuration files:

1. Click **Backup**.

The Backup Devices Configuration page opens in the main window.

2. Select the devices to back up from the device tree.

3. To back up configuration files immediately, click **Backup Now**.

The backup job runs. When it finishes, the Manage Device Configuration table shows updated information for the devices you backed up.

4. To schedule the backup to run later, click **Schedule Backup**.

The Schedule Backup window opens.

- a. Select the **Schedule at a later time** check box.

- b. Specify when the backup will run using the **Date and Time** fields.

- c. Optionally, configure the backup job to repeat by selecting the **Repeat** check box, then specifying the backup schedule using the provided fields.

Optionally, you can specify when repeated backups will stop by selecting the **End Time** check box, then specifying the last date on which the repeated backup job will run using the **Date and Time** fields.

- d. Click **Schedule Backup**.

Restoring Device Configuration Files

You can restore a backed up configuration file to the device from which it was backed up.



CAUTION: Restoring a configuration file to a device is considered an out-of-band configuration change, which can cause some unexpected results. For more information, see [“Out-of-Band Configuration Changes” on page 188](#).

To restore backed up configuration files to devices:

1. Select the devices to restore from the Manage Device Configuration list.
2. Click **Restore**.

The Restore Device Configuration File(s) window opens.

3. To restore a configuration file that is older than the most recent version, click in the **Latest Version** cell and select the version to restore.
4. Click **Restore**.

Viewing Device Configuration Files

To view the backed up configuration files for a device:

1. Select the device from the Manage Device Configuration list.
2. Click **View Configuration File**.

The Device Configuration Summary window opens, displaying the most recently backed up configuration file.

3. To view an older stored configuration file version, select a version number from the **Config File Version** list.

Comparing Device Configuration Files

To compare backed up device configuration files:

1. Select the configuration files to compare from the Manage Device Configuration list.
2. Click **Compare Configuration Files**.

The Compare Configuration Files window opens.

3. Select a source device from the **Source Device** list and a configuration file version from the **Config File Version** list.

4. Select a target device from the **Target Device** list and a configuration file version from the **Config File Version** list.
5. The configuration file versions you selected are displayed in the window. The file name and version appears at the top of each file. The differences between the configuration files are color-coded. The color-coding legend appears at the top of the window.

Deleting Device Configuration Files

When you delete a device's backed up configuration, all of the configuration file versions for the device are deleted.

To delete device configuration files:

1. Select the configuration files to delete from the Manage Device Configuration list.
2. Click **Delete**.
The Delete Device Configuration File(s) window opens.
3. Verify that the correct devices are listed, then click **Delete**.

Managing Device Configuration File Management Jobs

Each time you back up or restore device configuration files, a device configuration file management job is created.

To manage device configuration file management jobs:

1. Click **Deploy** in the Network Director banner.
2. In the Tasks pane, select **Device Configuration Files > View Configuration File Mgmt Jobs**.
The Device Configuration Jobs page opens in the main window, listing the device configuration file management jobs.

Managing these jobs is similar to managing other types of jobs using the System mode. The advantage of accessing the jobs this way is that the jobs list show only configuration file management jobs. See [“Managing Jobs” on page 105](#) for more information.

RELATED DOCUMENTATION

[Understanding the Deploy Mode Tasks Pane | 1177](#)

[Understanding Deploy Mode in Network Director | 1172](#)

[Managing Jobs | 105](#)

[Network Director Documentation home page](#)

Creating and Managing Baseline of Device Configuration Files

IN THIS SECTION

- [Selecting Baseline Management Options | 1231](#)
- [Baselining Device Configuration Files | 1231](#)
- [Restoring Baseline Device Configuration Files | 1232](#)
- [Viewing Baseline Configuration Files | 1233](#)
- [Comparing Baseline Configuration with Current Configuration | 1233](#)
- [Deleting Baseline | 1233](#)
- [Managing Baseline Management Jobs | 1234](#)

You can create a baseline device configuration and the device Junos (OS) version on the Network Director server. By creating a baseline configuration file for a device you define a reference point to save the device configuration and its OS version to a particular known state and later restore the configuration to that known state. You can select the devices at the scope level, custom grouping, or for individual devices and create baseline configuration files and images for all or for the selected devices. The baseline configuration file includes the entire configuration and image files. When you restore a device configuration, you restore both the baseline configuration file and the image of the file. However, restoring image is optional.

An alarm is triggered if there are any changes to the baseline configuration. The alarm contains the delta information for later reference. For example, when you add a new device an alarm of with minor severity is generated to inform user about the device addition. When you move a device from an unassigned category to a specific category, an alarm of major severity is generated.

To start baseline file management:

1. Click **Deploy** in the Network Director banner.
2. In the Tasks pane, select **Baseline Management > Manage Baseline**.

The Manage Device Baseline page opens in the main window.

This topic describes:

Selecting Baseline Management Options

From the Manage Device Baseline page, you can:

- Create device baseline configuration files by clicking **Baseline**. See [“Baselining Device Configuration Files” on page 1231](#) for more information.
- Restore baselined device configuration files to devices by selecting devices and clicking **Restore**. See [“Restoring Baseline Device Configuration Files” on page 1232](#) for more information.
- View the baselined configuration files by selecting the device and clicking **View Configuration File**. See [“Viewing Baseline Configuration Files” on page 1233](#) for more information.
- Compare baseline device configuration files with current configuration files by selecting devices and clicking **Compare With Current Config**. See [“Comparing Baseline Configuration with Current Configuration” on page 1233](#) for more information.
- Delete baseline configuration files for devices by selecting devices and clicking **Delete**. See [“Deleting Baseline” on page 1233](#) for more information.

[Table 289](#) describes the information provided in the Manage Device Baseline table.

Table 289: Manage Device Baseline Table

Table Column	Description
Device Name	Name of baseline device.
Baseline Label	Name of the baseline configuration.
Baseline Update Time	Date and time when the baseline configuration file for a device is last updated.
Baseline State	Indicates whether the baseline configuration is same as or out of sync with the current configuration.

Baselining Device Configuration Files

To create baseline configuration for devices:

1. Click **Baseline**.

The Create Baseline for Devices page opens in the main window.

2. Type a baseline label name.

3. Select the devices for which you want to create a baseline configuration files from the device tree.
4. To back up configuration files immediately, click **Create Baseline Now**.

The device baseline job runs. When it finishes, the Manage Device Baseline table shows updated information for the devices for which you performed the baselining.

5. To schedule the backup to run later, click **Schedule Baseline**.

The Schedule Baseline window opens.

- a. Select the **Schedule at a later time** check box.
- b. Specify when the baseline will run using the **Date and Time** fields.
- c. Optionally, configure the baseline job to repeat by selecting the **Repeat** check box, and then specifying the backup schedule by using the fields provided.

Optionally, you can specify when repeated backups will stop by selecting the **End Time** check box, and then specifying the last date on which the repeated backup job will run, using the **Date and Time** fields.

- d. Click **Schedule Baseline**.

Restoring Baseline Device Configuration Files

To restore the baseline configuration file and the OS image of devices:

1. Select the devices from the Manage Device Baseline list.
2. Click **Restore**.

The Restore Baseline(s) window opens.

3. Click **Restore Image** to restore the OS image and select **Reboot device after successful installation** if you want to reboot the device.
4. To view the device configuration summary, click **Click to View Config File**.
5. Click **Restore**.

The device baseline job runs. When it finishes, the Manage Device Baseline table shows updated information for the devices for which you performed the restore.

Viewing Baseline Configuration Files

To view the baseline configuration files for a device:

1. Select the device from the Manage Device Baseline list.
2. Click **View**.

The Device Configuration Summary window opens, displaying the most recently baseline configuration file.

3. To compare the baseline configuration with the current configuration, click **Compare With Current Config**.

Comparing Baseline Configuration with Current Configuration

To compare the baseline configuration with the current configuration:

1. Select the configuration files to compare from the Manage Device Baseline list.
2. Click **Compare With Current Config**.

The Compare Baseline With Current Configuration window opens.

3. The baseline versions and the current version are displayed in the window. The device name and version appear at the top of each file. The differences between the configuration files are color-coded. The color-coding legend appears at the top of the window.
4. Click **Close**.

Deleting Baseline

When you delete a baseline configuration file, all its corresponding versions are also deleted.

To delete baseline configuration file for a device:

1. Select the device name from the Manage Device Baseline list.
2. Click **Delete**.

The Delete Baseline(s) window opens.

3. Verify that the correct devices are listed, then click **Delete** to delete the device configuration.

Managing Baseline Management Jobs

Each time you create baseline configuration files, a baseline management job is also created.

To manage device configuration file management jobs:

1. Click **Deploy** in the Network Director banner.
2. In the Tasks pane, select **Device Configuration Files > View Baseline Mgmt Jobs**.

The Device Configuration Jobs page opens in the main window, listing the device configuration file management jobs.

3. To view the details of a job, select the job name and click **Show Details**.

Managing these jobs is similar to managing other types of jobs using the System mode. The advantage of accessing the jobs this way is that the jobs list shows only configuration file management jobs. See [“Managing Jobs” on page 105](#) for more information about managing jobs.

RELATED DOCUMENTATION

[Understanding the Deploy Mode Tasks Pane | 1177](#)

[Understanding Deploy Mode in Network Director | 1172](#)

[Managing Jobs | 105](#)

Deploying and Managing Software Images

IN THIS CHAPTER

- [Managing Software Images | 1235](#)
- [Deploying Software Images | 1238](#)
- [Managing Software Image Deployment Jobs | 1243](#)

Managing Software Images

This topic describes how to manage software images for managed devices.

To start managing software images:

1. Click **Deploy** in the Network Director banner.
2. In the Tasks pane, select **Image Management > Manage Image Repository**.

The Device Image Repository page opens in the main window. The table lists the software images in the repository.

Starting with the Network Director 4.0R1 release, images uploaded on the Junos Space Platform from the **Images and Scripts > Images** page will also be available on the **Manage Image Repository** page.

NOTE:

- Only images uploaded on Junos Space Platform are available in Network Director. However, images uploaded in Network Director will NOT be available in Junos Space Platform.
- If you delete an image from Network Director, the image is NOT deleted in Junos Space Platform.

3. In the Tasks pane, select **Device Configuration File Management > Manage Device Configuration**.

The Manage Device Configuration page opens in the main window. The table lists the devices that have configuration files backed up software images in the repository.

This topic describes:

- [Selecting Software Image Management Options | 1236](#)
- [Adding Software Images to the Repository | 1236](#)
- [Using the Device Image Upload Window | 1237](#)
- [Viewing Software Image Details | 1237](#)
- [Using the Device Image Summary Window | 1237](#)
- [Deleting Software Images | 1238](#)

Selecting Software Image Management Options

From the Device Image Repository page, you can:

- Add a software image to the repository by clicking Add.
- View details about a software image by selecting it and clicking Details.
- Delete software images from the repository by selecting them and clicking Delete.

[Table 290](#) describes the information provided in the Device Image Repository table.

Table 290: Device Image Repository Table

Table Column	Description
Check box	Select to perform an action on the software image in that row.
Name	Software image name.
Version	Software version.
Series	Device series that uses the software image.
Uploaded By	User who uploaded the software image.
Created On	Time when the software image was uploaded to the server.
Size(MB)	Size of the software image in megabytes.

Adding Software Images to the Repository

Software images are stored in a repository on the Network Director server.

To add a software image to the repository:

1. Click **Add**.

The Device Image Upload window opens.

2. Use the Device Image Upload window to upload a device software image. See [“Using the Device Image Upload Window” on page 1237](#) for a description of the window.

Using the Device Image Upload Window

To use the Device Image Upload window to add a software image to the repository:

1. Click **Browse** and browse to the software image file.
2. Click **Upload** to add the file to the repository.

Viewing Software Image Details

To view details about a software image:

1. Select the software image file in the table.
2. Click **Details**.

The Device Image Summary window opens. See [“Using the Device Image Summary Window” on page 1237](#) for information about this window.

Using the Device Image Summary Window

Use the Device Image Summary window to view detailed information about a software image. [Table 291](#) describes the fields in this window.

Table 291: Device Image Summary Window

Field	Description
Name	Software image filename.
Version	Software version (release number).
Series	Device series on which the software is supported.
Supported Platforms	Platforms on which the software is supported.

Table 291: Device Image Summary Window (continued)

Field	Description
Uploaded By	User who uploaded the image to the server.
Created On	Date and time when the software image was uploaded.
Size (MB)	Size of the software image file, in megabytes.
OK	Click to close the window.

Deleting Software Images

To delete software image files:

1. Select the check box in the rows of the software image files that you want to delete.
2. Click **Delete**.

RELATED DOCUMENTATION

[Managing Software Images | 1174](#)

[Deploying Software Images | 1238](#)

[Managing Software Image Deployment Jobs | 1243](#)

[Network Director Documentation home page](#)

Deploying Software Images

This topic describes how to deploy software images to managed devices. You must upload software images to the Network Director server before you can deploy them to devices. See “[Managing Software Images](#)” on [page 1235](#) for more information.

To start deploying software images:

1. Click **Deploy** in the Network Director banner.
2. Select a node in the View pane that contains the devices on which you want to deploy software images.

3. In the Tasks pane, select **Image Management > Deploy Images to Devices**.

The Select Devices page of the Deploy Images to Devices wizard opens in the main window.

This topic describes:

- [Specifying Software Deployment Job Options | 1239](#)
- [Selecting Software Images To Deploy | 1240](#)
- [Selecting Options for Software Deployment | 1241](#)
- [Summary of Software Deployment | 1243](#)

Specifying Software Deployment Job Options

To specify software deployment job options in the Select Devices page:

1. In the Job name field, enter a job name.
2. From the Device and deployment options list, select an option:
 - Select **Staging only (Download image to the device)** to download the software image to the device but not install it.
 - Select **Upgrade only (Install previously staged image on device)** to upgrade the device to a software image that was previously staged on the device.
 - Select **Staging and Upgrade (Download and Install image on device)** to download the software image and install it on the device.

Devices are not automatically rebooted after upgrade to make the device begin running the new software version. You can select the option to reboot the device automatically after the upgrade in a later wizard page.

3. Click **Next** to continue to the next page.

The Select Images page opens. Select a software image as described in [“Selecting Software Images To Deploy” on page 1240](#).

Selecting Software Images To Deploy

The Select Images page includes a table listing each device group and device that you selected for deployment. See [Table 292](#) for a description of the table columns.

If you selected the Upgrade only (Install previously staged image on device) option, only devices that contain a previously staged software image appear in the table. You cannot select a different image to install on these devices.

To select the software images to deploy, perform the following steps on the table row for each device group or individual device that you want to upgrade:

1. In the Proposed Image Version/Profile column, click **Select Image/Profile**.

The Select Image/Profile list is displayed.

2. From the Select Image/Profile list, select a software image.

TIP: To clear this field, select **Select Image/Profile** from the list.

3. After you finish selecting software images, click **Next** to continue to the next page.

The Select Options page opens.

TIP: A message notifies you if you do not select a software image for all the listed devices. This is just for your information. No action is taken on devices for which you do not select a software image. In effect, this removes those devices from the job.

Select options for software deployment as described in [“Selecting Options for Software Deployment” on page 1241](#).

Table 292: Select Images for Devices Table Description

Table Column	Description
Device Family	<p>Device family to which the device belongs. Devices are grouped by family. To display the devices within a device family, click the arrow next to the device family name.</p> <p>For example, all the EX4200 devices are listed under <i>EX4200</i> device group and all the Junos Fusion fabric devices are listed under <i>Fusion Enterprise</i> or <i>Fusion Datacenter</i> group. The Junos Fusion device family, when expanded displays the aggregation devices, satellite devices, and software upgrade groups, if defined.</p>

Table 292: Select Images for Devices Table Description (continued)

Table Column	Description
Count	Number of devices contained within a device family.
IP Address	Device's IP address.
Device Name	Device's name.
State	Device's state: <ul style="list-style-type: none"> • UP—Network Director can communicate with the device. • DOWN—Network Director cannot communicate with the device.
Running Image Version	Software version the device is running.
Proposed Image Version/Profile	Software version that is installed on the device when the job runs successfully. You can select to upgrade the software image on one or more individual device, device family, aggregation devices, satellite devices, or on a software upgrade group.

Selecting Options for Software Deployment

The options that you can configure in the Select Options page are described in [Table 293](#). The options that are available depend on the job flow you chose in the Select Images page.

After you finish selecting options, click **Next** to continue to the next page. The Summary page opens. Review the job summary as described in [“Summary of Software Deployment” on page 1243](#).

Table 293: Image Management Job Options

Option	Action
Select Options	
All Device Types	
Delete any existing image before download	Select to delete any existing software images on devices before downloading the new software image.

Table 293: Image Management Job Options (continued)

Option	Action
Reboot device after successful installation	<p>Select to reboot the device after the software image is installed. A reboot is required to begin running the new software version on the device.</p> <p>NOTE: This option might get disabled based on your details that you specify in the remaining fields. This indicates that for the options that you specified, the system automatically reboots the device as per the requirement during or after the image upgrade.</p>
Wired Devices	
Check compatibility with current configuration	Select to validate the software package or bundle against the current configuration as a prerequisite to adding the software package or bundle.
ISSU/NSSU	<p>Select if you want to perform an in-service software upgrade (ISSU) or a nonstop software upgrade (NSSU).</p> <p>ISSU enables you to upgrade between two different Junos OS releases with minimal disruption on the control plane and with minimal disruption of traffic.</p> <p>NSSU enables you to upgrade the software running on an EX Series switch with redundant Routing Engines or on most EX Series Virtual Chassis by using a single command and with minimal disruption to network traffic</p>
Archive data (Snapshot)	Select to take an archive snapshot of the files currently used to run the switch and copy them to an external USB storage device connected to the switch.
Copy to alternate slice	<p>Select to copy the new Junos OS image into the alternate root partition. This ensures that the resilient dual-root partitions feature operates correctly.</p> <p>This option is available only if you select Reboot device after successful installation.</p>
Wireless Devices	
Use Hitless Upgrade	Select to use the hitless upgrade process to upgrade the devices. Applies only to WLC controllers in cluster mode.
Select Schedule	
Stage now	Select Stage now to start staging software images to devices as soon as the job runs.
Stage later time	Select Stage later time to schedule the staging for a later time.
Staging Schedule	If you selected Stage later time, enter the date and time for staging to start.

Table 293: Image Management Job Options (continued)

Option	Action
Upgrade now	Select Upgrade now to start upgrading software images on devices as soon as staging finishes.
Upgrade later time	Select Upgrade later time to schedule the software upgrade for a later time.
Deployment Schedule	<p>If you selected Upgrade later time, enter the date and time for upgrade to start.</p> <p>If you scheduled staging, you must schedule the upgrade for at least 10 minutes after staging, to ensure that staging completes before upgrade starts.</p>

Summary of Software Deployment

On the Summary page, review the selections you made for the job. To change selections, click **Edit** in the area that you want to change. You can also click the boxes in the process flowchart above the wizard page to navigate between pages. When you are done making selections, click **Finish** on the Summary page to save the job, and run it if you configured the job to run immediately.

RELATED DOCUMENTATION

[Managing Software Images | 1174](#)

[Managing Software Image Deployment Jobs | 1243](#)

[Managing Software Images | 1235](#)

[Network Director Documentation home page](#)

Managing Software Image Deployment Jobs

This topic describes how to manage software image jobs. A software image job is created each time you deploy software images to devices or schedule a software image deployment. You can check the status of jobs, see job details, and cancel scheduled jobs.

To start managing software image jobs:

1. Click **Deploy** in the Network Director banner.
2. In the Tasks pane, select **Image Management > View Image Deployment Jobs**.

The Image Deployment Jobs page opens in the main window.

This topic describes:

- [Selecting Software Image Management Options | 1244](#)
- [Viewing Software Image Job Details | 1245](#)
- [Using the Device Image Staging Window | 1245](#)
- [Canceling Software Image Jobs | 1246](#)

Selecting Software Image Management Options

From the Image Deployment Jobs page, you can:

- Show deployment job details by selecting a job and clicking Show Details. See [“Viewing Software Image Job Details” on page 1245](#) for more information.
- Cancel a pending job by selecting the job and clicking Cancel Job. See [“Canceling Software Image Jobs” on page 1246](#) for more information.

[Table 294](#) describes the information provided in the of the Image Deployment Jobs table.

Table 294: Image Deployment Jobs Table

Table Column	Description
Job Id	An identifier assigned to the job.
Check box	Select to perform an action on the job in that row.
Job Name	Job name.
Percent	Percentage of the job that is complete.
Status	<p>Job status. The possible states are:</p> <ul style="list-style-type: none"> ● CANCELLED—The job was cancelled by a user. ● SCHEDULED—The job is scheduled but has not run yet. ● INPROGRESS—The job is running. ● SUCCESS—The job completed successfully. This state is applied if all of the devices in the job completed successfully. ● FAILURE—The job failed. This state is applied if any of the devices in the job failed. But some of the devices might have completed successfully. View the job details for the status of each device.
Summary	Job summary.
Scheduled Start Time	Job’s scheduled start time.

Table 294: Image Deployment Jobs Table (continued)

Table Column	Description
Actual Start Time	Time when the job started.
End Time	Time when the job ended.
User	User who created the job.
Recurrence	This field is not used for software image management jobs.

Viewing Software Image Job Details

To view the details of a software image job:

1. Select the job in the table.
2. Click **Show Details**.

The Device Image Staging window opens. See [“Using the Device Image Staging Window” on page 1245](#) for a description of the window.

Using the Device Image Staging Window

Use the Device Image Staging window to view information about software image jobs. [Table 295](#) describes this window.

Table 295: Device Image Staging Window Description

Field	Description
Job Name	Job name.
Start Time	Job's scheduled start time.
End Time	Time when the job ended.
% Complete	Percentage of the job that is complete.

Table 295: Device Image Staging Window Description (*continued*)

Field	Description
Status	Job status. The possible statuses are: <ul style="list-style-type: none"> ● CANCELLED—The job was cancelled by a user. ● SCHEDULED—The job is scheduled but has not run yet. ● INPROGRESS—The job is running. ● SUCCESS—The job completed successfully. ● FAILURE—The job failed.
Host Name	Host name of device.
Status	Device status. The possible statuses are: <ul style="list-style-type: none"> ● INPROGRESS—The job is running. ● SUCCESS—The job completed successfully. ● FAILURE—The job failed.
% Complete	Percentage of the job that is complete on the device.
Start Time	Time when the job started on the device.
End Time	Time when the job ended on the device.
Description	Description of the job on the device. Can include error messages for failed devices.
Close	Click to close the window.

Canceling Software Image Jobs

To cancel a software image job:

1. Select the job in the table.
2. Click **Cancel**.

RELATED DOCUMENTATION

[Managing Software Images | 1174](#)

[Deploying Software Images | 1238](#)

Managing Software Images | **1235**

[Network Director Documentation home page](#)

Managing Devices

IN THIS CHAPTER

- [Converting Automatically Discovered Access Points to Manually Configured Access Points | 1248](#)
- [Enabling or Disabling Network Ports on Switches | 1250](#)
- [Understanding Node Groups for a QFabric System | 1251](#)
- [Creating and Managing Node Groups for a QFabric System | 1252](#)
- [Converting the QSFP+ Ports on QFX Series and QFabric Devices | 1256](#)

Converting Automatically Discovered Access Points to Manually Configured Access Points

When Auto AP mode is enabled on controllers, all access points are recognized but not persistently configured on the controller. Instead, the access points are dynamically added to the controller each time the access points boot up.

NOTE: For information about enabling Auto AP mode, see [“Creating and Managing Wireless Auto AP Profiles” on page 980](#).

You can convert access points that were dynamically added to a controller using an Auto AP profile to a persistent access point configuration on the controller. This topic describes converting dynamic access point configurations to persistent access point configurations.

To convert automatically discovered access points to persistent access points:

1. Click **Deploy** mode in the Network Director banner.
2. Under Wireless Network in the View pane, select either a controller or a cluster of controllers.
3. Click **Convert Auto AP** under Device Management in the Tasks pane.

The Convert Auto AP page opens, displaying a list of automatically discovered access points and the information listed in [Table 296](#).

Table 296: Automatically Discovered Access Point Information

Field	Description
AP Number	Temporary access point number assigned by the controller.
AP Name	Temporary access point name assigned by the controller.
Model	Access point model number discovered by the controller.
Serial Number	Access point model number discovered by the controller.
IP Address	Temporary IP address assigned by the controller.

- From the list of automatically discovered access points, select one to convert to a persistent access point, and then click **Convert Auto AP**.
- This message is displayed: **This will convert Auto APs to configured APs. Do you want to continue?**
Click **Yes**.

The Convert Auto AP Job details window is displayed with the converted access points.

RELATED DOCUMENTATION

[Understanding Auto AP Profiles | 883](#)

[Adding and Managing an Individual Access Point | 1156](#)

[Creating and Managing Wireless Auto AP Profiles | 980](#)

[Network Director Documentation home page](#)

Enabling or Disabling Network Ports on Switches

Network ports connect switches to the network and carry network traffic. You can enable or disable network ports of switches that are part of your network. When you enable or disable a port, the administrative status of the port changes to UP or DOWN respectively. When you disable a port, the system marks that port as administratively down, without removing the port configurations.

You can enable or disable one or more ports at a time using the Manage Port Admin State page. The status of the port is indicated by the Admin State and the Link State fields. The administrative status of a port is indicated by the Admin State field.

To enable or disable a network interface:

1. Do one of the following:

- In the topology view, locate the device for which you want to enable or disable ports and click **Device Management > Manage Port Admin State** from the Tasks pane.
- While in the Deploy mode, select the device for which you want to enable or disable ports in the View pane and click **Device Management > Manage Port Admin State** from the Tasks pane.

The Manage Port Admin State page appears displaying all the physical ports available on the selected device and the current status of each port. This page also displays the port mode of each interface, if any. Port mode can be access, tagged-access, or trunk mode.

2. Do one of the following:

- Select the check box adjacent to the ports that you want to enable and click **Change Admin State UP**.
- Select the check box adjacent to the interfaces that you want to disable and click **Change Admin State DOWN**.

3. Click **Done**. Network Director changes the administrative status of the ports and displays a confirmation message confirming the changes.

RELATED DOCUMENTATION

[Understanding the Network Director User Interface | 85](#)

[Network Director Documentation home page](#)

Understanding Node Groups for a Qfabric System

IN THIS SECTION

- [Network Node Groups | 1251](#)
- [Server Node Groups | 1252](#)

Node groups help you combine multiple QFabric Node devices into a single virtual entity within the QFabric system to enable redundancy and scalability at the edge of the data center.

This topic covers:

Network Node Groups

A set of one or more Node devices that connect to an external network is called a *network Node group*. The network Node group also relies on two external Routing Engines running on the Director group. These redundant *network Node group Routing Engines* run the routing protocols required to support the connections from the network Node group to external networks.

When configured, the Node devices within a network Node group and the network Node group Routing Engines work together in tandem as a single entity. By default, network Node group Routing Engines are part of the **NW-NG-0** network Node group but no Node devices are included in the group. As a result, you must configure Node devices to be part of a network Node group.

In a QFabric system deployment that requires connectivity to external networks, you can modify the automatically generated network Node group by including its preset name **NW-NG-0** in the Node group configuration. Within a network Node group, you can include a minimum of one Node device up to a maximum of eight Node devices. By adding more Node devices to the group, you provide enhanced scalability and redundancy for your network Node group.

NOTE: The QFabric system creates a single **NW-NG-0** network Node group for the default partition. You cannot configure a second network Node group inside the default partition. The remaining Node devices within the default partition are reserved to connect to servers, storage, or other endpoints internal to the QFabric system. These Node devices either can be retained in the automatically generated server Node groups or can be configured as part of a redundant server Node group.

Server Node Groups

A *server Node group* is a set of one or more Node devices that connect to servers or storage devices. Unlike Node devices that are part of a network Node group and rely on an external Routing Engine, a Node device within a server Node group connects directly to endpoints and implements the Routing Engine functions locally, using the local CPU built into the Node device itself.

By default, each Node device is placed in its own self-named autogenerated server Node group to connect to servers and storage. You can override the default assignment by manually configuring a redundant server Node group that contains a maximum of two Node devices. You can use a redundant server Node group to provide multihoming services to servers and storage, as well as configure aggregated LAG connections that span the two Node devices.

NOTE: The Node devices in a redundant server Node group must be of the same type, either a QFX3500 Node device or a QFX3600 Node device. You cannot add a QFX3500 and a QFX3600 Node device to the same redundant server Node group.

RELATED DOCUMENTATION

[Creating and Managing Node Groups for a QFabric System | 1252](#)

[Setting Up QFabric Systems | 780](#)

[Network Director Documentation home page](#)

Creating and Managing Node Groups for a QFabric System

IN THIS SECTION

- [Managing Node Groups | 1253](#)
- [Creating Node Groups | 1254](#)
- [Specifying Settings for a Node Group | 1254](#)

Node groups help you combine multiple QFabric Node devices into a single virtual entity within the QFabric system to enable redundancy and scalability at the edge of the data center.

To start managing nodes groups:

1. Click **Deploy** in the Network Director banner.
2. Select the fabric to manage in the network view pane.
3. Click **Manage Node Group** under Device Management in the Tasks pane.

The Manage Node Groups page appears.

This topic describes:

Managing Node Groups

Use the Manage Node Groups page to manage existing node groups and to create new ones. The configured node groups are listed in expandable and collapsible groups:

- Group: NNG—The network node group. Each fabric has only one network node group, so you cannot create another. You can edit the group's membership.
- Group: RSNG—Redundant server node groups. You can create new redundant server node groups and edit existing ones.
- Group: SNG—Server node groups. You can create new server node groups and edit existing ones.

From the Manage Node Groups page, you can:

- Create a new node group by clicking **Add**.
- Modify an existing node group by selecting it and clicking **Edit**.
- Delete a node group by selecting it and clicking **Delete**.

[Table 297](#) describes the information provided about node groups on the Manage Node Groups page. This page lists all node groups within the fabric that is selected in the network view.

Table 297: Manage Node Groups Information

Column	Description
Node Group Name	Name given to the node group when it was created.
Description	Description of the node group entered when it was created. TIP: To display the entire description, you might need to resize the Description column by clicking the column border in the heading and dragging it.

Table 297: Manage Node Groups Information (continued)

Column	Description
Selected Nodes	<p>Nodes that belong to the node group.</p> <p>The connection state of the node is indicated by an icon: The green upward pointing arrow indicates the node is connected. The red downward pointing arrow indicates the node is not connected.</p>
Deployment State	Shows whether the node group was deployed.

Creating Node Groups

To create a node group:

1. Click **Deploy** in the Network Director banner.
2. Select the fabric to manage in the network view pane.
3. Click **Manage Node Group** under Device Management in the Tasks pane.
The Manage Nodes Groups page appears.
4. Click **Add**.
The Create Node Group page appears.
5. Configure the new node group as described in [“Specifying Settings for a Node Group” on page 1254](#).
6. Click **OK**.
The node group is listed on the Manage Node Groups page with the deployment state Pending Deployment.
7. To deploy all undeployed node group changes, click **Deploy Now**.

Specifying Settings for a Node Group

Use the Create Node Group page to define the devices in the node group.

[Table 298](#) describes the settings available on the Create Node Group page.

Table 298: Create Node Group Settings

Field	Action
Node Group Type	<p>Select a node group type:</p> <ul style="list-style-type: none"> • SNG—Server node group. • RSNG—Redundant server node group.
Node Group Name	<p>Type the name of the group.</p> <p>The following rules apply to QFabric Node group naming:</p> <ul style="list-style-type: none"> • Node group names must use alphabetic (A through Z and a through z), numeric (0 through 9), or dash (-) characters. • The maximum length of a Node group name is 30 characters. • Node group names are case sensitive. For example, MY-NG-1 and my-ng-1 refer to different components. • You cannot use the reserved names all, fabric, or director-group as a Node group name.
Description	Type a description of the Node group, which will appear on Manage Node Groups page.
Node Selection	
Add	<p>Click to add a device to the node group:</p> <ol style="list-style-type: none"> 1. Click Add. 2. Select devices from the Node Device Selection page. <p>You can select one device for a server node group, and two devices for a redundant server node group.</p> 3. Click OK.
Remove	<p>Click to remove a device from the node group:</p> <ol style="list-style-type: none"> 1. Select a device from the list on the Node Device Selection page. 2. Click Remove.

RELATED DOCUMENTATION

Converting the QSFP+ Ports on QFX Series and QFabric Devices

IN THIS SECTION

- [Selecting Devices | 1256](#)
- [Converting Ports | 1258](#)
- [Reviewing and Deploying Port Conversions | 1259](#)

The 40-Gbps QSFP+ ports on QFX Series and QFabric devices can be configured to operate as four 10-Gigabit Ethernet (*xe*) ports, one 40-Gigabit Ethernet (*xle*) port, or one 40-Gbps data uplink (*fte*) port (for QFabric devices).

To start converting QSFP+ ports:

1. Click **Deploy** in the Network Director banner.
2. Select the node in the View pane that contains the ports you want to convert.
3. Select the task **Device Management > Convert Ports** in the Tasks pane.

The Ports Conversion wizard opens to the Device Selection page. Continue with [“Selecting Devices” on page 1256](#).

This topic describes:

Selecting Devices

Use the Device Selection page to select the devices that contain the ports you want to convert.

To select devices that contain the ports you want to convert:

1. Select the device family to filter and display devices from that device family. You can select either **Data Center ELS** or **Data Center non ELS**.

If you selected Data Center ELS, proceed to step 3.

2. If you selected Data Center non ELS, select the **Standalone** radio button or the **QFabric** radio button. The device list displays only devices of the selected type.

3. Select the devices that contain the ports you want to convert by selecting their check boxes.

4. Click **Next**.

The Convert Ports page opens. Continue with section [“Converting Ports” on page 1258](#).

[Table 299](#) describes the information provided about devices on the Device Selection page. This page lists all the devices in the selected scope that contain QSFP+ ports.

Table 299: Port Conversion Device Selection Page

Column	Description
Hostname	Configured name of the device or IP address if no hostname is configured.
IP Address (Standalone devices only)	IP Address of the device.
NodeGroup Name (QFabric devices only)	Name of the node group the port belongs to.
Serial Number	Serial number on device chassis.
Platform	Model number of the device.
Connection State	Connection status of the device in Network Director: <ul style="list-style-type: none"> ● UP—Device is connected to Network Director. ● DOWN—Device is not connected to Network Director. ● N/A—Connection status is unavailable to Network Director.

Table 299: Port Conversion Device Selection Page (continued)

Column	Description
Config State (Standalone devices only)	<p>Displays the configuration status of the device:</p> <ul style="list-style-type: none"> ● In Sync—The configuration on the device is in sync with the Network Director configuration for the device. ● Out Of Sync—The configuration on the device does not match the Network Director configuration for the device. This state is usually the result of the device configuration being altered outside of Network Director. You cannot deploy configuration on a device from Network Director when the device is Out Of Sync. To resolve this state, use the Resynchronize Device Configuration task in Deploy mode. ● Sync failed—An attempt to resynchronize an Out Of Sync device failed. ● Synchronizing—The device configuration is in the process of being resynchronized. ● N/A—The device is down.

Converting Ports

Use the Convert Ports page to convert QSFP+ ports between port types.

The page contains a table in which you configure the port conversion. The Port Name (Default) column displays the default port name.

To convert QSFP+ ports:

1. To convert a port, click its **Convert to Port** column.
2. Select an option from the list that opens:
 - **No Change**—Does not change the port type.
 - **fte**—Configures the port as one 40-Gbps data uplink port (for QFabric nodes only).
 - **xle**—Configures the port as one 40-Gigabit Ethernet port.
 - **xe**—Configures the port as a group of 10-Gigabit Ethernet ports.
3. The Port Name (After Conversion) column displays what the port name will be if you commit the current settings.
4. When you finish making port type settings, click **Next**.

The Review page opens. Continue with [“Reviewing and Deploying Port Conversions”](#) on page 1259

Reviewing and Deploying Port Conversions

Use the Review page to review settings and deploy the port conversion:

1. To change settings from the Review page, click **Back** to return to previous wizard pages.
2. When you finish making changes, click **Deploy** to deploy the port conversion to the selected devices.

RELATED DOCUMENTATION

[Understanding Deploy Mode in Network Director | 1172](#)

[Network Director Documentation home page](#)

Setting Up Zero Touch Provisioning for Devices

IN THIS CHAPTER

- [Understanding Zero Touch Provisioning in Network Director | 1260](#)
- [Configuring and Monitoring Zero Touch Provisioning | 1261](#)

Understanding Zero Touch Provisioning in Network Director

Zero touch provisioning allows you to provision new Juniper Networks switches in your network automatically—without manual intervention. When you physically connect a switch to a network and boot it with the factory-default configuration, the switch attempts to upgrade the Junos OS software automatically and autoinstall a configuration file from the network. Use the Zero Touch Provisioning wizard to create a profile that applies all the configurations to a Dynamic Host Configuration Protocol (DHCP) server that you configure. You can apply one or more profiles to a DHCP server.

After you enable zero touch provisioning for a DHCP server that is part of a given subnet in your network, and connect a new switch to that subnet, the following series of events occurs:

1. The switch contacts the DHCP server to obtain an IP address. The DHCP server assigns an IP address to the switch. The DHCP server also passes on the location of the software image, and the configuration file to the switch. This information is passed on to the DHCP server from Network Director when you create and save a zero touch provisioning profile.
2. The switch uses this information to locate the software image, and the configuration file. These files are stored in an FTP, TFTP, or an HTTP server.
3. The switch then upgrades the operating system version by using the software image and loads the configuration file.

NOTE: You can use zero touch provisioning to provision EX Series switches to run Junos OS Release 12.3R5 and 13.3 only. If a switch is provisioned with any other Junos OS Release, then Step 4 is not applicable. You must manually discover the switch from Network Director to be able to manage it.

4. After a successful upgrade, the switch sends out an trap message to Network Director to announce that a new switch has been deployed in the network. If the trap message is successfully received, Network Director adds the switch to the Network Director's inventory. This eliminates the need to manually discover new devices that are added to your switching network.

NOTE: if the SNMP trap that the switch sends to Network Director does not reach the destination, then Network Director does not know about the new device and the device will not be added to the Network Director's inventory. In such a scenario, you must manually discover the new device from Network Director.

For more information on zero touch provisioning for switches, see [Understanding Zero Touch Provisioning](#).

RELATED DOCUMENTATION

[Configuring and Monitoring Zero Touch Provisioning | 1261](#)

[Network Director Documentation home page](#)

Configuring and Monitoring Zero Touch Provisioning

IN THIS SECTION

- [Configuring Zero Touch Provisioning | 1262](#)
- [Specifying the Server Details | 1263](#)
- [Specifying the Software Image and Configuration Details | 1265](#)
- [Reviewing and Modifying Zero Touch Provisioning Settings | 1266](#)
- [What To Do Next | 1266](#)
- [Configuration Statements for Custom Configuration of DHCP Server | 1266](#)
- [Monitoring Zero Touch Provisioning Profiles | 1267](#)

Zero touch provisioning (ZTP) allows you to provision new switches in your network automatically—without manual intervention. When you physically connect a switch to a network and boot it with the factory-default configuration, the switch attempts to upgrade the Junos OS software automatically and autoinstall a configuration file from the network.

The switch uses information that you configure on a Dynamic Host Control Protocol (DHCP) server to determine whether to perform these actions and to locate the necessary software image and configuration files on the network. You can configure the DHCP server by using a zero touch provisioning profile. If you do not configure a DHCP server, the switch boots with the preinstalled software and the default configuration.

The type of DHCP server that you want to use determines whether Network Director configures the DHCP server for you or whether you must manually configure the DHCP server. If you select CentOS or Ubuntu DHCP servers, Network Director configures the DHCP server by using the details that you specified in the zero touch provisioning profile. If you use any other DHCP server, you must manually configure the DHCP server. For such DHCP servers, you can use Network Director only to monitor the switches once they are provisioned. For details on configuring a DHCP server manually, see the DHCP server documentation.

For more information on zero touch provisioning for switches, see [Understanding Zero Touch Provisioning](#).

Before you begin, ensure that you have the necessary privileges on the FTP and the file server that Network Director uses for zero touch provisioning. For more details, see “[User Privileges Required for the DHCP and File Server While Using Zero Touch Provisioning](#)” on page 743.

NOTE: For detailed information about DHCP and DHCP options, see RFC2131 (<http://www.ietf.org/rfc/rfc2131.txt>) and RFC2132 (<http://www.ietf.org/rfc/rfc2132.txt>). These documents refer to Internet Systems Consortium (ISC) DHCP version 4.2. For more information about this version, see <http://www.isc.org/software/dhcp/documentation>.

Configuring Zero Touch Provisioning

Before you begin:

Ensure that the switch has access to the following network resources:

- The DHCP server that provides the location of the software image and configuration files on the network
See your DHCP server documentation for configuration instructions.
- The File Transfer Protocol (anonymous FTP), the Hypertext Transfer Protocol (HTTP) server, or the Trivial File Transfer Protocol (TFTP) server on which the software image and configuration files are stored. If you are using an FTP server, ensure that the FTP server is configured to enable anonymous access. Refer to your FTP server documentation to know more about this.

NOTE: Although TFTP is supported, we recommend that you use FTP or HTTP instead, because these transport protocols are more reliable.

- (Optional) A Network Time Protocol (NTP) server to perform time synchronization on the network
- (Optional) A system log (syslog) server to manage system log messages and alerts

Identify the type of DHCP server that you will be using for zero touch provisioning:

- CentOS DHCP Server—If your DHCP server uses the following command to restart the server, then select **CentOS** as the DHCP server type:

```
service dhcpd restart
```

- Ubuntu DHCP Server—If your DHCP server uses the following command to restart the server, then select **Ubuntu** as the DHCP server type:

```
service isc-dhcp-server restart
```

- Other—If your server is not an ISC DHCP server running on Linux operating system, then you must select **Other** and configure the DHCP server manually.

NOTE: CentOS 6.10 is the supported or qualified version of CentOS for the DHCP server in Network Director 4.0 release.

For information about the CentOS and Ubuntu versions supported by Network Director Release 4.0, see the *Supported Platforms* section of the [Network Director Release Notes](#).

To configure zero touch provisioning:

1. While in the Deploy mode, select **Zero Touch Provisioning** > **Manage ZTP** from the Tasks pane.
The Manage ZTP Profiles page appears.
2. Specify the server details in the Server Setup wizard page as described in “[Specifying the Server Details](#)” on page 1263.

Specifying the Server Details

To configure the server settings:

1. Enter the settings described in [Table 300](#). Required settings are indicated in the user interface by a red asterisk (*) that appears next to the field label.

Table 300: Server Details

Field	Description
Profile Name	Name of the zero touch provisioning profile.
DHCP Server Info	
DHCP Server Type	<p>The type of DHCP server that provides the necessary information to the switch. You can choose to configure a CentOS DHCP server, an Ubuntu DHCP server, or any other DHCP server.</p> <p>If you select Other, Network Director also selects the Manually Configure Server check box and hides all the other details except the File Server Details. You must configure the DHCP server manually.</p>
Manually Configure Server	<p>Select to indicate that you want to manually configure the DHCP server. You can configure the CentOS and Ubuntu DHCP servers manually or from Network Director.</p> <p>If you select Manually Configure Server check box, Network Director hides all the other details except the File Server Details.</p>
DHCP Server	IP address or the hostname of the DHCP server.
DHCP User	<p>Username for the DHCP server.</p> <p>NOTE: This user must have write permission for the <code>dhcpcd.conf</code> file.</p>
DHCP Password	Password for the specified username.
Confirm Password	Confirm the password.
File Transfer Server Info	
File Server	The type of file server where the software images and the configuration files are to be stored. You can choose to use an FTP, HTTP, or a TFTP file server.
File Server IP	IP address or the hostname of the file server.
File Server Root Dir	The root directory of the file server.
Optional Settings	
Syslog Server IP	IP address of the system log server, if you want to perform data logging for zero touch provisioning.
NTP Server IP	IP address of the NTP server, if you want to use time synchronization.

2. Click **Next** and proceed to specify the software image, configuration file, and the IP address range to be configured on the DHCP server. For more details, see [“Specifying the Software Image and Configuration Details” on page 1265](#).

Specifying the Software Image and Configuration Details

To specify the software image, configuration file, and the IP address range to be configured on the DHCP server:

1. Enter the password that you want to set for the root user on the switch, in the ZTP Devices Root User Password field and confirm the password in the Confirm Password field.

NOTE: Once the switch is successfully provisioned, Network Director uses this password for discovering the device.

2. In the Configure Settings table, click **Add** to specify details for a switch model.
Network Director adds a row to the Configure Settings table.
3. In the Device Model field, select the switch model for which you want to specify the image and configuration file details.
4. (Only for the CentOS DHCP server) In the Image File field, select the image file that you want to upload for the selected switch model. This field lists the software images that you have uploaded to Network Director from the Device Image Repository page. For details about uploading a software image, see [“Managing Software Images” on page 1235](#).
5. Do one of the following to upload the configuration file to the DHCP server:
 - Select the factory-default configuration file for the selected switch model in the Config File field. Network Director ships with a factory-default configuration for all supported switch models.
 - If you want to upload a custom configuration file for the given switch model, click **Upload Config** and select a configuration file. When you upload a custom configuration file, ensure that the configurations mentioned in [“Configuration Statements for Custom Configuration of DHCP Server” on page 1266](#) are included in the configuration file.
6. In the Subnet field, specify the subnet that the DHCP server caters to.
7. In the From IP and To IP fields, specify the range of IP addresses that the DHCP server can assign to new switches.

8. (Only for the CentOS or Ubuntu DHCP server) Click **Export DHCP Config** if you want to view the configuration that Network Director sends to the DHCP server.

Network Director downloads the configuration and you can view it using any text editor. If you chose to configure the DHCP server manually in the Server Details page, you can use this configuration file to complete the manual configuration.

9. Click **Next** to review the details of the zero touch provisioning profile that you created.

Reviewing and Modifying Zero Touch Provisioning Settings

From this page, you can save or make changes to a zero touch provisioning profile:

- To make changes to the profile, click the **Edit** button associated with the configuration you want to change.

Alternatively, you can click the appropriate buttons in the zero touch provisioning workflow at the top of the page that corresponds to the configuration you want to change.

When you are finished with your modifications, click **Review** to return to this page.

- To save a zero touch provisioning profile or to save modifications to the settings of an existing profile, click **Finish**.

What To Do Next

- For manual configuration, use the DHCP configuration file to manually configure the DHCP server. If you selected the DHCP server as CentOS or Ubuntu, Network Director uploads the software image to the file server that you specified. If you selected any other DHCP server, you must manually upload the software image to the file server and specify the path when you configure the DHCP server.
- (Only for the CentOS or Ubuntu DHCP servers) For automatic configuration, Network Director configures the DHCP server with the details that you specified in the zero touch provisioning profile and uploads the software image to the file server that you specified.

Configuration Statements for Custom Configuration of DHCP Server

Insert the following configuration statements to the configuration file, if you want to upload a custom configuration file to the DHCP server:

```
system {
  root-authentication {
    encrypted-password "PASSWORD"; ## SECRET-DATA
  }
}
```

```

}
event-options {
policy target_add_test {
  events snmpd_trap_target_add_notice;
  then {
    raise-trap;
  }
}
}
trap-group networkdirector_trap_group {
version all;
destination-port NDPORT;
categories {
link;
services;
authentication;
}
targets{
NDIP;
}
}

```

Monitoring Zero Touch Provisioning Profiles

You can use the Monitor ZTP Profiles page to view details about the switches that were provisioned using a given zero touch provisioning profile and added successfully to the Network Director inventory.

To monitor a zero touch provisioning profile:

1. While in the Deploy mode, select **Zero Touch Provisioning > Monitor** from the Tasks pane. The Monitor ZTP Profiles page appears.
2. In the Choose ZTP Profile box, select the zero touch provisioning profile that you want to monitor.
Network Director displays the zero touch provisioning summary and details of switches that were discovered using the selected profile.

RELATED DOCUMENTATION

[Understanding Zero Touch Provisioning in Network Director | 1260](#)

[Managing Software Images | 1235](#)

[Network Director Documentation home page](#)

5

PART

Monitoring Devices and Traffic

[About Monitor Mode | 1269](#)

[Monitoring Traffic | 1283](#)

[Monitoring Client Sessions | 1314](#)

[Monitoring Radio Frequency | 1322](#)

[Monitoring Devices | 1347](#)

[Monitoring and Analyzing Fabrics | 1360](#)

[Monitoring Virtual Networks | 1371](#)

[General Monitoring | 1378](#)

[Monitor Reference | 1383](#)

About Monitor Mode

IN THIS CHAPTER

- [Understanding Monitor Mode in Network Director | 1269](#)
- [Understanding the Monitor Mode Tasks Pane | 1277](#)

Understanding Monitor Mode in Network Director

IN THIS SECTION

- [Scope and Monitor Tab Availability | 1270](#)
- [Monitors and Tasks | 1271](#)
- [Scope and Data Aggregation | 1272](#)
- [How Network Director Collects and Displays Monitoring Data | 1273](#)
- [How Network Director Displays and Stores Trend Data | 1273](#)
- [More About the Monitor Tabs | 1274](#)

Monitor mode in Network Director provides you visibility into your network status and performance. Network Director monitors its managed devices and maintains the information it collects from the devices in a database. Monitor mode displays this information in easy-to-understand graphs and in tables that you can sort and filter, allowing you to quickly visualize the state of your network, spot trends developing over time, and find important details.

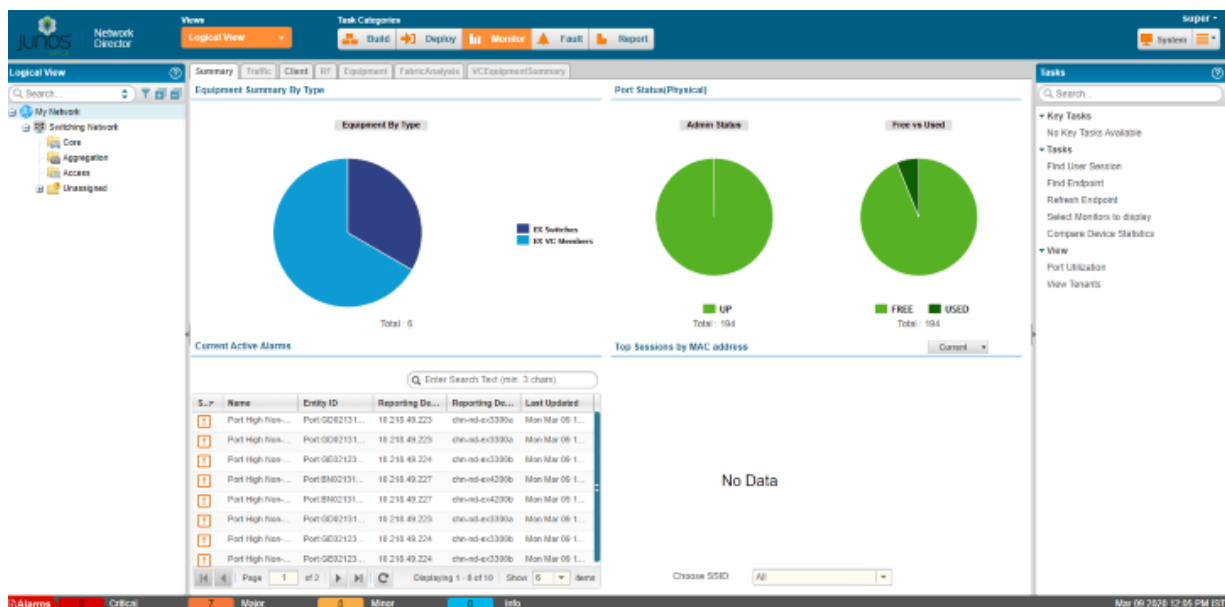
NOTE: Monitor mode for the Datacenter View displays information specific to the virtual networks, virtual switches, and hosts. To know more about Monitor mode for Datacenter View, see [“Using Monitor Mode for Virtual Devices” on page 1371](#).

Monitor mode divides monitoring activity into the following categories:

- Traffic—Provides information about traffic on switches, wireless LAN controllers, and interfaces.
- Client—Provides session information about clients connected to wireless access points and to 802.1X authenticator switch ports.
- RF—Provides information about the wireless environment and signal performance.
- Equipment—Provides information about the state of switches, wireless LAN controllers, interfaces, wireless access points, and radios.
- Fabric Analysis—Displays the results of running the Run Fabric Analyzer task on a QFabric or Virtual Chassis Fabric (VCF). It shows information about the health, connectivity, and topology of the fabric.
- VC Equipment Summary—Provides information about the operational status of the virtual chassis devices.

You can access these categories through tabs on the Monitor mode landing page, as shown in [Figure 52](#). An additional tab, the Summary tab, is available that provides a high-level dashboard for the scope selected in the View pane. The monitoring information displayed in the Summary tab also appears on other tabs.

Figure 52: Monitor Mode Landing Page and Tabs



This topic describes:

Scope and Monitor Tab Availability

Your current scope—that is, your view and node selection in the View pane—affects which Monitor tabs are available. For example, if you select a switch, the RF tab is not available.

The shading of the tabs indicate whether a tab is selected, available, or not available:

- The currently selected tab has dark text on a light background.

- Tabs that are available but not selected have dark text on a dark background.
- Tabs that are not available for your current scope have light text on a light background.

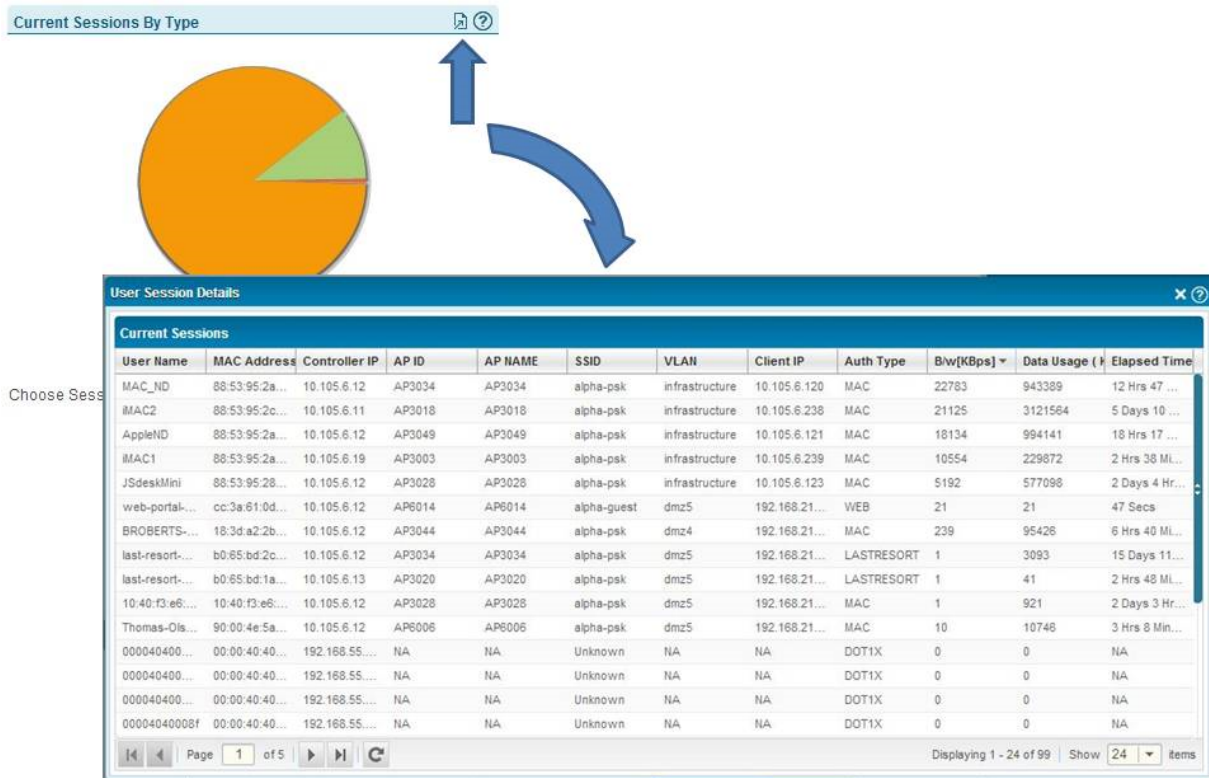
When you enter Monitor mode from another mode, the Summary tab is selected for all scopes. If you have selected a tab and then change scope, the tab remains selected if it is supported in the new scope. If it is not supported in the new scope, Network Director selects a default tab for that scope.

Monitors and Tasks

When you click a Monitor tab, the landing page for that tab is displayed, which contains a set of monitors. These monitors enable you to see at a glance important information about the aspect of your network being monitored. For example, the monitors in the Client tab present high-level information about the sessions in the selected scope: the users and client sessions consuming the most bandwidth, the distribution of active sessions by type, and the trend in session count over time.

Detailed information is also available from many monitors when you click the Details icon on the monitor. If the Details icon is not visible in the title bar of a monitor, mouse over the monitor to make it visible. For example, if you click the Details icon from the Current Sessions By Type monitor, you can view detailed information about the current sessions, as shown in [Figure 53](#).

Figure 53: Accessing Session Details from the Current Session by Type Monitor



In addition to monitors, each tab provides a set of tasks available from the Tasks pane. These tasks enable you to perform additional monitoring functions. Some tasks enable you to view more specialized monitoring data; others enable you to perform an operation, such as pinging a host. For a complete list of tasks available in Monitor mode, see [“Understanding the Monitor Mode Tasks Pane”](#) on page 1277.

The scope you select affects which monitors are displayed and which tasks are available. In the Equipment tab, for example, you see a different set of monitors for an EX Series switch than you see for a Wireless LAN controller.

Scope and Data Aggregation

Network Director enables you to monitor one or more devices. It provides a broader network view by aggregating data from devices and making that data available for viewing at higher scopes within the network.

A typical example is RF interference data. Network Director associates RF interference data with the radio that reported it. You can select a radio in the View pane to view the interference data reported by that radio. However, you can also view the RF interference data for the entire wireless network or for a particular

location (floor, building, or site). At each of these scopes, Network Director combines or aggregates the data associated with all the radios included in that scope.

Not all data is aggregated at higher scopes. For example, it does not make sense to provide power supply status at any higher scope than the device itself. Whenever monitors are available at a scope higher than the device scope, however, the data presented is aggregated data from all devices contained in that scope.

How Network Director Collects and Displays Monitoring Data

Network Director collects monitoring data from all its managed devices at regular intervals known as polling intervals. These polling intervals can vary according to the type of data being collected. Network Director sets default polling intervals for each type of data—you can, however, change these polling intervals in Preferences.

The polling intervals are aligned to clock time. For example, if the polling interval is set to 5 minutes, then within every hour, Network Director collects data at :00, :05, :10, :15, and so on. If the polling interval is set to 15 minutes, Network Director collects data within every hour at :00, :15, :30, and :45.

Network Director uses the Juniper Networks Device Management Interface (DMI) to the managed devices to collect the data. If you have a Junos Space fabric, Network Director balances the load of polling the managed devices across the nodes in the fabric.

When you display a monitor, the current data is from the last polling interval. Displaying or refreshing a monitor does not trigger Network Director to collect data. However, Network Director automatically refreshes monitors with new data after a polling interval completes. Each monitor displays the time that the data was last refreshed.

The detail windows for monitors are not automatically refreshed after a polling period completes. You must manually refresh them to obtain new polling data.

How Network Director Displays and Stores Trend Data

In addition to displaying current data, Network Director also displays historical data in trend graphs so that you can view trends in network performance over time.

When you display a trend graph, you can select the time period over which the data is displayed—usually 1 hour, 8 hours, 1 day, 1 week, 1 month, 3 months, 6 months, or 1 year. These predefined periods are always relative to the current time and date—that is, if you select a week, the data is from the last 7 days. You can also define a custom time period, which enables you to display data for a period between specific dates and times.

For a trend graph displaying a predefined period of 1 hour, the number of data points depends on the configured polling interval. For periods greater than an hour, the number of data points displayed depends on the time period selected and how Network Director consolidates data over time.

To allow storing of monitoring data for a long period of time, Network Director consolidates older data. Consolidation involves deriving a single value from a set of shorter term values, generally by averaging the shorter term values, and then using that value as a data point in a longer term data set. After the shorter term data is consolidated into longer term data, it is discarded to save storage space. For example, if a value is polled every 5 minutes, the set of 12 values is consolidated into a single value after an hour has passed. That value then becomes one of the 24 data points that makes up the data set for a day. Similarly, after a day has passed, data is consolidated into one data point that represents that day; after a month has passed, data is consolidated into a one data point that represents that month. Data is not kept for more than a year. You can, however, run reports on some monitoring data in Report Mode and archive the reports to maintain a history that is longer than a year.

For all trend graphs, Network Director will not display data until it has more than two data points to display. This means that after you discover a device, trend data will not appear until three polling periods have passed.

More About the Monitor Tabs

IN THIS SECTION

- [The Summary Tab | 1274](#)
- [The Traffic Tab | 1275](#)
- [The Client Tab | 1275](#)
- [The RF Tab | 1275](#)
- [The Equipment Tab | 1275](#)
- [The Fabric Analysis Tab | 1276](#)
- [VC Equipment Summary Tab | 1276](#)

The following sections provide more information about each tab in Monitor mode.

The Summary Tab

The Summary tab is displayed whenever you enter Monitor mode. It serves as a high-level dashboard for the current selected scope in the View pane.

The monitors displayed in the Summary tab can belong to any of the Monitor categories. Each scope has a predefined set of monitors that are displayed. For example, if your scope is the Wireless Network, the monitors on the Summary tab summarize the status of wireless equipment in the network, the interference sources in the network, the alarms active on the wireless devices, and the number of sessions in the wireless network.

When you select an individual device in the View pane, the Summary tab itself displays an arrow that indicates whether the device is up (green up arrow) or down (red down arrow).

For the My Network scope, you can customize what monitors appear on Summary tab, giving you the ability to view at a glance those aspects of network health and performance that are most important to you.

The Traffic Tab

The Traffic tab provides information for analyzing traffic on switches and wireless LAN controllers. The four monitors provide an aggregated view of all network traffic on a device, such as proportion of current proportion of multicast, unicast, broadcast traffic or the trend in packet errors. Tasks provide more detailed looks at traffic, such as traffic statistics for individual ports or the degree in which a port's bandwidth is being used.

The Client Tab

The Client tab provides information about clients and sessions on the network. A client is any device that is connected to the network through a wireless access point or through an access port on a switch that is an 802.1X authenticator port. Examples of clients include VoIP phones, laptops, printers, security cameras, and so on. When a client connects to the network, a session starts, which is uniquely identified by the MAC address of the client.

The Client tab monitors provide a view of overall client session activity in the selected scope. They show the total number of sessions, sessions consuming the most bandwidth, and trends in the number of sessions. Detailed views provide information about each client, such as MAC address, IP address, username, client VLAN, and port or wireless access point the client is connected to. You can also search for a particular client session or sessions using a variety of search criteria and view client history.

NOTE: Because traffic information is unavailable for sessions connected to access ports on switches, monitors that show session traffic, such as the Top Sessions by MAC Address monitor, are not displayed for scopes that contain switches only.

The RF Tab

The RF tab provides information about the wireless environment and signal performance, allowing you to identify problems that affect wireless connectivity. Monitors provide information about throughput, retransmissions, packet errors, signal-to-noise ratio, and interference sources. Tasks enable you to determine a radio's neighbors and to display spectrograms for troubleshooting interference.

The Equipment Tab

The Equipment tab provides information about the operational status of individual devices. Monitors display CPU and memory use, power supply and fan status, port status, and general device information for switches and wireless LAN controllers. The status of access point and radios is displayed when you select their wireless LAN controller. Additional information provided by this tab includes the state of logical

Ethernet switching interfaces on standalone switches, the topology of Virtual Chassis, and the list of access points that use a selected controller as a secondary controller.

The Fabric Analysis Tab

The Fabric Analysis tab displays the results of running the Fabric Analyzer on a QFabric or Virtual Chassis Fabric (VCF). It shows information about the health, connectivity, and topology of the fabric. For information about analyzing fabrics, see [“Analyzing QFabric Devices” on page 1361](#) and [“Analyzing Virtual Chassis Fabrics” on page 1367](#).

VC Equipment Summary Tab

The VC Equipment Summary tab displays the operational status of the virtual chassis members associated with Logical View, Location View, Device View and Custom Group View.

NOTE: The VC Equipment Summary tab is enabled only for members that are at the virtual chassis container level, virtual chassis level or virtual chassis device level in the applicable view panes.

When you select any virtual chassis or virtual chassis members in the View pane, the VC Equipment Summary tab displays a pie chart that indicates the total number of virtual chassis members and the connection status of the virtual chassis members. Mouse over a pie segment to view the actual number of VC members and the percentage represented by that pie segment. The connection states are as follows:

- UP (green)—Device is connected to Network Director
- DOWN (red)—Device is not connected to Network Director

To see detailed information about the virtual chassis members, click the Details icon on the top right corner of the VC Equipment Member Status title bar. The details shown depends on the type of the virtual chassis device selected. In the VC Equipment Member Status window, the connection status is shown as green up arrow when the virtual chassis members are up and red down arrow when the virtual chassis members are down.

RELATED DOCUMENTATION

[Understanding the Monitor Mode Tasks Pane | 1277](#)

[Understanding the Network Director User Interface | 85](#)

[Network Director Documentation home page](#)

Understanding the Monitor Mode Tasks Pane

The Tasks pane in Monitor mode displays a list of tasks that are available for the currently selected Monitor tab. These tasks provide monitoring functions in addition to the monitors available under each tab.

The tasks listed in the Tasks pane vary according to the selected tab—that is, Summary, Traffic, Client, RF, or Equipment—and the scope you have selected in the View pane. For example, the VC Protocols Statistics task is available only when you select the Traffic tab and a Virtual Chassis or Virtual Chassis member in the View pane.

For each Monitor mode tab, the following tables list each task and provide a short description of the task:

- [Table 301](#): Summary Tab Tasks
- [Table 302](#): Traffic Tab Tasks
- [Table 303](#): Client Tab Tasks
- [Table 304](#): RF Tab Tasks
- [Table 305](#): Equipment Tab Tasks
- [Table 306](#): VC Equipment Summary Tab Tasks
- [Table 307](#): Device Management Tasks
- Key Tasks—Network Director enables you to group the tasks that you perform frequently and create a list of key tasks. You can add any task from the Tasks pane to the Key Tasks list by selecting a task and clicking the plus (+) sign that appears adjacent to the task. For some modes, you can see that Network Director has predefined some key tasks for you. You can modify this set of tasks to suit your requirements. This feature is available in Task pane irrespective of your current mode, scope, or view.
- For information about the Fabric Analysis tab, see [“Analyzing QFabric Devices” on page 1361](#) and [“Analyzing Virtual Chassis Fabrics” on page 1367](#).

Table 301: Summary Tab Tasks

Task	Description
Backed-Up APs	Displays information about the access points for which the selected controller is the secondary controller in a cluster.
Compare Device Statistics	Compares statistics from multiple devices in real time.
Find End Point	Finds end points that match supplied attributes and returns information about the location of the end points on the network.
Find User Session	Finds client sessions that match supplied attributes and returns information about the sessions.

Table 301: Summary Tab Tasks (continued)

Task	Description
View Tenants	Shows details of tenants and overlay networks.
View Congestion Events	Shows latency congestion information based on high-frequency statistics data.
Ping To a Host	From the selected device, pings the host you specify and returns the results.
Port Utilization	Displays port utilization trend information for the devices in the selected scope. You can view overall port utilization for the selected device or can view individual port utilization.
Refresh End Point	Refreshes end point location information for the Find End Point task.
Run Fabric Analyzer	Analyzes a QFabric device and provides information about its health, connectivity, and topology on the Fabric Analysis tab.
Select Monitors to display	Selects the monitors that are displayed in the Summary tab
Show ARP Table	Shows Address Resolution Protocol (ARP) table information for a device.
Show PoE Interfaces	Shows details such as interface name, admin status, maximum power limit, power consumption and priority, of PoE interfaces in satellite devices of Junos Fusion Enterprise and EX devices.
Show Routing Instances	Shows the routing instances configured on an MX Series router.
Traceroute To a Host	From the selected device, traces the route to the host you specify and returns the results.

Table 302: Traffic Tab Tasks

Task	Description
Compare Device Statistics	Compares statistics from multiple devices in real time.
Find End Point	Finds end points that match supplied attributes and returns information about the location of the end points on the network.
Find User Session	Finds client sessions that match supplied attributes and returns information about the sessions.
L3 VLAN Statistics	Displays packet in and out statistics for Layer 3 VLANs on the selected device.

Table 302: Traffic Tab Tasks (continued)

Task	Description
View Congestion Events	Shows latency congestion information based on high-frequency statistics data.
Ping To a Host	From the selected device, pings the host you specify and returns the results.
Port Statistics	Displays packet and error statistics for all ports on the selected device.
Port Utilization	Displays port utilization trend information for the devices in the selected scope. You can view overall port utilization for the selected device or can view individual port utilization.
Run Fabric Analyzer	Analyzes a QFabric device and provides information about its health, connectivity, and topology on the Fabric Analysis tab.
Show ARP Table	Shows Address Resolution Protocol (ARP) table information for a device.
Show PoE Interfaces	Shows details such as interface name, admin status, maximum power limit, power consumption and priority, of PoE interfaces in satellite devices of Junos Fusion Enterprise and EX devices.
Show Routing Instances	Shows the routing instances configured on an MX Series router.
Traceroute To a Host	From the selected device, traces the route to the host you specify and returns the results.
VC Protocol Statistics	Displays Virtual Chassis Control Protocol (VCCP) statistics for the selected Virtual Chassis or Virtual Chassis member, such as the kind and number of protocol data units (PDUs) sent and received.

Table 303: Client Tab Tasks

Task	Description
Compare Device Statistics	Compares statistics from multiple devices in real time.
Find End Point	Finds end points that match supplied attributes and returns information about the location of the end points on the network.
Find User Session	Finds client sessions that match supplied attributes and returns information about the sessions.
View Congestion Events	Shows latency congestion information based on high-frequency statistics data.

Table 303: Client Tab Tasks (continued)

Task	Description
Ping To a Host	From the selected device, pings the host you specify and returns the results.
Refresh End Point	Refreshes end point location information for the Find End Point task.
Run Fabric Analyzer	Analyzes a QFabric device and provides information about its health, connectivity, and topology on the Fabric Analysis tab.
Show ARP Table	Shows Address Resolution Protocol (ARP) table information for a device.
Show PoE Interfaces	Shows details such as interface name, admin status, maximum power limit, power consumption and priority, of PoE interfaces in satellite devices of Junos Fusion Enterprise and EX devices.
Show Routing Instances	Show the routing instances configured on an MX Series router.
Traceroute To a Host	From the selected device, traces the route to the host you specify and returns the results.

Table 304: RF Tab Tasks

Task	Description
Compare Device Statistics	Compares statistics from multiple devices in real time.
Find User Session	Finds client sessions that match supplied attributes and returns information about the sessions.
Interference Sources	Displays the sources of interferences detected by the selected radio.
Spectrogram	Displays a spectrum analysis of the 2.4-GHz and 5-GHz bands.
RF Neighborhood	Displays a list radios that are in the vicinity of the selected radio. These radios are either radios that the selected radio detects or radios that detect the selected radios.

Table 305: Equipment Tab Tasks

Task	Description
Backed-Up APs	Displays information about the access points for which the selected controller is the secondary controller in a cluster.

Table 305: Equipment Tab Tasks (continued)

Task	Description
Compare Device Statistics	Compares statistics from multiple devices in real time.
Find End Point	Finds end points that match supplied attributes and returns information about the location of the end points on the network.
Find User Session	Finds client sessions that match supplied attributes and returns information about the sessions.
View Congestion Events	Shows latency congestion information based on high-frequency statistics data.
Logical Interfaces	Displays the status of the Ethernet switching interfaces on the device, including aggregated Ethernet interfaces. Information includes VLAN membership, STP state, and port mode.
Ping to a Host	From the selected device, pings the host you specify and returns the results.
Run Fabric Analyzer	Analyzes a QFabric device and provides information about its health, connectivity, and topology on the Fabric Analysis tab.
Show ARP Table	Shows Address Resolution Protocol (ARP) table information for a device.
Show PoE Interfaces	Shows details such as interface name, admin status, maximum power limit, power consumption and priority, of PoE interfaces in satellite devices of Junos Fusion Enterprise and EX devices.
Show Routing Instances	Show the routing instances configured on an MX Series router.
Traceroute To a Host	From the selected device, traces the route to the host you specify and returns the results.

Table 306: VC Equipment Summary Tab Tasks

Task	Description
Find User Session	Finds client sessions that match supplied attributes and returns information about the sessions.
Compare Device Statistics	Compares statistics from multiple devices in real time.

Table 307: Device Management Tasks in Monitor Mode

Task	Description
View Aruba Wireless Device Inventory	Displays the list of Aruba devices connected to Juniper Networks switches.
Launch Aruba Airwave	Launches the APs/Devices > List page of the Aruba Airwave application to monitor Aruba devices.

RELATED DOCUMENTATION

[Understanding Monitor Mode in Network Director | 1269](#)

[Understanding the Network Director User Interface | 85](#)

[Network Director Documentation home page](#)

Monitoring Traffic

IN THIS CHAPTER

- [Monitoring Traffic on Devices | 1283](#)
- [Monitoring Port Traffic Statistics | 1284](#)
- [Monitoring Traffic on Layer 3 VLANs | 1287](#)
- [Monitoring Routing Instances | 1289](#)
- [Monitoring Port Utilization | 1300](#)
- [Monitoring Tenant Details | 1304](#)
- [Monitoring Virtual Chassis Protocol Statistics | 1310](#)
- [Viewing Congestion Events | 1312](#)

Monitoring Traffic on Devices

The monitors on the Traffic tab provide information about the traffic traversing QFabric systems, switches, routers, Virtual Chassis, Virtual Chassis Fabrics (VCFs), Layer 3 Fabrics, and wireless controllers.

To monitor traffic on a device:

1. Click **Monitor** in the Network Director banner.
2. Select the device in the View pane that contains the traffic you want to monitor.
3. Select the **Traffic** tab to open the traffic monitors.
4. To get help for a monitor, click the Help button in its title bar.

The available monitors include:

- [“Unicast vs Broadcast/Multicast Monitor” on page 1440](#): shows the distribution of unicast, broadcast, and multicast traffic entering and leaving the device.
- [“Unicast vs Broadcast/Multicast Trend Monitor” on page 1441](#): shows trend data about the distribution of unicast, broadcast, and multicast traffic entering and leaving the device.

- [“Traffic Trend Monitor” on page 1439](#): shows trend data about the amount of traffic entering and leaving the device.
- [“Error Trend Monitor” on page 1391](#): shows trend data about the amount of errors on the device.
- [“Top APs by Traffic Monitor” on page 1432](#) shows the APs that are handling the most traffic.
- [“Top Talker - Wired Devices Monitor” on page 1435](#) shows the hosts that are using the most bandwidth.

RELATED DOCUMENTATION

[Understanding Monitor Mode in Network Director | 1269](#)

[Monitoring Port Traffic Statistics | 1284](#)

[Monitoring Virtual Chassis Protocol Statistics | 1310](#)

[Monitoring Traffic on Layer 3 VLANs | 1287](#)

[Network Director Documentation home page](#)

Monitoring Port Traffic Statistics

IN THIS SECTION

- [Procedure for Monitoring Port Traffic Statistics | 1284](#)
- [Port on Device Window | 1285](#)
- [Port Traffic Stats Window | 1286](#)

This topic describes how to monitor port traffic statistics on a device. You can monitor port traffic statistics for a switch, router, Virtual Chassis, wireless controller, Virtual Chassis Fabric, QFabric system or Layer 3 Fabric.

This topic describes:

Procedure for Monitoring Port Traffic Statistics

1. Click **Monitor** in the Network Director banner.

2. Do one of the following:

To view the port statistics:

- a. Select a node in the View pane that contains the port traffic you want to monitor.
- b. Select the **Traffic** tab.
- c. In the Tasks pane, select **Port Statistics**.

The Port Traffic Stats window opens. For information about this window, click the Help button in the title bar of the window or see [“Port Traffic Stats Window” on page 1286](#).

To view the port status and analyze network traffic for devices that are configured for network traffic analysis:

- a. Select a device from the view pane.
- b. In the Tasks pane, select **Traffic Analysis**.

The Port on Device window opens. For information about this window and the tasks that you can perform from this window, see [“Port on Device Window” on page 1285](#).

Port on Device Window

Port on Device window displays the details of all the ports on a device. [Table 308](#) describes the fields that are displayed in the Port on Device window.

Table 308: Port on Device table field descriptions

Field Name	Description
Port Name	Identification of the port.
Admin State	The administrative state of the port: enabled (UP) or disabled (DOWN).
Operational State	The operational status—link up (UP) or link down (DOWN).
Max Bandwidth	The actual bandwidth available on the port, in megabits (Mb).
Negotiated Bandwidth	The negotiated bandwidth based on the speed that is configured or auto-negotiated for the interface.

To view more details about the traffic on any port, select the port and click View Traffic. The Traffic on Port window opens. For more details and field descriptions, see [“Device & Port Utilization Widget” on page 151](#).

Port Traffic Stats Window

The Port Traffic Stats window displays information about the port traffic on the node you selected in the View pane. It contains the following elements:

- **Port Traffic Trend graph**—This line graph shows trends in the data and error rates on the port selected in the ports table below it. The horizontal axis shows the times when samples were taken. The vertical axis shows the data rate on the left side (in packets per second) and the error rate on the right side (in errors per second).

To display traffic for a different port, select the port from the table below the graph. To change the time period over which to display the traffic trends, select a time period from the list in the upper right corner.

NOTE: After a device is discovered, trend data does not appear immediately—it appears only after three polling periods have elapsed.

To highlight a line on the graph, mouse over the line legend. To remove or restore a line, click the line legend. To see numeric values, mouse over where a data line intersects with a dotted vertical grid line.

- **Ports table** (on the lower left side of the window)—This table provides information about the ports as described in [Table 309](#). Selecting a port from this table updates the Port Traffic Trend graph to display traffic information about the selected port.
- **Counter selection table** (on the lower right side of the window)—This table enables you to select which counters to display on the Port Traffic Trend graph. It includes separate tabs for packet counters and error counters. Select the check box in the Show column of each counter that you want to display on the graph. The Per/Sec column shows the rate per second of that row’s counter.

Table 309: Port Traffic Window

Table Column	Description
Serial Num	Serial number of the device to which the port belongs.
Port Name	Port number. Channelized ports are indicated by the port number followed by <i>:<channelized port number></i> . For example, xe/0/0/1:2 indicates that this channelized port is a part of the xe/0/0/1 port with a channelized port number of 2.
Port Usage Type	Port mode—either ACCESS or UPLINK.

Table 309: Port Traffic Window (continued)

Table Column	Description
MAC Addresses	Port MAC address.
Link Type	Full duplex, half duplex, or unspecified.
In Packets/Sec.(Current)	Current rate of inbound packets.
Out Packets/Sec.(Current)	Current rate of outbound packets.

RELATED DOCUMENTATION

[Understanding the Monitor Mode Tasks Pane | 1277](#)

[Network Director Documentation home page](#)

Monitoring Traffic on Layer 3 VLANs

IN THIS SECTION

- [Procedure for Monitoring Layer 3 VLAN Traffic Statistics | 1287](#)
- [L3 VLAN Traffic Stats Window | 1288](#)

This topic describes how to monitor Layer 3 VLAN traffic statistics on a device. You can monitor Layer 3 VLAN statistics for a switch, router, Virtual Chassis, Virtual Chassis Fabric, Layer 3 Fabric, and the aggregation devices in a Junos Fusion fabric.

This topic describes:

Procedure for Monitoring Layer 3 VLAN Traffic Statistics

1. Click **Monitor** in the Network Director banner.
2. Select a node in the View pane that contains the Layer 3 VLAN traffic you want to monitor.

3. Select the **Traffic** tab.
4. In the Tasks pane, select **View > L3 VLAN Statistics**.

The L3 VLAN Traffic Stats window opens. For information about this window, click the Help button in the title bar of the window or see [“L3 VLAN Traffic Stats Window” on page 1288](#).

L3 VLAN Traffic Stats Window

The L3 VLAN Traffic Stats window displays information about the Layer 3 VLAN traffic on the node you selected in the View pane. You can monitor Layer 3 VLAN statistics for a switch, router, Virtual Chassis, Virtual Chassis Fabric, Layer 3 Fabric, and the aggregation devices in a Junos Fusion fabric.

NOTE: For a Junos Fusion fabric, you must select an aggregation device to view the Layer 3 VLAN statistics. The L3 VLAN Statistics option is not available if you select a Junos Fusion satellite device.

The L3 VLAN Traffic Stats window contains two panes:

- **VLAN Traffic line graph**—This graph shows the data transmission rate on the Layer 3 VLAN selected in the table beneath the graph. The horizontal axis shows the times when samples were taken. The vertical axis shows the data rate, in bytes per second.

To show a Layer 3 VLAN on the VLAN Traffic line graph, select the Layer 3 VLAN from the table beneath the graph. To highlight a line on the graph, mouse over the line legend. To remove or restore a line, click the line legend. To see numeric values, mouse over a data point.

- **Layer 3 VLAN traffic statistics table**—This table provides information about the Layer 3 VLANs as described in [Table 310](#). Selecting a Layer 3 VLAN from this table updates the VLAN Traffic graph to display the traffic information for the selected Layer 3 VLAN.

Table 310: Layer 3 VLAN Traffic Statistics Table

Table Column	Description
L3 Interface	Layer 3 interface assigned to the VLAN.
SerialNo	The serial number of the device containing the Layer 3 VLAN.
VLAN Name	VLAN name.
VLAN ID	VLAN ID.
Description	VLAN description.

Table 310: Layer 3 VLAN Traffic Statistics Table (continued)

Table Column	Description
In Packet	Number of packets entering the VLAN.
Out Packet	Number of packets leaving the VLAN.

RELATED DOCUMENTATION

[Understanding the Monitor Mode Tasks Pane | 1277](#)

[Network Director Documentation home page](#)

Monitoring Routing Instances

IN THIS SECTION

- [Procedure for Monitoring Routing Instances | 1290](#)
- [Show Routing Instances Window | 1290](#)
- [Show Interfaces Window | 1291](#)
- [Show Bridge Domains Window | 1292](#)
- [Show Connections | 1293](#)
- [Show Routing Tables | 1296](#)
- [Show MAC Table | 1299](#)

This topic describes how to monitor VPN routing instances on MX Series routers by using Network Director. Using Network Director, you can determine which interfaces and bridge domains belong to the routing instances and view traffic statistics for those interfaces and bridge domains. You can also display connection information for Layer 2 VPN and virtual private LAN service (VPLS) routing instances.

Network Director can be used to monitor the following types of Layer 2 routing instances:

- Default routing instance
- Ethernet VPN (EVPN)
- Layer 2 VPN

- VPLS
- Virtual switch

Network Director can be used to monitor the following types of Layer 3 routing instances:

- Layer 3 VPN

This topic describes:

Procedure for Monitoring Routing Instances

Use the options in the Show Routing Instances window to monitor routing instances.

1. Click **Monitor** in the Network Director banner.
2. Select an MX Series router in the View pane that contains the port traffic you want to monitor.
3. In the Tasks pane, select **Tasks > Show Routing Instances**.

The Show Routing Instances window opens. For information about this window, click the Help button in the title bar of the window or see [“Show Routing Instances Window” on page 1290](#).

Show Routing Instances Window

The Show Routing Instances window lists the routing instances configured on a selected device. Use this window to display the interfaces or bridge domains belonging to a routing instance and obtain traffic statistics for the interfaces. You can also display information about the VPLS and Layer 2 VPN connections. [Table 311](#) describes the fields in this window.

Table 311: Fields in the Show Routing Instances Window

Field	Description
Routing Instance Name	Name of the routing instance. The default routing instance is named <i>default-switch</i> .



Table 311: Fields in the Show Routing Instances Window (continued)

Field	Description
Type	<p>Identifies the routing instance type:</p> <ul style="list-style-type: none"> • EVPN • L2VPN • L3VPN • Virtual Switch <p>The default routing instance is of this type.</p> <ul style="list-style-type: none"> • VPLS • VRF (L3VPN)
Details	<p>Provides the following information (if configured for the routing instance):</p> <ul style="list-style-type: none"> • Route Distinguisher—Used to identify all routes that are part of the VPN. The route distinguisher makes IP addresses globally unique, so that the same IP address prefixes can be used for different VPNs. • Target—Extended BGP community used to match routes for import and export.
Interfaces	<p>Displays the number of interfaces belonging to the routing instance. Click the number to open the Show Interfaces window, described in “Show Interfaces Window” on page 1291.</p>
Bridge Domains	<p>Displays the number of bridge domains belonging to the routing instance. Click the number to open the Show Bridged Domains window, described in “Show Bridge Domains Window” on page 1292.</p>
Actions	<ul style="list-style-type: none"> • Click Show Connections to display information about Layer 2 VPN and VPLS connections. The information described in “Show Connections” on page 1293 is displayed. This link is available only for Layer 2 VPN and VPLS routing instances. • Click Show MAC Table to display the MAC table for the selected routing instance. For details, see “Show MAC Table” on page 1299. • Click Show Routing Table to view the routing table information for the selected routing instance. For details, see “Show Routing Tables” on page 1296.

Show Interfaces Window

The Show Interfaces window lists the logical interfaces configured on the routing instance and provides the information about the interfaces as described in [Table 312](#).

Table 312: Show Interfaces Information

Field	Description
Interface Name	The interface name.
Port Mode	Indicates one of two modes—Access or Trunk: <ul style="list-style-type: none"> • Access—The interface can be in a single VLAN only. • Trunk—The interface can be in multiple VLANs and accept tagged packets from multiple devices.
Interface State	Indicates whether the interface is  UP or  DOWN.
STP State	Indicates whether the interface is in a discarding (blocked) or in forwarding (unblocked) state. (Not shown for interfaces belonging to Layer 2 VPN and Layer 3 VPN routing instances.)
Local IP Address	Local IP address. (Shown only for interfaces belonging to Layer 2 VPN and Layer 3 VPN routing instances.)
Remote IP Address	Remote IP address. (Shown only for interfaces belonging to Layer 2 VPN and Layer 3 VPN routing instances.)
Actions	<ul style="list-style-type: none"> • Click View Statistics to display traffic statistics for the interface. The Show Interface Statistics window opens, which charts the number of input and output packets and the number of input and output bytes. • Click Show MAC Table to display the MAC table for the interface. For more details, see “Show MAC Table” on page 1299.



Show Bridge Domains Window

The Show Bridge Domains window lists the bridge domains configured on the routing instance. To display information about the VLAN IDs and interfaces configured on a bridge domain, select the bridge domain. [Table 313](#) describes the information provided in the Show Bridge Domains window.

Table 313: Show Bridge Domains Information

Field	Description
Bridge Domains	The bridge domain name.
Actions	Click Show MAC Table to display the MAC table for the selected bridge domain. For details, see “Show MAC Table” on page 1299 .

Table 313: Show Bridge Domains Information (continued)

Field	Description
VLAN ID	The VLAN ID or IDs assigned to the bridge domain.
Interface Name	The name of a logical interface assigned to the VLAN ID.
Port Mode	Indicates one of two modes—access or trunk: <ul style="list-style-type: none"> • Access—The interface can be in a single VLAN only. • Trunk—The interface can be in multiple VLANs and accept tagged packets from multiple devices.
Interface State	Indicates whether the interface is  UP or  DOWN .
STP State	Indicates whether the interface is in a discarding (blocked) or in forwarding (unblocked) state.
Actions	<ul style="list-style-type: none"> • Click View Statistics to display traffic statistics for the interface. The Show Interface Statistics window opens, which charts the number of input and output packets and the number of input and output bytes. • Click Show MAC Table to display the MAC table for the interface. For details, see “Show MAC Table” on page 1299.

Show Connections

The Show Connections window provides information about the VPN connections for Layer 2 VPN and VPLS routing instances as described in [Table 314](#).

Table 314: Show Connections Information



Field	Description
Local Site Name	Name of the local site.
Local Site ID	Identifier for the local site.
Local Interface Name	Name of the local interface.
Interface Status	Indicates whether the local interface is  UP or  DOWN .
Remote Site ID	Identifier for the remote site.
Remote IP	IP address of the remote provider edge device (PE device).

Table 314: Show Connections Information (*continued*)

Field	Description
Connection Status	

Table 314: Show Connections Information (continued)

Field	Description
	<p>Status of the connection:</p> <ul style="list-style-type: none"> ● EI—The local VPN interface is configured with an encapsulation that is not supported. ● EM—The encapsulation type received on this connection from the neighbor does not match the local connection interface encapsulation type. ● VC-Dn—The virtual circuit is currently down. ● CM—The two routers do not agree on a control word, which causes a control word mismatch. ● CN—The virtual circuit is not provisioned properly. ● OR—The label associated with the virtual circuit is out of range. ● OL—No advertisement has been received for this virtual circuit from the neighbor. There is no outgoing label available for use by this virtual circuit. ● LD—All of the CE-facing interfaces to the local site are down. Therefore, the connection to the local site is signaled as down to the other PE routers. No pseudowires can be established. ● RD—All the interfaces to the remote neighbor are down. Therefore, the remote site has been signaled as down to the other PE routers. No pseudowires can be established. ● LN—The local site has lost path selection to the remote site and therefore no pseudowires can be established from this local site. ● RN—The remote site has lost path selection to a local site or to a remote site and therefore no pseudowires are established to this remote site. <p>In a multihoming configuration, one multihomed PE site displays the state LN, and the other multihomed PE site displays the state RN in the following circumstances:</p> <ul style="list-style-type: none"> ● The multihomed links are both configured to be the backup site. ● The two multihomed PE routers have the same site ID, but have a peering relationship with a route reflector (RR) that has a different site ID. <ul style="list-style-type: none"> ● XX—The connection is down for an unknown reason. This is a programming error. ● MM—The MTUs for the local site and the remote site do not match. ● BK—The router is using a backup connection. ● PF—Profile parse failure. ● RS—The remote site is in a standby state. ● NC—The interface encapsulation is not configured as an appropriate CCC (circuit cross-connect), TCC (translational cross-connect), Layer 2 VPN, or VPLS encapsulation. ● WE—The encapsulation configured for the interface does not match with the encapsulation configured for the associated connection within the routing instance. ● NP—The router detects that interface hardware is not present. The hardware might be offline, a PIC might not be of the compatible type, or the interface might be configured in a different routing instance. ● ->—Only the outbound connection is up.

Table 314: Show Connections Information (*continued*)

Field	Description
	<ul style="list-style-type: none"> ● <--Only the inbound connection is up. ● Up—The connection is operational. ● Dn—The connection is down. ● CF—The router cannot find enough bandwidth to the remote router to meet the connection bandwidth requirement. ● SC—The local site identifier is the same as the remote site identifier. No pseudowire can be established between these two sites. You must configure different values for the local and remote site identifiers. ● LM—The local site identifier is not the minimum designated, which means it is not of the lowest value. There is another local site with a lower value for site identifier. Pseudowires are not being established to this local site and the associated local site identifier is not being used to distribute Layer 2 VPN or VPLS label blocks. However, this is not an error state. Traffic continues to be forwarded to the PE router interfaces connected to the local sites when the local sites are in this state. ● RM—The remote site identifier is not the minimum designated, which means it is not the lowest. There is another remote site connected to the same PE router which has lower site identifier. The PE router cannot establish a pseudowire to this remote site and the associated remote site identifier cannot be used to distribute VPLS label blocks. However, this is not an error state. Traffic continues to be forwarded to the PE router interface connected to this remote site when the remote site is in this state. ● IL—The incoming packets for the connection have no MPLS label. ● MI—The configured mesh group identifier is in use by another system in the network. ● ST—The router has switched to a standby connection. ● PB—Profile is busy. ● SN—The neighbor is static.
Time Last Up	The time when the connection was last in the Up condition.

Show Routing Tables

The Routing Tables window enables you view the routing table information for the selected virtual routing instance. For L3VPN and EVP services, you can determine which LSPs or tunnels are being used by looking at the routing tables.

- **Routing Tables**—The Routing Tables table shows the routing tables associated with the virtual instance and the number of active routes in each table. Click on a routing table to display the actual contents of the routing table.

- Details—The Details table shows the contents of the selected routing table. [Table 315](#) displays the fields that are displayed in the Details table.

Table 315: Show Routing Table Field Descriptions

Name	Description
Routing Instance	Name of the routing instance.
Number of Destinations	Number of destinations for which there are routes in the routing table.
Active Routes	Number of routes that are active.
Hidden Routes	Number of routes that are not used because of routing policy.
Hold-down Routes	Number of routes that are in the hold-down state before being declared inactive.
Total Routes	Total number of routes.
Destination Prefix	<p>Route destination (for example:10.0.0.1/24). Sometimes the route information is presented in another format, such as:</p> <ul style="list-style-type: none"> • MPLS-label (for example, 80001). • interface-name (for example, ge-1/0/2). • neighbor-address:control-word-status:encapsulation type:vc-id:source (Layer 2 circuit only. For example, 10.1.1.195:NoCtrlWord:1:1:Local/96): <ul style="list-style-type: none"> • neighbor-address—Address of the neighbor. • control-word-status—Whether the use of the control word has been negotiated for this virtual circuit: NoCtrlWord or CtrlWord. • encapsulation type—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport. • vc-id—Virtual circuit identifier. • source—Source of the advertisement: Local or Remote.
State	State of the route.
Protocol	Name of the protocol from which the route was learned. For example, OSPF , RSVP , and Static .
Protocol Preference	Preferred protocol for this routing instance. Junos OS uses this preference to choose which routes become active in the routing table.
Age	Displays how long since the route was learned.

Table 315: Show Routing Table Field Descriptions (*continued*)

Name	Description
Metric	Cost value of the indicated route. For routes within an AS, the cost is determined by the IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
BGP Local Preference	A metric used by BGP sessions to indicate the degree of preference for an external route. The route with the highest local preference value is preferred.
Route Learned From	Interface from which the route was received.
AS Path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: <ul style="list-style-type: none"> • I—IGP • E—EGP • ?—Incomplete; typically, the AS path was aggregated.
Validation State	(BGP-learned routes) Validation status of the route: <ul style="list-style-type: none"> • Invalid—Indicates that the prefix is found, but either the corresponding AS received from the EBGp peer is not the AS that appears in the database, or the prefix length in the BGP update message exceeds the maximum length permitted in the database. • Unknown—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • Valid—Indicates that the prefix and autonomous system pair are found in the database.
Next Hop Type	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route. If the destination is Discard, traffic is dropped.
Local Interface	The local interface used to reach the next hop.
Address	IP address of the interface.
Via Interface	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word Selected.
MPLS Label	MPLS label and operation occurring at the next hop. The operation can be pop (where a label is removed from the top of the stack), push (where another label is added to the label stack), or swap (where a label is replaced by another label).

Show MAC Table

The Show MAC table window displays the MAC table for the selected routing instance. [Table 316](#) describes the fields that are displayed in the Show MAC Table window.

Table 316: Show MAC Table Field Descriptions

Field Name	Description
Routing Instance	Name of the routing instance.
Type	Identifies the routing instance type: <ul style="list-style-type: none"> • EVPN • L2VPN • L3VPN • Virtual Switch <p>The default routing instance is of this type.</p> <ul style="list-style-type: none"> • VPLS • VRF (L3VPN)
Bridge Domain	Name of the bridging domain.
VLAN ID	VLAN ID of the routing instance or bridge domain in which the MAC address was learned.
MAC Address	MAC address or addresses learned on a logical interface.
MAC Flags	Status of MAC address learning properties for each interface: <ul style="list-style-type: none"> • S—Static MAC address is configured. • D—Dynamic MAC address is configured. • L—Locally learned MAC address is configured. • C—Control MAC address is configured. • SE—MAC accounting is enabled. • NM—Non-configured MAC. • R—Remote PE MAC address is configured.
Logical Interface	Name of the logical interface.

RELATED DOCUMENTATION

[Understanding the Monitor Mode Tasks Pane](#) | 1277

Monitoring Port Utilization

IN THIS SECTION

- [How to Access the Port Utilization Task | 1300](#)
- [Port Utilization Details Window | 1301](#)
- [Utilization for Device Window | 1301](#)
- [Utilization for Fabric Devices Window | 1303](#)

Network Director provides information about port utilization in either one of two places, depending on the node you select in the View pane:

- **Port Utilization monitor**—This monitor, available in the Summary tab, provides a bar chart that shows the aggregate utilization of the ports on a device or devices over a period of time that you select. For more information about using the Port Utilization monitor, see [“Port Utilization Monitor” on page 1399](#).
- **Port Utilization task**—This task, available from **View > Port Utilization** in the Tasks pane of the Summary or Traffic tabs, provides a bar chart similar to the Port Utilization monitor bar chart. Unlike the Port Utilization monitor, it also enables you to obtain information on individual port utilization over time when you have selected an individual device or Layer 3 Fabric in the View pane.

This topic describes the Port Utilization task. It describes:

How to Access the Port Utilization Task

1. Click **Monitor** in the Network Director banner.
2. Select a node in the View pane that contains the ports whose utilization you want to monitor.
3. Select the **Summary** or **Traffic** tab.
4. In the Tasks pane, select **View > Port Utilization**.

If you have selected a node that contains more than one device, the Port Utilization Details window opens. For information about this window, see [“Port Utilization Details Window” on page 1301](#).

If you have selected an individual device, the Utilization for Device window opens. For information about this window, see [“Utilization for Device Window” on page 1301](#).

If you have selected a fabric device such as Junos Fusion, Layer 3 Fabric, Virtual Chassis Fabric, or a QFabric device, the Utilization for Fabric Device window opens. For information about this window, see [“Utilization for Fabric Devices Window” on page 1303](#).

Port Utilization Details Window

This window provides a bar chart showing the aggregate port utilization trend for the devices within the selected scope.

Each bar in the bar chart represents the overall port utilization for all the devices at a polling interval. The vertical axis shows the number of ports polled. The horizontal axis shows the time when each poll was taken.

Each bar is divided into the following colored sections to indicate the distribution of port utilization at the polling interval:

- Green indicates ports that operated at less than 50% of negotiated speed.
- Yellow indicates ports that operated at between 50% and 80% of negotiated speed.
- Red indicates ports that operated at more than 80% of negotiated speed.

You can perform the following actions on the bar chart:

- Change the time period over which to display the trend data by selecting a time period from the list in the upper right corner.
- Remove or restore a utilization category (bar color) by clicking its legend.
- Display a numeric value by mousing over a bar.

Utilization for Device Window

The Utilization for Device window shows the port utilization trend for individual devices and ports. It is available when you select a individual device in the View pane.

The Utilization for Device window provides two views of port utilization:

- Device—This view provides a trend chart of overall port use on the device over time.
- Port—This view provides a heat map of all the ports on the device, enabling you to view the utilization of individual ports. You can choose a port from the heat map to view a utilization trend chart for that particular port.

The Device view is the default view. Click **Port** to change to the Port view.

Device View

The Device view provides a bar chart that shows the trend of overall port use on the device. Each bar represents the overall port utilization at a polling interval. The vertical axis shows the number of ports polled. The horizontal axis shows the time when each poll was taken.

Each bar is divided into the following colored sections to indicate the distribution of port utilization at the polling interval:

- Green indicates ports that operated at less than 50% of negotiated speed.
- Yellow indicates ports that operated at between 50% and 80% of negotiated speed.
- Red indicates ports that operated at more than 80% of negotiated speed.

You can perform the following actions in Device view:

- Change the time period over which to display the trend data by selecting a time period from the list in the upper right corner.
- Remove or restore a utilization category (bar color) by clicking its legend.
- Display a numeric value by mousing over a bar.

Port View

The Port view provides utilization heat maps of the ports on the device—one heat map for access ports and another for uplink ports. In the heat maps, each port on the device is represented by a box that is color-coded to indicate the level of port utilization. Cooler colors (for example, green) indicate lower port utilization, while hotter colors (for example, red) indicate higher port utilization.

Click a port box to display a utilization trend chart for that individual port.

You can perform the following actions in the Port view:

- On a heat map:
 - Mouse over a port box to see more information about the port such as the port utilization percent, port type, MAC address of the port, duplex mode, device serial number, admin status and operational status of the port, negotiated speed, and the last flap time.
 - Change the time period over which the port utilization percentage is derived.
 - Click a port box to display the utilization trend chart for that port.
 - Use the percentage slider under the port heat map to display only those ports for which utilization falls within a certain percentage range.
- On the port utilization trend chart:
 - Change the time period over which to display the trend data.
 - Display the percentage utilization and polling time by mousing over a data point.

Utilization for Fabric Devices Window

The Utilization for Fabric Devices window provides information about port utilization for the devices and ports within a fabric device such as a Junos Fusion, Layer 3 Fabric, Virtual Chassis Fabric, or a QFabric device. It is available when you select a fabric device from the Fabrics container in the View pane.

The top part of the Utilization for Fabric Devices window displays a heat map of the devices in the fabric. Each device in the fabric is shown as either a spine or leaf device (for Layer 3 Fabric or Virtual Chassis Fabric) or aggregation device and satellite device (for Junos Fusion devices) and is color-coded to show the overall port utilization on the device.

You can interact with this fabric-level heat map as follows:

- Mouse over a box representing a device. Information about that device is displayed, such as IP address, model, overall port utilization, and a list of the five ports with the highest utilization.
- Click a box representing a device. The information in the remainder of the window is changed to reflect the port utilization of the selected device. For example, for a Junos Fusion device, clicking an aggregation device displays the heat map for the cascade ports and the uplink ports whereas clicking the satellite device displays the heat map for the extended ports.

You can select two different views of the port utilization on the device:

- **Device**—This view provides a trend chart of overall port use on the device over time.
- **Port**—This view provides a heat map of all the ports on the device, enabling you to view the utilization of individual ports. You can choose a port from the heat map to view a utilization trend chart for that particular port.

The Device view is the default view. Click **Port** to change to the Port view.

Device View

The Device view provides a bar chart that shows the trend of overall port use on the selected device. Each bar represents the overall port utilization at a polling interval. The vertical axis shows the number of ports polled. The horizontal axis shows the time when each poll was taken.

Each bar is divided into the following colored sections to indicate the distribution of port utilization at the polling interval:

- Green indicates ports that operated at less than 50% of negotiated speed.
- Yellow indicates ports that operated at between 50% and 80% of negotiated speed.
- Red indicates ports that operated at more than 80% of negotiated speed.

You can perform the following actions on the bar chart:

- Change the time period over which to display the trend data by selecting a time period from the list in the upper right corner.

- Remove or restore a utilization category (bar color) by clicking its legend.
- Display a numeric value by mousing over a bar.

Port View

The Port view provides utilization heat maps of the ports on the device—one heat map for access ports and another for uplink ports. In the heat maps, each port on the device is represented by a box that is color-coded to indicate the level of port utilization. Cooler colors (for example, green) indicate lower port utilization, while hotter colors (for example, red) indicate higher port utilization.

You can perform the following actions on the device heat map:

- Mouse over a port box to see more information about the port, such as the port utilization percent, port type, MAC address of the port, duplex mode, device serial number, admin status and operational status of the port, negotiated speed, and the last flap time.
- Change the time period over which the port utilization percentage is derived.
- Use the percentage slider under the port heat map to display only those ports whose percent utilization falls within a certain range.
- Click a port box to display the utilization trend chart for that port.

The port utilization trend chart shows the utilization trend for the selected port. You can:

- Change the time period over which to display the trend data.
- Display the percentage utilization and polling time by mousing over a data point.

RELATED DOCUMENTATION

[Port Utilization Monitor | 1399](#)

[Network Director Documentation home page](#)

Monitoring Tenant Details

IN THIS SECTION

- [Viewing the List of Tenants | 1306](#)
- [View Port Details of Tenants | 1307](#)
- [View Endpoints | 1307](#)
- [View the Port Utilization Trend for a VXLAN Port | 1309](#)

This topic describes how to monitor details about the tenants that are part of your overlay network. You can create overlay networks and tenants in Network Director by using Layer 3 Fabrics that are created and managed from Network Director. You can monitor the tenant details for the entire network, for a specific Layer 3 Fabric that acts as the underlay for the tenant overlay network, or for a data center that uses a Layer 3 Fabric.

To view the tenant details:

1. Click **Monitor** in the Network Director banner.
2. Select the **Summary** tab.
3. For the Monitor life cycle mode, select a combination of view and a device or container as shown in [Table 317](#).

Table 317: Scopes that You Can Use to View Tenant Details

View Selector	Selection from the View Pane
Logical or Device	My Network
Logical or Device	My Network > Fabric > Layer 3 Fabric My Network > Fabrics > Layer 3 Fabric > Spine Device My Network > Fabrics > Layer 3 Fabric > Leaf Device
Logical or Device	My Datacenters > Data center My Datacenters > Data center > Layer 3 Fabric My Datacenters > Data center > Layer 3 Fabric > Spine Device My Datacenters > Data center > Layer 3 Fabric > Leaf Device

From the Tasks pane, select **Tasks > View > View Tenants**.


The View Tenants window opens, displaying a list of tenants.

NOTE: The View Tenant task is enabled for devices that are at the virtual chassis level.

You can use the filters available on this window to sort ports by **Port Utilization** percentage or **Port Status**.

Use the breadcrumbs at the top of the window to navigate to the various views within the View Tenant window. For example, in [Figure 54](#), from the Endpoint view, you can click **All Ports** to view details of all the ports that Tenant 1 uses, click **Tenant 1** to view the port details summary for Tenant 1, or click **All Tenants** to view the details of all the tenants in the View Tenants window.

Figure 54: Breadcrumbs on the View Tenants window



You can perform the following tasks from the View Tenants window:

Viewing the List of Tenants

The View Tenants window enables you to view details about the tenants, the number of ports used by a tenant, and the status of the ports. The level of information that Network Director displays in this window depends on the scope that you select. If you select a Layer 3 Fabric, this window displays the tenants that are part of that fabric, whereas if you select from My Network, this window displays all the tenants that are part of the network—tenants from multiple data centers and fabrics.

The View Tenant details table displays the details of all the tenants and their port status and utilization. See [Table 318](#) for a description of the fields in this table.

Table 318: View Tenant Details Table Field Descriptions

Field	Description
Tenants	Name of the tenant.
VXLAN ID	VXLAN ID of the overlay networks for the tenant.
Total Ports	Number of ports that the tenant uses.
Number of Ports with Status	Displays the number of ports that are up and ports that are down in separate columns.
Ports with Utilization (%)	Displays the number of ports that utilize high, medium, and low bandwidth in separate columns.

View Port Details of Tenants

To view more details about the ports that are used by a tenant:

1. In the View Tenants window, click the number of ports field corresponding to a tenant for which you want to view more details.

The Port Details view opens.

[Table 319](#) describes the fields of the Port Details view.

NOTE: You can click the number of ports in any of fields—Total Ports, Number of Ports with Status, or Ports with Utilization (%)—to open the Ports Details view. Network Director filters the ports based on the column that you clicked. For example, if you click the port number in the Total Ports column, the Port Details window displays all the ports that the tenant uses; If you click the port number in the Ports with Utilization (%) > High, the Port Details view filters and displays only the ports that have high utilization.

Table 319: Port Details View Field Descriptions

Field	Description
Device	The device on which the given port exists.
Port	The port number.
Port Status	Indicates whether the port is up or down.
Port Utilization %	Utilization (percentage) of the selected port.
Actions	Click to view details about the VXLAN endpoints. The View All Endpoints window opens.

View Endpoints


Endpoints are the hosts on which a tenant network terminates. You can view the endpoint details for each VXLAN overlay network.

To view endpoint details:

From the Port Details view, click **Show Endpoints** in the Actions column to view the end point details for a tenant on a specific device and port. The Show Endpoints window opens displaying the port number and the corresponding MAC address on a host.

View the Port Utilization Trend for a VXLAN Port

To view the port utilization trend for a port:


1. From the Port Details view, click  in the Port Utilization % column to view the utilization trend for the port.

The Port Utilization Trend window is divided into three sections—packets details graph, port details, packet counter and error count table.

The packet details graph displays the number of packets that passed through the port per second plotted during a certain time period as shown in [Figure 55](#). The vertical axis on the left shows the number of packets per second. The horizontal axis shows the time when the packet count was taken. The dotted lines indicate the number of errors. You can read the error count for any time period by using the vertical axis on the right.

NOTE: The default polling interval for the packet details graph is 1 hour. You can modify this by selecting an appropriate value from the polling interval drop down list that is displayed above the packet details graph.

Figure 55: Packet Details Graph



The port details section displays the port name, port usage, MAC address of the port, and the number of packets that passed through the port.

The packet counter and error count table consists of two sub-tabs. The Packet Counter sub-tab displays the distribution of unicast, broadcast, and multicast traffic types on the port you selected. The port traffic is classified as *Unicast In*, *Broadcast In*, *Multicast In*, *Unicast Out*, *Broadcast Out*, and *Multicast Out*.

The Error Count sub-tab displays the count of packet errors under each major error category.

You can select or clear each counter to view or hide details about a specific counter in the packet details graph.

RELATED DOCUMENTATION

[Understanding the Monitor Mode Tasks Pane | 1277](#)

[VXLAN–EVPN Overlay Overview | 828](#)

[Network Director Documentation home page](#)

Monitoring Virtual Chassis Protocol Statistics

IN THIS SECTION

- [Procedure for Monitoring Virtual Chassis Protocol Statistics | 1310](#)
- [Virtual Chassis Protocol Statistics Window | 1310](#)

This topic describes how to monitor Virtual Chassis protocol statistics on a device. You can monitor Virtual Chassis protocol statistics for a Virtual Chassis node in any view.

This topic describes:

Procedure for Monitoring Virtual Chassis Protocol Statistics

1. Click **Monitor** in the Network Director banner.
2. Select a node in the View pane that contains the Virtual Chassis protocol traffic you want to monitor.
3. Select the **Traffic** tab.
4. In the Tasks pane, select **View > VC Protocol Statistics**.

The Virtual Chassis Protocol Statistics window opens. For information about this window, click the Help button in the title bar of the window or see [“Virtual Chassis Protocol Statistics Window” on page 1310](#).

Virtual Chassis Protocol Statistics Window

The Virtual Chassis Protocol Statistics window displays information about the Virtual Chassis protocol statistics on the Virtual Chassis node you selected in the View pane. It contains these panes:

- The top pane of the window lists the Virtual Chassis members and provides the information about each member that is described in [Table 320](#).

Select a member's table row to see information about that member in the other panes.

- The middle and bottom panes provide the information described in [Table 321](#).

Table 320: Virtual Chassis Protocol Statistics Window Top Pane

Table Column	Description
Member	Virtual Chassis member's ID.
Role	Member's Virtual Chassis role. Roles include Master, Backup, and LineCard.
FPC Slot	Member's FPC slot in the Virtual Chassis.
Member Serial Number	Member's serial number.

Table 321: Virtual Chassis Protocol Statistics Window Middle and Bottom Panes

Field or Table Column	Description
System Name	Member system name.
Purges initiated	Number of purges that the system initiated. A purge is initiated if the software determines that a link-state PDU must be removed from the network.
Shortest-path-first runs	Number of shortest-path-first (SPF) calculations that have been performed.
Link-state PDUs queue length	Number of link-state PDUs waiting in the queue for processing. This value is almost always 0.
Link-state PDU fragments computed	Number of link-state PDU fragments that the local system has computed.
Link-state PDUs regenerated	Number of link-state PDUs that have been regenerated. A link-state PDU is regenerated when it is nearing the end of its lifetime and it has not changed.
Protocol data unit type (PDU)	Protocol data unit type.
PDUs Received	Number of PDUs received since VCCP started or since the statistics were set to zero.
PDUs Processed	Number of PDUs received minus the number of PDUs dropped.
PDUs Dropped	Number of PDUs dropped.

Table 321: Virtual Chassis Protocol Statistics Window Middle and Bottom Panes (*continued*)

Field or Table Column	Description
PDU's Transmitted	Number of PDU's transmitted after VCCP started or after the statistics were set to zero.
PDU's Retransmitted	Number of PDU's retransmitted after VCCP started or after the statistics were set to zero.

RELATED DOCUMENTATION

[Understanding the Monitor Mode Tasks Pane | 1277](#)

[Network Director Documentation home page](#)

Viewing Congestion Events

This topic describes how to view congestion events on a device. A congestion event occurs when congestion on a device port exceeds the configured threshold.

You can view congestion events only for devices that support Cloud Analytics Engine and that have the high-frequency traffic statistics feature enabled in Network Director. For information about Cloud Analytics Engine, see [“Understanding Cloud Analytics Engine and Network Director” on page 83](#).

To view congestion events on a device, you must first do the following:

- Configure the Data Learning Engine (DLE) settings under **Preferences > Monitoring > Data Learning Engine Settings**. The DLE is a component of Cloud Analytics Engine. For information on configuring the DLE settings,, see [“Specifying the Data Learning Engine \(DLE\) Settings” on page 118](#).
- Enable high-frequency traffic statistics on the device and optionally configure thresholds. For information, see [“Enabling High-Frequency Traffic Statistics Monitoring on Devices” on page 1202](#).

To view congestion events on a device:

1. In the View pane, select a device on which the high-frequency traffic statistics feature is enabled.
2. Click **Monitor** in the Network Director banner to open Monitor mode.
3. In the Tasks pane, select **Tasks > View Congestion Events**. The View Congestion Events window opens.

The View Congestion Events window lists congestion events that occurred on the device during the time span of 1 minute. The table column headings are the seconds within the selected minute. Each row represents a device interface. Each cell represents the activity on that interface during that second. When congestion events occurred during that second, a bubble appears in the cell. The size of the bubble indicates how many congestion events occurred during that second. The color of the bubble indicates the severity of the congestion during that second: cooler colors indicate lower severity, and hotter colors indicate higher severity.

You can perform these actions in the View Congestion Events window :

- Use the Select Hour and Select Minute lists to select the minute in which to display congestion events and then click **Submit**.
- Mouse over a port name to change the bubbles in its row into the number of congestion events that occurred during each second.
- Click a bubble to open a bar chart that shows detailed information about the congestion events that occurred during that second.

RELATED DOCUMENTATION

[Understanding Cloud Analytics Engine and Network Director | 83](#)

[Enabling High-Frequency Traffic Statistics Monitoring on Devices | 1202](#)

[Understanding Monitor Mode in Network Director | 1269](#)

[Network Director Documentation home page](#)

Monitoring Client Sessions

IN THIS CHAPTER

- Finding User Sessions | 1314
- Finding End Points | 1319
- Monitoring Client Sessions | 1321

Finding User Sessions

IN THIS SECTION

- Procedure for Finding User Sessions | 1314
- Search User Session Window | 1315

This topic describes how to find user sessions on the network. You can search for sessions based on several session attributes. When you find a session, you can view its current and historical bandwidth usage.

This topic describes:

Procedure for Finding User Sessions

1. Click **Monitor** in the Network Director banner.

You can search for user sessions in any tab in Monitor mode.

2. In the Tasks pane, select **Tasks > Find User Session**.

The Search User Session window opens. For information about this window, click the Help button in the title bar of the window or see [“Search User Session Window”](#) on page 1315.

Search User Session Window

The Search User Session window enables you to search for and view information about user sessions. The search scope is the entire managed network, regardless of which node is selected in the View pane. You can view current and historical session information.

To find user sessions:

1. Enter search text in the text box. The search looks for the search text in these session attributes:

- MAC address
- IP address (IPv4 or IPv6)
- User name

2. Click **Search**.

The found user sessions appear in a table. See [Table 322](#) for a description of this table.

3. To view more information about a session, click its table row.

Detailed information about the session appears. The MAC address appears at the top of the page. The page contains these sections:

- **Current Session Information**—Displays information about the current session. [Table 323](#) describes the information shown for sessions connected to the network by a wireless connection. [Table 324](#) describes the information shown for sessions connected to the network by a wired connection.
- **Past Session Information**—Displays information about the MAC addresses' past sessions. This information is not shown for sessions connected to the network by a wired connection. You can select the time period to view from the list above the table. [Table 325](#) describes the information shown. You can expand the record of a past session to see more information about it by using the plus and minus buttons in the left column. [Table 323](#) describes the detailed information shown for sessions connected to the network by a wireless connection.

4. When you are done viewing a session's details, to return to the search results, click **Back** in the top left corner of the window.

Table 322: User Session Details Table for Found User Sessions

Table Column	Description
MAC Address	MAC address of the connected device.
Client IPv4	IPv4 address of the connected device.
User Name	User name of the connected user.

Table 322: User Session Details Table for Found User Sessions (*continued*)

Table Column	Description
Session Type	Shows whether the session is connected by wired or wireless connection.
Client IPv6	IPv6 address of the connected device.
Link-local	Link-local address.

Table 323: Detailed Session Information for Found Wireless User Sessions

Field	Description
User Name	User name of the connected user.
Session Started On	Time when the current session started.
Elapsed Time	Length of time the session has been active.
Client IP	IP address of the connected device. Includes IPv4 and IPv6 addresses.
Controller IP	IP address of the controller to which the client is connected.
VLAN	Name of the VLAN the session is using.
SSID	SSID of the wireless LAN to which the session is connected.
RSSI	Received signal strength indication (RSSI), in decibels referred to 1 milliwatt (dBm). A higher value indicates a stronger signal.
SNR Value	Signal-to-noise ratio (SNR). A measure of the level of a desired signal against the level of background noise, measured in decibels (dB).
Session Location	The physical location of the wireless access point serving the session. The physical location has the hierarchy site–building–floor.
Client Device Type	Client's device type.
Client Device Group	Client's device group.

Table 323: Detailed Session Information for Found Wireless User Sessions (*continued*)

Field	Description
Client Device Profile	Client's device profile.
Unicast KBytes Received	Unicast bytes received by the session.
Unicast KBytes Transmitted	Unicast bytes transmitted by the session.
Multicast KBytes Received	Multicast bytes received by the session.
Link-local	Link-local address.
Data Usage (KBytes)	Data transmitted and received by the client, in kilobytes.
Roaming History	Shows the session's roaming history in a table with these columns: <ul style="list-style-type: none"> • Session Start Time—Time when the session connected to the wireless access point. • AP Name—Name of the wireless access point to which the session connected.
Trends (Statistics)	Shows statistical trends for the session. Select the statistic type to view by using the buttons above the chart.

Table 324: Current Session Information for Found Wired User Sessions

Field	Description
Username	Username of the connected user.
Device IP	IP address of the device.
Authentication Type	Type of authentication used to authenticate the session.
VLAN	Name of the VLAN the session is using.
Device Serial	Device's serial number.
Port	Port to which the device is connected.

Table 325: Past Session Information for Found User Sessions

Table Column	Description
Session Start Time	Time when the current session started.
Elapsed Time	Length of time the session has been active.
Client IPv4	IPv4 address of the connected device.
Client IPv6	IPv6 address of the connected device.
Link-local	Link-local address.
Controller IP	IP address of the controller to which the client is connected.
SSID	SSID of the wireless LAN to which the session is connected.
VLAN	VLAN to which the client is connected.
RSSI	Received signal strength indication (RSSI), in decibels referred to 1 milliwatt (dBm). A higher value indicates a stronger signal.
SNR	Signal-to-noise ratio (SNR),. A measure of the level of a desired signal against the level of background noise, measured in decibels (dB).
RxUniKBytes Value	Unicast bytes received by the session.
RxMultiKBytes Value	Unicast bytes transmitted by the session.
TxUniKBytes Value	Multicast bytes received by the session.
Data Usage (KBytes)	Data transmitted and received by the client, in kilobytes.

RELATED DOCUMENTATION

[Understanding the Monitor Mode Tasks Pane | 1277](#)

[Network Director Documentation home page](#)

Finding End Points

IN THIS SECTION

- [Procedure for Finding End Points | 1319](#)
- [Find End Point Window | 1319](#)
- [Refreshing End Point Information | 1320](#)

This topic describes how to find end points on the network. End points are computing devices that are connected to the network. You can search for end points based on several attributes. When you find an end point, you can see its last known location in the network.

This topic describes:

Procedure for Finding End Points

1. Click **Monitor** in the Network Director banner.

You can search for end points in any tab in Monitor mode.

2. In the Tasks pane, select **Tasks > Find Endpoint**.

The Find End Point window opens. For information about this window, click the Help button in the title bar of the window or see [“Find End Point Window” on page 1319](#).

Find End Point Window

The Find End Point window enables you to search for end points and see their last known location in the network. The search scope is the entire managed network, regardless of which node is selected in the View pane.

To find end points:

1. Enter search text in the text box. The search looks for the search text in these end point attributes:
 - MAC address
 - IP address
2. Click **Search**.

The found end points appear in the Search Results table. See [Table 326](#) for a description of this table.

Table 326: Table of Found End Points

Table Column	Description
MAC Address	MAC address of the connected end point.
IP Address	IP address of the connected end point.
Device Name	Name of the networking device that last saw the end point on the network.
Interface Name	Name of the device interface that last saw the end point on the network.
VLAN	VLAN on which the end point was last seen.
Last Seen	When the end point was last seen on the network.
Actions	Click Verify Current Location to verify the information shown for the end point. If any information changed since the last poll, it is updated in the table.

Refreshing End Point Information

Information about all end points connected to the managed network is polled automatically once every 24 hours. You can refresh this information manually.

NOTE: Refreshing end point information can consume significant system resources and take several minutes to complete, depending on the size of the network and the number of connected end points.

To refresh information about all end points connected to the managed network:

1. Click **Monitor** in the Network Director banner.
2. In the Tasks pane, select **Tasks > Refresh Endpoint**.

A confirmation window opens, listing the job ID of the refresh job.

The endpoint refresh runs as a job. You can monitor the job status in System mode by selecting **Tasks > Manage Jobs** in the Task pane.

RELATED DOCUMENTATION

Monitoring Client Sessions

The Client tab in Monitoring mode provides information about clients and sessions on the network. It is available when the node you select in the View pane contains client and session data. The types of available monitoring data vary depending on the node or node type selected.

To monitor client sessions:

1. Click **Monitor** in the Network Director banner.
2. Select a node in the View pane that contains the client sessions you want to monitor.
3. Select the **Client** tab.
4. To get information about a monitor, click the Help button in its title bar.

The Client monitors include:

- “[Top Users Monitor](#)” on page 1436: shows the clients that use the most bandwidth.
- “[Top Sessions by MAC Address Monitor](#)” on page 1438: shows the sessions that use the most bandwidth.
- “[Session Trends Monitor](#)” on page 1415: shows trends about the number of active sessions.
- “[Current Sessions by Type Monitor](#)” on page 1388 shows the current active sessions by their type.
- “[Current SSID Statistics Monitor](#)” on page 1389 shows SSID statistics information.
- “[SNR SSID Statistics Monitor](#)” on page 1420 shows information about signal-to-noise ratio (SNR) statistics.
- “[Top APs by Session Monitor](#)” on page 1431 shows information about the wireless access points with the most active sessions.

RELATED DOCUMENTATION

Monitoring Radio Frequency

IN THIS CHAPTER

- Monitoring RF 802.11 Packet Errors | **1323**
- Monitoring RF Interference Sources on One Radio | **1325**
- Monitoring RF Interference Sources For Radios on One Access Point | **1328**
- Monitoring RF Interference Sources on Wireless Devices | **1329**
- Troubleshooting Excessive Wireless Interference | **1332**
- Monitoring RF Signal-to-Noise Ratio | **1334**
- Monitoring RF Throughput | **1336**
- Monitoring the Percentage of RF Packet Retransmissions | **1338**
- Monitoring the RF Neighborhood | **1340**
- Monitoring the RF Spectrum of a Radio | **1343**

Monitoring RF 802.11 Packet Errors

Damaged packets cannot be read by devices, therefore they are discarded and re-sent if they are data packets. (See [“Monitoring the Percentage of RF Packet Retransmissions” on page 1338.](#)) This increases the number of re-sent packets and decreases the throughput on the network. (See [“Monitoring RF Throughput” on page 1336.](#)) Voice and video packets are dropped but not re-sent, decreasing the quality of VoIP and streamed media. Packet errors can occur when:

- Noise causes spurious packets.
- Signal degradation occurs causing a weak signal.
- Channels are too congested, causing packet collision and corruption.
- Hardware or drivers are faulty.
- Excessive packets are being received from one source—this could be a flood attack.

In Network Director, at the configured interval (set in [“Setting Up User and System Preferences” on page 108](#)), the total number of 802.11 packet errors is compiled and plotted on a line chart monitor. You can monitor packet errors for the following:







- An access point
- A radio
- An entire site
- A building
- A floor of a building
- A wiring closet

NOTE: Unlike access points and radios, which appear on the list automatically, sites, buildings, floors and wiring closets must be configured by you.

To view RF packet errors over a fixed period of time:

1. Select **Monitor** Mode in the Network Director banner.
2. Select any view from the View pane.
3. Expand the list in the View pane, then select one of the objects listed in [Table 327](#):

Table 327: Objects in the View Pane

Icon	Object
	Individual radio
	Individual access point
	Wiring closet—to create a wiring closet, see “Setting Up Closets” on page 233 .
	Selecting a floor in logical view displays all access points on that floor—to create a floor, see “Configuring Floors” on page 231 .
	Selecting a building in logical view displays all access points in that building—to create a building, see “Configuring Buildings” on page 229 .
	Selecting a site from the logical view displays all access points in that building—to create a site, see “Creating a Site” on page 228 .

The Monitor mode RF tab becomes available when you select one of the objects listed in [Table 327](#) .

4. Click the Monitor mode **RF** tab to view the four basic monitors, which includes the RF packet error monitor.

The populated RF Packet Errors monitor opens, displaying the number of packet errors that the selected object has experienced over the past hour.

5. Optionally, change the timeframe covered by the monitor by selecting a different time from the list.
6. Click **Help** (?) for help interpreting the throughput chart or see [“802.11 Packet Errors Monitor” on page 1384](#).

RELATED DOCUMENTATION

[802.11 Packet Errors Monitor | 1384](#)

[Monitoring the Percentage of RF Packet Retransmissions | 1338](#)

[Configuring Floors | 231](#)

[Configuring Buildings | 229](#)

[Creating a Site | 228](#)

[Monitoring RF Throughput | 1336](#)

Monitoring RF Interference Sources on One Radio

IN THIS SECTION

- [Monitoring RF Radio Interference Sources | 1325](#)
- [RF Interference Sources Pie Chart for a Radio | 1326](#)

Because the 2.4-GHz band includes radio transmissions from devices other than wireless networks, interference is a common problem. Network Director detects, classifies, and displays radio interference in several monitors. This topic describes monitoring the interference of one radio displayed in a pie chart.

NOTE: You can also monitor interference by [“Monitoring RF Interference Sources on Wireless Devices”](#) on page 1329 and [“Monitoring RF Interference Sources For Radios on One Access Point”](#) on page 1328.

Monitoring RF Radio Interference Sources

To view a radio’s RF interference sources in a pie chart over a fixed period of time:

1. Select **Monitor** Mode in the Network Director banner.
2. Select any view from the View pane.
3. Expand the wireless list in the View pane, then select a radio.

The monitor mode RF tab becomes available when you select a radio.

4. In the Tasks pane on the right, click **Interference Sources**.

A pie chart is displayed with a breakdown of the interference sources detected on the selected radio.

5. Optionally, change the timeframe covered by the monitor by selecting a different time from the list.
6. Click **Help (?)** for information about the radio interference chart or see [“RF Interference Sources Pie Chart for a Radio” on page 1326](#).

NOTE: To change the polling interval for monitors, see [“Setting Up User and System Preferences” on page 108](#).

RF Interference Sources Pie Chart for a Radio

The RF Interference Sources pie chart for a single radio reflects all devices that have interfered with the traffic of the radio selected in the View pane. Network Director tracks and monitors interference from these sources:

- Microwave ovens—Most domestic microwave ovens use 2.45 GHz, and can interfere with Wi-Fi channels from 8 to 10 (or even 7 to 11). Interference varies depending on the model of the oven—for example, commercial restaurant microwave ovens sweep over a wider spectrum and have a higher duty cycle.
- Continuous wave devices continuously transmit at a particular frequency without attempting to share the radio frequency medium with other devices. Devices that use continuous wave technology in the same frequency bands as wireless LAN networks will interfere with wireless communications, reducing performance or totally preventing communication. Several examples of devices that use continuous wave transmission that interferes with WiFi are video surveillance cameras and baby monitors.
- Bluetooth devices
- Phone FHSS from cordless phones
- Unknown devices

To track these interference devices, Network Director polls the access point's controller at the standard interval. The categories with the largest sections of the pie cause the most radio interference.

You can perform the following actions on the pie chart:

- Change the time period over which to display interference by selecting a time period from the list in the upper right corner.
- Display a numeric value for interference objects by mousing over a section of the chart.
- Click the monitor's title to see a list of interference incidents along with the information listed in [Table 328](#).

Table 328: Information on RF Interference Sources for a Radio

Information	Description
Last Seen	Date and time the interference was last detected.
Transmitter ID	If the interference is caused by an object with a MAC address, the MAC address is displayed. If the object has no MAC address, MSS calculates a MAC address, using the characteristics of the object. This way, you can correlate interference events over time.
Listener MAC	MAC address of the access point that detected the interference.
AP	Name of the access point that detected the interference.
Controller	Name of the controller that reported the interference.
Channel	Channel the interference affected.
RSSI	Received signal strength indication (RSSI), in decibels referred to 1 milliwatt (dBm). A higher value indicates a stronger signal.
Duty Cycle	Reported fraction of time that the source is emitting RF.
Source Type	Possible sources of interference include Bluetooth, Continuous Wave, Microwave Oven, Unknown, and Phone FHSS.
CIM (%)	Estimated severity of interference on this channel caused by the source.

Interference is frequently not a problem on wireless networks with light traffic, but as traffic becomes heavier, throughput and capacity decrease and other problems become apparent. RF interference can cause packet retransmission (see [“Monitoring the Percentage of RF Packet Retransmissions” on page 1338](#)). Interference is also a security concern because jamming can bring down the network .

Ideally, interference retransmission does not cause more than 10% of the total number of packets sent. If your retransmission percentage is higher, you can try to lower it by:

- Locating and eliminating offending devices. If the item cannot be removed, you can add electromagnetic interference (EMI) shielding such as grounded mesh, foils, insulating foams, or insulating paint. This will limit the interference to a small area.
- Moving the affected access point.
- Moving clients to channels with less interference. Keep in mind, however, that Bluetooth devices, cordless phones, 802.11FH devices, and jamming emissions are broadband, so it's not possible to change channels

away from them—they are everywhere in the band. For more information, see [“Understanding Wireless Radio Channels”](#) on page 856 and [“Understanding Adaptive Channel Planner”](#) on page 861.

For more information about wireless interference, see [“Understanding Wireless Interference”](#) on page 914.

RELATED DOCUMENTATION

[Understanding Wireless Interference | 914](#)

[Monitoring the Percentage of RF Packet Retransmissions | 1338](#)

[Monitoring RF Interference Sources on Wireless Devices | 1329](#)

[Monitoring RF Interference Sources For Radios on One Access Point | 1328](#)

[Network Director Documentation home page](#)

Monitoring RF Interference Sources For Radios on One Access Point

Because the 2.4-GHz band includes radio transmissions from devices other than wireless networks, interference is a common problem. Network Director detects, classifies, and displays radio interference in several monitors. This topic describes the access point interference monitor that can be applied only to access points. If the access point has two radios, interference for each is displayed separately.

NOTE: You can also monitor interference by [“Monitoring RF Interference Sources on One Radio”](#) on page 1325 and [“Monitoring RF Interference Sources on Wireless Devices”](#) on page 1329.

To view an access point’s RF interference sources over a fixed period of time:

1. Select **Monitor** Mode in the Network Director banner.
2. Select any view from the View pane.
3. Expand the wireless list in the View pane, then select an access point.
The RF monitor tab becomes available when you select an access point.
4. Click the **RF** tab.

The RF Interference Sources monitor bar chart for access point radios is displayed as one of the four default monitors.

5. Optionally, change the timeframe covered by the monitor by selecting a different time from the list.
6. Optionally, add or remove a radio on the chart by clicking **Radio1** or **Radio2** in the legend.
7. Click **Help** (?) for information about the access point interference chart or see “[RF Interference Sources Monitor For an Access Point](#)” on page 1412.

RELATED DOCUMENTATION

[RF Interference Sources Monitor For an Access Point | 1412](#)

[Monitoring RF Interference Sources on One Radio | 1325](#)

[Monitoring RF Interference Sources on Wireless Devices | 1329](#)

[Understanding Wireless Interference | 914](#)

[Network Director Documentation home page](#)

Monitoring RF Interference Sources on Wireless Devices










Because the 2.4-GHz band includes radio transmissions from devices other than wireless networks, interference is a common problem. Network Director detects, classifies, and displays radio interference in several monitors. This topic describes the Summary interference monitor that can be applied to various wireless objects such as a radio, an access point, a controller cluster, or an entire wireless network.

NOTE: You can also monitor interference by “[Monitoring RF Interference Sources on One Radio](#)” on page 1325 and “[Monitoring RF Interference Sources For Radios on One Access Point](#)” on page 1328.

To view a wireless object’s RF interference sources over a fixed period of time:

1. Select **Monitor** Mode in the Network Director banner.
2. Select any view from the View pane.
3. Expand the wireless list in the View pane, then select one of the objects listed in [Table 329](#).

Table 329: Wireless Objects in the View Pane

Icon	Object
	Entire Wireless Network in any view.
	Wireless Mobility Domain in any view.
	Controller Cluster in any view. NOTE: You cannot see interference for a single controller.
	Individual access point in any view.
	Individual radio in any view.
	Selecting a floor in logical view displays all access points on that floor—to create a floor, see “Configuring Floors” on page 231 .
	Selecting a building in logical view displays all access points in that building—to create a building, see “Configuring Buildings” on page 229 .
	Selecting a site from the logical view displays all access points in that building—to create a site, see “Creating a Site” on page 228 .
	Wiring closet—to create a wiring closet, see “Setting Up Closets” on page 233 .

The Summary monitor tab becomes available when you select one of the objects listed in [Table 329](#).

4. Click the **Summary** tab.

The RF Interference Sources Summary monitor is displayed as one of the four default monitors.

5. Optionally, change the timeframe covered by the monitor by selecting a different time from the list.
6. Optionally, click the monitor’s title to see a list of interfering objects along with the information listed in [Table 330](#).

Table 330: Information on RF Interference Sources for a Radio

Information	Description
Last Seen	Date and time the interference was last detected.

Table 330: Information on RF Interference Sources for a Radio (continued)

Information	Description
Transmitter ID	If the interference is caused by an object with a MAC address, the MAC address is displayed. If the object has no MAC address, MSS calculates a MAC address, using the characteristics of the object. This way, you can correlate interference events over time.
Listener MAC	MAC address of the access point that detected the interference.
AP	Name of the access point that detected the interference.
Controller	Name of the controller that reported the interference.
Channel	Channel the interference affected.
RSSI	Received signal strength indication (RSSI), in decibels referred to 1 milliwatt (dBm). A higher value indicates a stronger signal.
Duty Cycle	Reported fraction of time that the source is emitting RF.
Source Type	Possible sources of interference include Bluetooth, Continuous Wave, Microwave Oven, Unknown, and Phone FHSS.
CIM (%)	Estimated severity of interference on this channel caused by the source.

7. Click **Help** (?) for information on the RF Interference Sources chart or see [“RF Interference Sources Monitor for Wireless Devices”](#) on page 1409.

NOTE: To change the polling interval for monitors, see [“Setting Up User and System Preferences”](#) on page 108.

RELATED DOCUMENTATION

[RF Interference Sources Monitor for Wireless Devices | 1409](#)

[Monitoring RF Interference Sources on One Radio | 1325](#)

[Monitoring RF Interference Sources For Radios on One Access Point | 1328](#)

Troubleshooting Excessive Wireless Interference

Problem

Description: The network is experiencing a high level of interference. Excessive wireless interference on the wireless network has been reported by one or more of these resources:

- Network Director generates the alarm *RF Interference Detected*. The monitoring feature of Network Director shows high interference—see [“Monitoring RF Interference Sources For Radios on One Access Point” on page 1328](#), [“Monitoring RF Interference Sources on One Radio” on page 1325](#), and [“Monitoring RF Interference Sources on Wireless Devices” on page 1329](#).
- The CLI command **show rfdetect data noise** displays a summary of the noise interference detected.
- In RingMaster, the Client Monitor, RF Monitor, and RF Trends windows show high interference—see the Explore and Status Summary windows of the Monitor tab.

Symptoms: The symptoms of wireless interference include:

- Users are experiencing reduced range for your WiFi network, generally much lower than what is stated in the hardware specifications.
- Users are experiencing sudden drops in transfer speeds, even without much traffic on the network.
- The wireless signal is dropping out in certain places or at certain times during the day.
- Wireless signal strength going up and down randomly.
- Steaming audio, video, or over-network file transfers halt, and then restart

Cause

Leaving the channel number on each radio set to the default value can result in crowding on that channel and high interference levels on the radios.

Overlapping channels can interfere with each other—for more information, see [“Understanding Wireless Radio Channels” on page 856](#).

A high number of CRC errors can indicate a hidden node. Hidden nodes in a wireless network refer to access points that are out of range. Due to the inherent coverage limitations of access points, AP2 may be able to see both AP1 and AP3 and receive data from both access points, but AP1 and AP3 cannot see each other. The problem occurs when both AP1 and AP3 send packets simultaneously to AP2. Since the nodes cannot sense the carrier, Carrier sense multiple access with collision avoidance (CSMA/CA) does not work, and collisions occur, corrupting the data at the access point.

Other wireless products, such as cordless phones, microwave ovens, Bluetooth devices, and test equipment, share the 2.4-gigahertz (GHz) radio frequency bands.

A knowledgeable attacker with the right tools can easily jam the 2.4 GHz frequency in a way that drops the signal to a level where the wireless network can no longer function. For more information, see [“Monitoring RF Interference Sources For Radios on One Access Point” on page 1328](#), [“Monitoring RF Interference Sources on One Radio” on page 1325](#), and [“Monitoring RF Interference Sources on Wireless Devices” on page 1329](#).

Solution

Because RF interference can happen at any time, it is prudent to use any automated responses provided. Automated responses to interference data include:

- Auto-tune channel switching, which selects channels to minimize co-channel interference between access points and also mitigates severe interference by temporarily tuning to another channel.
Use auto-tuning to balance traffic on all channels.
- Rate adaptation, which adjusts the modulation scheme at the access point's transmitter in response to variations in signal-to-noise ratio (SNR) at the receiver. The rate adaptation feature is always enabled to use the best rate to reach associated clients. There are no settings to enable or disable this feature, or change its behavior.

Other methods for mitigating interference include:

- Assign some individual radios to other specific channels. See [“Understanding Wireless Radio Channels” on page 856](#).
- Move interfering devices out of range.
- Choose a cordless phone that uses the 5.8-GHz, 1.9-GHz, or 900-megahertz (MHz) band.
- Overcome the hidden node CRC error problem, implement handshaking in conjunction with the CSMA/CA scheme.

RELATED DOCUMENTATION

[Monitoring RF Interference Sources For Radios on One Access Point | 1328](#)

[Monitoring RF Interference Sources on One Radio | 1325](#)

[Monitoring RF Interference Sources on Wireless Devices | 1329](#)

Monitoring RF Signal-to-Noise Ratio

Signal-to-noise ratio (SNR) is a measure of the level of a desired signal against the level of background noise, measured in decibels (dB). Think of having a conversation in an empty restaurant where your signal (your voice) is clearly heard. Now, think of that same restaurant at noon, when the background noise (all other conversations) is at a peak. You will need to talk louder and lean closer to be heard at noon. In other words, you must increase your signal to overcome the background noise.

Signal-to-noise ratio, along with the bandwidth and channel capacity of a communication channel, affects throughput (see [“Monitoring RF Throughput” on page 1336](#)), especially video throughput.

To view the signal-to-noise ratio on a radio over a fixed period of time:

1. Select **Monitor** Mode in the Network Director banner.
2. Select any view from the View pane.
3. Expand the wireless list in the View pane, then select a radio.

The Monitor mode RF tab becomes available when you select a radio. The Signal-to-Noise Ratio monitor becomes available only when you select a radio.

4. Click the Monitor mode **RF** tab to view the four basic monitors for radios, which includes the Signal-to-Noise Ratio monitor.

By default, the populated Signal-to-Noise Ratio monitor graph uses throughput that the selected object has experienced over the last hour.

5. Optionally, change the timeframe covered by the monitor by selecting a different time from the list.
6. Click **Help** (?) for help interpreting the throughput chart or see [“Signal-to-Noise Ratio Monitor” on page 1418](#).

NOTE: To change the polling interval for monitors, see [“Setting Up User and System Preferences” on page 108](#).

If you are using MSS version 8.0 or higher, you can try using Transmit beam-forming (TxBF) to improve your SNR values. TxBF is a technique that uses an array of transmitting antennas to send radio signals with adjusted magnitude and phase at each antenna to achieve a focused beam target to the receiver. TxBF can increase the Signal-to-Noise Ratio (SNR) at the receiver and improve performance. This feature is supported on the WLA532, WLA321, and WLA322, and is configured from the MSS CLI.

RELATED DOCUMENTATION

[Signal-to-Noise Ratio Monitor | 1418](#)

[Monitoring RF Throughput | 1336](#)

[Monitoring the Percentage of RF Packet Retransmissions | 1338](#)

[Network Director Documentation home page](#)

Monitoring RF Throughput

Network throughput is the average rate of successful message delivery over a communication channel. The larger the channel capacity or bandwidth, the greater the potential throughput. There are also factors that reduce WLAN throughput, such as:

- Network load and congestion
- Encryption
- Transmission error correction and packet retransmission of a given packet
- Quality of Service priority settings
- Configuration of the WLAN Service profiles and Radio profiles in use
- Building construction, internal walls and floors, metallic objects
- Mobile clients
- Overhead in the WLAN itself (Layer 2 and below) and overhead in network protocols
- Interference from nearby transmitters using the same frequency

Network Director tracks and monitors the throughput of wireless data for the following wireless objects:






- An access point
- A radio
- An entire site
- A building
- A floor of a building
- A wiring closet

NOTE: Unlike access points and radios, which appear on the list automatically, sites, buildings, floors and wiring closets must be configured by you.

To view throughput over a fixed period of time for a selected object:

1. Select **Monitor** Mode in the Network Director banner.
2. Select any view from the View pane.
3. Expand the list in the View pane, and then select one of the objects listed in [Table 331](#):

Table 331: Objects in the View Pane With Throughput Data

Icon	Object
	Individual radio
	Individual access point
	Site—Selecting a site in logical view displays all access points on that site—to create a site, see “Creating a Site” on page 228 .
	Building—Selecting a building in logical view displays all access points in that floor—to create a building, see “Configuring Buildings” on page 229 .
	Floor—Selecting a floor in logical view displays all access points on that floor—to create a floor, see “Configuring Floors” on page 231 .

The Monitor mode RF tab becomes available when you select any of the objects listed in [Table 331](#) .

4. Click the Monitor mode **RF** tab to view the RF monitors, which includes the RF Throughput or Packet Throughput monitor.

The populated RF Throughput or Packet Throughput monitor opens, displaying the throughput that the selected object has experienced over the last hour.

5. Optionally, change the timeframe covered by the monitor by selecting a different time from the list.
6. Optionally, select which SSID to monitor from the **Choose SSID** list.
7. Click **Help (?)** for help interpreting the throughput chart or refer to [“RF Throughput or Packet Throughput Level Monitor” on page 1413](#).

NOTE: To change the polling interval for monitors, see [“Setting Up User and System Preferences” on page 108](#).

RELATED DOCUMENTATION

[RF Throughput or Packet Throughput Level Monitor | 1413](#)

[Monitoring the Percentage of RF Packet Retransmissions | 1338](#)

[Configuring Floors | 231](#)

[Configuring Buildings | 229](#)

[Creating a Site | 228](#)

[Network Director Documentation home page](#)

Monitoring the Percentage of RF Packet Retransmissions

IN THIS SECTION

- [Procedure for Viewing RF Packet Transmission | 1338](#)

This monitor reflects the percentage of data packets (but not voice packets) that are retransmitted when they do not reach their destination. The higher the percentage of retransmitted data packets, the slower the throughput. (See “[Monitoring RF Throughput](#)” on page 1336.) Data packets can be retransmitted when:

- Noise causes spurious packets
- Signal degradation occurs causing a weak signal
- Channels are too congested, causing packet collision and corruption
- Hardware or drivers are faulty

Network Director tracks and monitors packet retransmissions for the devices listed in [Table 332](#).

This topic describes:

Procedure for Viewing RF Packet Transmission

You can monitor packet transmission for the following:







- An AP
- A radio
- An entire site
- A building
- A floor of a building
- A wiring closet

NOTE: Unlike access points and radios, which appear on the list automatically, sites, buildings, floors and wiring closets must be configured by you.

To view RF packet retransmission over a fixed period of time:

1. Select **Monitor** Mode in the Network Director banner.
2. Select any view from the View pane.
3. Expand the list in the View pane, and then select one of the objects listed in [Table 332](#):

Table 332: Objects in the View Pane With Packet Retransmission Statistics

Icon	Object
	Individual radio.
	Individual access point.
	Wiring closet—to create a wiring closet, see “Setting Up Closets” on page 233 .
	Selecting a floor in logical view displays all access points on that floor—to create a floor, see “Configuring Floors” on page 231 .
	Selecting a building in logical view displays all access points in that building—to create a building, see “Configuring Buildings” on page 229 .
	Selecting a site from the logical view displays all access points in that building—to create a site, see “Creating a Site” on page 228 .

The Monitor mode RF tab becomes available when you select any of the objects in [Table 332](#).

4. Click the Monitor mode **RF** tab to view the four basic monitors, which includes the RF packet retransmission monitor.

The populated RF packet retransmission monitor opens, displaying the number of packet retransmissions that the selected object has experienced in the last hour.

5. Optionally, change the timeframe covered by the monitor by selecting a different time from the list.
6. Click **Help** (?) for help interpreting the throughput chart or see [“Percentage of Packets Retransmitted Monitor” on page 1396](#).

NOTE: To change the polling interval for monitors, see [“Setting Up User and System Preferences”](#) on page 108.

RELATED DOCUMENTATION

[Percentage of Packets Retransmitted Monitor | 1396](#)

[Monitoring RF Throughput | 1336](#)

[Configuring Floors | 231](#)

[Configuring Buildings | 229](#)

[Creating a Site | 228](#)

[Network Director Documentation home page](#)

Monitoring the RF Neighborhood

IN THIS SECTION

- [Procedure for Viewing a Radio's Neighbors | 1340](#)
- [RF Neighborhood List | 1341](#)

View a radio's neighbors with this list. There are two options for viewing the neighboring devices of a radio. You can either view the results of a selected radio's scan for neighbors, or you can view the result when all other radios in the network found the selected radio. This topic describes both options.

This topic describes:

Procedure for Viewing a Radio's Neighbors

You can monitor the access points operating in the neighborhood of a specified radio.

To view a radio's neighbors:

1. Select **Monitor** Mode in the Network Director banner.
2. Select any View from the View pane.
3. Expand the list in the View pane, and then select a radio.

The RF monitor tab becomes available when an access point or a radio is selected. The neighbors data is available only for radios.

4. Click the Monitor mode **RF** tab. The four basic radio monitors are displayed.
5. Click **Show Neighbors** in the Tasks pane on the right.

The current neighbors for the selected radio are displayed with a list of transmitters heard by this radio.

6. Select either of the options: **Transmitters heard by this radio** or **Listeners that heard this radio**. [Table 333](#) describes these options.

Table 333: Neighbor Tracking Options

Field	Action
Transmitters heard by this radio (default)	Select to view the selected radio's scan results for other transmitters in the neighborhood of the radio.
Listeners that heard this radio	Select to view a list of other radios on the network that can hear the selected radio.

7. Click **Help** (?) for information about the RF Neighborhood list or refer to the "[RF Neighborhood List](#)" on page 1341.

RF Neighborhood List

The RF Neighborhood list includes either the neighbors located by the radio you indicated in the View pane or the neighbors that can hear the radio you indicated in the View pane. Determine which way the data will be displayed by selecting either **Transmitters heard by this radio** or **Listeners who heard this radio**. Either way, the neighbor details reported are described in [Table 334](#).

Table 334: Data For Neighbors Located by the Radio Scan

Field	Description
Neighbor	Wireless access point radio in close enough proximity to be detected by another access point radio.
BSSID	Identifier for an access point.
Channel	Number of the channel used by the neighbor radio—in the 2.4-GHz band, this is usually 1,6, an 11 in the US. In the rest of the world, channels 1, 5, 9, and 13 are used most often. The 5-GHz band has 24 usable channels, (36,1) (40,-1) (44,1) (48,-1) (52, 1) (56,-1) (60,1) (64,-1) (100,1) (104,-1) (108,1) (112,-1) (116,1) (120,-1) (124,1) (128,-1) (132,1) (136,1) (149,1) (153,-1) (157,1) (161,-1). The +1 and -1 indicated for some channels above indicate channel bonding, where a channel bonds with the one above or below it.
RSSI	Received signal strength indicator (RSSI) is the relative received strength of a signal in a wireless environment. RSSI is basically an indication of the power level being received by the antenna—therefore, the higher the RSSI number, the stronger the signal. (Because the numbers are negative, -50 represents a stronger signal than -88.)

You can perform the following actions on this list:

- Re-sort the list based on the values in any column by mousing over the column title, and then selecting one of these options from the list that appears:
 - Sort Ascending
 - Sort Descending

NOTE: The RSSI column also has a built-in arrow in the title that sorts the list by RSSI order.

- Remove any of the displayed columns by mousing over the column title, selecting **Columns** from the list that appears, and then adding or removing the check marks from Neighbor, BSSID, Channel, or RSSI.

RF neighbor information is available only for radios.

NOTE: To change the polling interval for monitors, see [“Setting Up User and System Preferences” on page 108](#).

RELATED DOCUMENTATION

[Understanding Wireless Scanning | 869](#)

[Network Director Documentation home page](#)

Monitoring the RF Spectrum of a Radio

IN THIS SECTION

- [Procedure for Viewing the Radio Spectrogram | 1343](#)
- [Spectrogram Charts | 1344](#)
- [Channel Spectrogram Chart | 1345](#)
- [Swept Spectrum and Duty Cycle Charts | 1346](#)

WLAN radios are continually scanning for potential clients. In addition to finding clients, these scans detect other electronic objects, such as other access points and various non-802.11 equipment. (For more information about scanning, see [“Understanding Wireless Scanning” on page 869.](#)) You can view some of the data collected by the Network Director scanning function by looking at a radio spectrogram in Network Director.

The Network Director radio spectrogram consists of two charts, the Channel Spectrogram on the top half of the screen, and the Spectrum Sweep on the bottom half of the screen. The Channel Spectrogram displays the selected radio’s channel activity, while the Spectrum Sweep displays objects detected in the selected radio’s scanned area.

NOTE: A spectrogram times out after five minutes. During this time, all clients are dropped.

This topic describes:

Procedure for Viewing the Radio Spectrogram

To view the status of the spectrum:

1. Select **Monitor** Mode in the Network Director banner.
2. Select **Logical View** from the View pane.

- Expand the list in the View pane, and then select a radio.

TIP: The Monitor mode RF tab becomes available when you select either an access point or a radio. The RF Spectrogram monitor becomes available only when you select a radio.

- Click the **RF** tab.
- Click **Spectrogram** from the Tasks pane. Two windows open, the RF 2.4-GHz Spectrogram window and the Swept Spectrum and Duty Cycle window.
- Configure the sweep criteria for the four lines on this chart by adding or deleting the plotted lines. Select or clear the check boxes for Max, Max Hold, Duty Cycle, or Channels. (For an explanation of these values, see [Table 335](#).)
- Click **Start** to begin the sweep using the criteria that you defined.



WARNING: Your clients are dropped from the selected radio during a sweep.

- Sweeps take place in spectral scan mode, which drops all client connections during the sweeps. You are asked if you want to continue. Click **Yes** to drop clients and continue the configured sweeps.
- The upper graph displays the real-time state of the radio's channels, while the lower graph displays the electronic state of the radio's spectrum. Click **Help** (?) for more information about the graphs or see "[Spectrogram Charts](#)" on page 1344.
- When you have seen the desired information, click **Stop**.
The radio returns to beaconing for potential clients.

Spectrogram Charts

Use the following information to interpret the spectrogram.

The Network Director radio spectrogram consists of charts representing signal strength. On the top half of the screen, the channel power spectrogram for either the 2.4-GHz channel or the 5-GHz channel is displayed. On the lower half of the screen, two tabs display the ongoing spectrum power sweep and the duty cycle sweep.



WARNING: Your clients are dropped during a spectrum sweep.

Channel Spectrogram Chart

The channel in use (2.4-GHz or 5-GHz) by the selected radio in the **View** pane is displayed as a grid with these parameters:

- Maximum power
- Average power
- Duty cycle
- Maximum hold

The 2.4-GHz channel has only one tab and all data is plotted on that tab. The larger 5-GHz channel is broken down into tabs for channels: CH-36 to CH-48, CH-52 to CH-64, CH-100 to CH-140, and CH-149 to CH-165. Click each 5-GHz tab to view the individual charts.

Four lines can be plotted on each chart, as listed in [Table 335](#):

Table 335: Spectrum Sweep Results

Line	Description
Max Power (red)	Indicates maximum power usage for this radio, plotted at the configured polling interval.
Avg Power (green)	Indicates average power usage for this radio, plotted at the configured polling interval.
Duty Cycle (yellow)	Indicates the amount of RF energy present in the spectrum as a result of an object emitting RF.
Max Hold (pink)	Peak power level that was seen across multiple samples.

NOTE: The polling interval can be reconfigured by [“Setting Up User and System Preferences”](#) on page 108.

You can perform the following actions on this chart:

- Highlight a line in the graph by mousing over the line's legend.
- Display a numeric value by placing the cursor where a vertical grid line bisects a data line.
- Remove any or all of the four lines on this chart by clearing the check boxes for Max, Max Hold, Duty Cycle, or Channels.
- Click **Start** to drop clients and begin the sweep using the criteria that you defined. Click **Stop** to stop the spectrum sweep and return the radios to normal operation.



WARNING: Your clients are dropped during a spectrum sweep.

Swept Spectrum and Duty Cycle Charts

The chart on the bottom half of the window is dynamic, with new sweep results added to the top of the chart after each polling period. The power detected in the sweep is indicated by color, with blue representing the lowest power and red representing the highest power. Typically, the sweep results are blue and green, with occasional yellow segments. Any red that appears in the sweep indicates a problem.

This chart has two tabs, one for average power used in the spectrum called *Swept Spectrum*, and one for the percentage of the *Duty Cycle* used. When you are scanning for either result, all other traffic is dropped. For this reason, you must click **Start** to view spectrum results or see the duty cycle plotted. When you click **Stop**, normal traffic is resumed.

Click the **Swept Spectrum** tab to view the plotted average power spectrum for the radio that you selected in the View pane. The parameters used for this chart are time and average power.

Click the **Duty Cycle** tab to view the duty cycle data for the radio that you selected in the View pane. The parameters used for this chart are time and duty cycle percentage.

This data in these charts is available only for radios.

RELATED DOCUMENTATION

[Understanding Wireless Scanning | 869](#)

[Network Director Documentation home page](#)

Monitoring Devices

IN THIS CHAPTER

- [Comparing Device Statistics | 1347](#)
- [Showing ARP Table Information | 1349](#)
- [Viewing PoE Information | 1350](#)
- [Monitoring Backed-Up Wireless Access Points on Wireless LAN Controllers | 1352](#)
- [Monitoring the Status of Aggregated Access Points and Radios | 1354](#)
- [Monitoring the Status of Logical Interfaces | 1354](#)
- [Monitoring the Status of Wireless Controllers, Access Points, and Radios | 1356](#)
- [Monitoring the Status of a Virtual Chassis | 1357](#)
- [Monitoring the Status of Virtual Chassis Members | 1358](#)

Comparing Device Statistics

IN THIS SECTION

- [Procedure for Comparing Device Statistics | 1347](#)
- [Compare Interfaces Window | 1348](#)

This topic describes how to compare statistics from multiple network devices and interfaces in real time. You select which devices, interfaces, and counters to compare, and how often to poll for new statistics.

This topic describes:

Procedure for Comparing Device Statistics

1. Click **Monitor** in the Network Director banner.

You can compare device statistics in any tab in Monitor mode.

2. In the Tasks pane, select **Tasks > Compare Device Statistics**.

The Compare Interfaces window opens. For information about this window, click the Help button in the title bar of the window or see [“Compare Interfaces Window” on page 1348](#).

Compare Interfaces Window

The Compare Interfaces window enables you to compare statistics from multiple device interfaces in real time. The search scope is the entire managed network, regardless of which node is selected in the View pane.

To compare device statistics:

1. Select the devices to compare from the device tree in the Select Devices section.
 2. Select a device in the Selected Devices section to select which of its interfaces to compare.

The Select Interfaces section lists the device’s interfaces. You can select up to two interfaces per device.
 3. Select an Interface in the Select Interfaces section to select which of its counters to compare.

The Select Counters section lists the interface’s counters.
 4. Select the counters to compare in the Select Counters section.
5. Repeat the process of selecting devices, interfaces, and counters to compare until you are finished selecting what to compare.
6. Select how often the data will be refreshed from the **Data Collection Frequency** list.
7. Click the **Compare** button to start comparing information.

A page opens containing a line graph for each counter you selected. Each graph displays all the interfaces for which its counter is selected.
8. To pause data collection, click the **Pause** button. To resume data collection, click the **Resume** button.
9. To change data collection settings, click the **Back** button.

RELATED DOCUMENTATION

[Understanding the Monitor Mode Tasks Pane | 1277](#)

[Network Director Documentation home page](#)

Showing ARP Table Information

IN THIS SECTION

- Procedure for Showing ARP Table Information | 1349
- Show ARP Table Information Window | 1349

This topic describes how to show Address Resolution Protocol (ARP) table information for a device. ARP table information is collected from the selected device when this task runs. You can search for ARP table records.

Procedure for Showing ARP Table Information

To show ARP table information for a device:

1. Click **Monitor** in the Network Director banner.
2. Select the device in the View pane that you want to monitor.
3. Select **Tasks > Show ARP Table** in the Task pane.

The Show ARP Table Information window opens. For information about this window, click the Help button in the title bar of the window or see [Table 336](#). You can click the Refresh button below the table to refresh the data from the device.

Show ARP Table Information Window

The Show ARP Table Information Window shows information from the selected device's ARP table.

Table 336: Show ARP Table Information Window

Control or Column	Description
Search controls	Search for ARP table records. Enter search text in the text box. The table of ARP records displays only matching records. Click the X button to clear the search and display all records.
MAC Address	MAC address.
IP Address	IP address.

Table 336: Show ARP Table Information Window (continued)

Control or Column	Description
Interface Name	Interface name.
Expiring in (sec)	Number of seconds until the record expires from the ARP table.

RELATED DOCUMENTATION

[Understanding the Monitor Mode Tasks Pane | 1277](#)

[Network Director Documentation home page](#)

Viewing PoE Information

IN THIS SECTION

- [Procedure for Viewing PoE Information | 1350](#)
- [Show PoE Information Window | 1351](#)

This topic describes how to view Power over Ethernet (PoE) information for EX devices.

Procedure for Viewing PoE Information

To view PoE information for a device:

1. Click **Monitor** on the Network Director banner.
2. Select the device in the View pane that you want to monitor.
3. Select **Tasks > Show PoE Interfaces** in the Tasks pane.

The Show PoE Information window opens. For information about this window, click the Help button in the title bar of the window or see [Table 337](#). You can click the **Refresh** button below the table to refresh the data from the device.

Show PoE Information Window

The Show PoE Information window shows information about the PoE ports of the selected device.

Table 337: Show PoE Information Window

Control or Column	Description
Search controls	Search for PoE records. Enter search text in the text box. The PoE table displays only the matching records. Click the X button to clear the search and display all records.
Interface Name	Name of the PoE interface.
Admin Status	Administrative state of the PoE interface—Enabled or Disabled. If the PoE interface is disabled, it can provide network connectivity, but it cannot provide power to connected devices.
Operational Status	Operational status of the PoE interface. The operational status can have one of the following values: <ul style="list-style-type: none"> • ON—The interface is currently supplying power to a powered device. • OFF—PoE is enabled on the interface, but the interface is not currently supplying power to a powered device. • FAULT—PoE interface is in the OFF state due to a fault condition. • Disabled—PoE is disabled on the interface.
Max Power Limit	Maximum power that can be provided by the interface.
Priority	Interface power priority—High or Low.
Power Consumption	Amount of power being used by the interface.
Class	Class of the powered device—IEEE 802.3af (PoE) or IEEE 802.3at (PoE+). Class 0 is the default class and is used when the class of the powered device is unknown. If no powered device is connected, this column displays Not Applicable.
LLDP Negotiation Priority	Interface power priority negotiated by LLDP.
LLDP Negotiation Power	Amount of power negotiated by LLDP, to be used by the interface.

RELATED DOCUMENTATION

Monitoring Backed-Up Wireless Access Points on Wireless LAN Controllers

IN THIS SECTION

- [Procedure for Monitoring Backed-Up Wireless Access Points on Wireless LAN Controllers | 1352](#)
- [Backed-Up APs Window | 1353](#)

This topic describes how to monitor backed-up wireless access points on a wireless LAN controller. You can monitor backed-up wireless access points on wireless LAN controller nodes in any view.

When wireless LAN controllers are configured in a cluster, each wireless access point in the cluster can have a primary controller and a secondary controller. You can see which wireless access points use the selected wireless LAN controller as their secondary controller. The term backed-up means that the wireless LAN controller is acting as a secondary (backup) controller for a wireless access point.

This topic describes:

Procedure for Monitoring Backed-Up Wireless Access Points on Wireless LAN Controllers

1. Click **Monitor** in the Network Director banner.
2. Select the wireless LAN controller in the View pane that contains the backed-up wireless access points you want to monitor.
3. Select the **Summary** tab or the **Equipment** tab.
4. In the Tasks pane, select **View > Backed-Up APs**.

The Backed-Up APs window opens. For information about this window, click the Help button in the title bar of the window or see [“Backed-Up APs Window” on page 1353](#).

Backed-Up APs Window

The Backed-Up APs window contains a table with information about the backed-up wireless access points on the node you selected in the View pane. See [Table 338](#) for a description of the table information.

Table 338: Backed-Up APs Table

Table Column	Description
AP Name	Name of the access point.
Serial Number	Serial number of the access point.
Model	The model number of the access point.
IP Address	The IP address assigned to the AP.
Status	Operational status of the access point: <ul style="list-style-type: none"> • Down—The access point is offline. • Up—The access point is online and enabled. • Up Redundant—The access point is online, reporting to this controller as redundant and another controller as primary.
Uptime	The length of time since the access point last booted.
Version	The version of the Mobility System Software (MSS) running on the access point.
Primary Controller	The primary controller for the access point.
Secondary Controller	The secondary controller for the access point.

RELATED DOCUMENTATION

[Understanding the Monitor Mode Tasks Pane | 1277](#)

[Network Director Documentation home page](#)

Monitoring the Status of Aggregated Access Points and Radios

Network Director provides two monitors that display an aggregate view of access points and radios at the network level. Aggregation is available for Controller Clusters, Wireless Mobility Domains, and the Wireless Network.

To locate these monitors:

1. Click **Monitor** in the Network Director banner to ensure that you are in Monitor mode. Monitor mode works in all views (Logical, Location, Device).
2. Select a supported node in the network tree (**Wireless Network, Wireless Mobility Domain, or Controller Cluster**). The Equipment tab is the default tab that displays the AP Status and Radio Status.
3. Click the Help icon on the monitor to learn more about the purpose or fields on a monitor.

The two monitors that display aggregate information for APs and Radios are:

- [“AP Status Monitor” on page 1386](#), summarizes of all of the network access points
- [“Radio Status Monitor” on page 1403](#), summarizes of all of the network radios

RELATED DOCUMENTATION

[AP Status Monitor | 1386](#)

[Radio Status Monitor | 1403](#)

[Network Director Documentation home page](#)

Monitoring the Status of Logical Interfaces

IN THIS SECTION

- [Locating Information about Logical Interfaces | 1355](#)
- [Show Logical Interface Information Table | 1355](#)

Network Director provides real-time statistics on logical Ethernet switching interfaces for switches, routers, Virtual Chassis, QFabric systems, and Layer 3 Fabrics.

This topic describes:

Locating Information about Logical Interfaces

Real-time logical interface statistics, including VLAN information, are available from the Show Logical Interfaces window in Monitoring mode. To find this information:

1. Select **Monitor** in the Network Director banner.
2. Select a node in the View pane that contains the logical interface you want to monitor.
3. Select the Equipment tab.
4. Click **Logical Interfaces** in the Tasks pane to open the Show Logical Interface Information table in main window.

Show Logical Interface Information Table

The Show Logical Interface Information table provides interface, VLAN, and spanning-tree status for an interface. The information is presented in a tabular format. The fields in the Show Logical Interface Information table are described in [Table 339](#).

Table 339: Show Logical Interface Information Fields

Field	Description
Logical Interface Name	The logical interface name.
Serial Number	The serial number of the device to which the logical interface belongs.
VLAN Membership ID	The VLAN to which the interface belongs.
Bridge Domain Membership	The bridge domain to which the interface belongs (for only devices that do not use the Enhanced Layer 2 Software (ELS)).
802.1Q Tag	The IEEE 802.1Q identifier for the VLAN.
Tagging	Indicates whether the packets entering the port are tagged or untagged.
Logical Interface State	Indicates whether the logical interface is up or down.

Table 339: Show Logical Interface Information Fields (*continued*)

Field	Description
STP State	Indicates whether the interface is discarding (blocked) or forwarding (unblocked).
Port Mode	Indicates one of three modes: access, tagged-access, or trunk. <ul style="list-style-type: none"> • Access—The interface can be in a single VLAN only. • Tagged-access—The interface can accept tagged packets from one access device. • Trunk—The interface can be in multiple VLANs and accept tagged packets from multiple devices.

RELATED DOCUMENTATION

[Monitoring Traffic on Devices | 1283](#)

[Monitoring Traffic on Layer 3 VLANs | 1287](#)

[Network Director Documentation home page](#)

Monitoring the Status of Wireless Controllers, Access Points, and Radios

When you select the node of a wireless controller in any view, Network Director presents four monitors. These monitors give you the overall workload and performance of the controller and the APs and radios under its control.

To locate the monitors that are specific to wireless controllers:

1. Click **Monitor** in the Network Director banner to ensure that you are in Monitor mode..
2. Expand the network tree to expose the wireless controller.
3. Select the wireless controller.
4. Click the **Equipment** tab. The page opens with the four monitors.

The available monitors are:

- [“Resource Monitor For Wireless LAN Controllers” on page 1407](#), that provides a graphical representation of CPU usage and memory consumption.
- [“Equipment Status Summary Monitor” on page 1393](#), that provides state information on ports, access points and radios.
- [“AP Status Monitor” on page 1386](#), that provides details of each AP including the current uptime.
- [“Radio Status Monitor” on page 1403](#), that provides details of each radio under the controller’s control.

RELATED DOCUMENTATION

[Resource Monitor For Wireless LAN Controllers | 1407](#)

[Equipment Status Summary Monitor | 1393](#)

[AP Status Monitor | 1386](#)

[Radio Status Monitor | 1403](#)

[Network Director Documentation home page](#)

Monitoring the Status of a Virtual Chassis

When you select a Virtual Chassis from the network tree in any view, four monitors are displayed that give at-a-glance information about the status and performance of the Virtual Chassis. Use this information to monitor the chassis as a whole, without reviewing each switch independently.

To locate the Virtual Chassis monitors:

1. Click **Monitor** in the Network Director banner to ensure that you are in Monitor mode.
2. Expand the network tree to expose the Virtual Chassis node.
3. Select the Virtual Chassis in the network tree.
4. Click the **Equipment** tab to display the four monitors.
5. Click the Help icon on the monitor learn more about the purpose or fields on a monitor.

The four monitors are:

- [“Resource Utilization Monitor for Switches, Routers, Virtual Chassis, Virtual Chassis Fabrics, and QFabric Systems” on page 1408](#), that provides information about the composition of the chassis, its members, and the location of neighboring switches.
- [“Status Monitor for Virtual Chassis” on page 1427](#), that provides information about the uptime, IP address, and hostname of the Virtual Chassis.
- [“Resource Monitor For Wireless LAN Controllers” on page 1407](#), that provides a graphical representation of CPU usage and memory consumption.
- [“Port Status Monitor” on page 1397](#), that provides port level status information.

RELATED DOCUMENTATION

[Monitoring the Status of Virtual Chassis Members | 1358](#)

[Resource Utilization Monitor for Switches, Routers, Virtual Chassis, Virtual Chassis Fabrics, and QFabric Systems | 1408](#)

[Status Monitor for Virtual Chassis | 1427](#)

[Resource Monitor For Wireless LAN Controllers | 1407](#)

[Port Status Monitor | 1397](#)

[Network Director Documentation home page](#)

Monitoring the Status of Virtual Chassis Members

When you select a member of a Virtual Chassis in the any view, Network Director displays four monitors. At the member node level, the information is highly specific to the equipment.

To locate these monitors:

1. Click **Monitor** in the Network Director banner to ensure that you are in Monitor mode.
2. Expand the network tree to expose the member of the Virtual Chassis node.
3. Select the Virtual Chassis in the network tree.
4. Click the **Equipment** tab to display the monitors.
5. Click the Help icon on the monitor to learn more about the purpose or fields on a monitor.

The monitors at this level are:

- [“Port Status Monitor” on page 1397](#), that provides summary and detailed information about the status of the physical network interfaces for the selected node in the View pane.
- [“Power Supply and Fan Status Monitor” on page 1400](#), that provides a graphical representation of the operating condition for this member. These graphs also show the ratio of filled slots to available power and fan slots.
- [“Status Monitor for Virtual Chassis Members” on page 1428](#), that provides status information for this member of the Virtual Chassis.

RELATED DOCUMENTATION

[Monitoring the Status of a Virtual Chassis | 1357](#)

[Port Status Monitor | 1397](#)

[Power Supply and Fan Status Monitor | 1400](#)

[Status Monitor for Virtual Chassis Members | 1428](#)

[Network Director Documentation home page](#)

Monitoring and Analyzing Fabrics

IN THIS CHAPTER

- Monitoring Junos Fusion Fabric Systems and Components | 1360
- Analyzing QFabric Devices | 1361
- Monitoring QFabric Devices and Components | 1365
- Analyzing Virtual Chassis Fabrics | 1367

Monitoring Junos Fusion Fabric Systems and Components

This topic describes how to monitor Junos Fabric systems and their components, and which Junos Fusion-specific monitors are available.

To monitor a Junos Fusion system or its components:

1. Click **Monitor** in the Network Director banner.
2. Select the fabric container or a specific Junos Fusion fabric that you want to monitor in the View pane.

The tabs and monitors that are available depends on your selection.

NOTE: You cannot select the Aggregation Device or Satellite Device nodes within a fusion fabric in the View pane. To monitor a fabric, select the fusion fabric parent node.

3. To get information about a monitor, click the Help button in its title bar, or refer to [Table 340](#) for information about the Junos Fusion-specific monitors that are available for each node type.

Table 340: Monitors Available for Junos Fusion Systems and Components

Selection in the View Pane	Monitoring Tab	Available Junos Fusion-Specific Monitors
Junos Fusion system	Summary	<ul style="list-style-type: none"> • Equipment Summary By Type Monitor on page 1394 • Port Status Monitor on page 1397 • Current Active Alarms Monitor on page 1464
	Traffic	<ul style="list-style-type: none"> • Unicast vs Broadcast/Multicast Monitor on page 1440 • Unicast vs Broadcast/Multicast Trend Monitor on page 1441 • Traffic Trend Monitor on page 1439 • Error Trend Monitor on page 1391
	Equipment	<ul style="list-style-type: none"> • Status Monitor for Junos Fusion Systems on page 1421 • Port Status Monitor on page 1397
Fabrics node	Summary	<ul style="list-style-type: none"> • Equipment Summary By Type Monitor on page 1394 • Port Status Monitor on page 1397 • Current Active Alarms Monitor on page 1464

RELATED DOCUMENTATION

[Understanding Monitor Mode in Network Director | 1269](#)

[Network Director Documentation home page](#)

Analyzing QFabric Devices

IN THIS SECTION

- [Procedure for Analyzing a QFabric Device Manually | 1362](#)
- [Using the Fabric Health Check Tab | 1362](#)
- [Using the Topology Tab | 1363](#)

This topic describes how to analyze QFabric devices. The Run Fabric Analyzer task analyzes a QFabric device and provides information about its health, connectivity, and topology.

NOTE: Fabric analysis is supported for QFabric devices running Junos Release 13.1R2 or later.

You must identify the QFabric's control plane switches by running the Setup QFabric task to get information about the control plane's health and topology.

The Fabric Analyzer runs automatically when you discover a QFabric device and when you change its configuration by using the Setup QFabric task. You can also run the Fabric Analyzer manually.

This topic describes:

Procedure for Analyzing a QFabric Device Manually

1. Click **Monitor** in the Network Director banner.
2. Select the QFabric device to analyze in the View pane.
3. In the Tasks pane, select **Tasks > Run Fabric Analyzer**.

The results of the analysis appear on the Fabric Analysis tab in Monitor mode, when the QFabric device is selected in the View pane. For information about using the tabs within this tab, see the following sections:

- [Using the Fabric Health Check Tab on page 1362](#)
- [Using the Topology Tab on page 1363](#)

Using the Fabric Health Check Tab

To check the health of a QFabric device, select the **Fabric Health Check** tab on the Fabric Analysis tab in Monitor mode. The Fabric Health Check tab contains these sections:

- The Control Plane Health Check section shows information about the health of the QFabric device's control plane. It displays a summary of the control plane health checks performed on the QFabric device. For each summary category, the number of failed (in red) and passed (in green) tests is shown.
- The Data Plane Health Check section shows connectivity information about the QFabric device's data plane in a table. Each node device is listed in the Node Device column, and each Interconnect device is a table column. Each cell indicates the connectivity status between the node device and the Interconnect device.

You can download a Fabric Analysis report by clicking the Download button in the tab's title bar. The report downloads as a zip file, which you unzip on your local machine. The report contains the following files:

- Fabric health check results in XML format.
- Connectivity check results in XML format.
- A Control Plane Topology diagram in SVG format.
- A Data Plane Topology diagram in SVG format.


Using the Topology Tab

Network Director enables you to view the data plane topology and the control plane topology. The data plane is a redundant, high-performance, and scalable data plane that carries QFabric system data traffic. The control plane is the network that carries control traffic between the various QFabric system components such as node devices, interconnect devices, director devices, and the control plane Ethernet switches.

The Topology tab displays the QFabric connectivity in the Data plane—with the interconnect in the center surrounded by the nodes— as shown in [Figure 56](#). Mouse over a device to view the port details and the connection state of that device.

By default, Network Director displays the Data plane topology for the selected QFabric. To view the control plane topology, click **Control Plane**.

Figure 56: Fabric Analyzer Topology View of a QFabric Switch



In the topology diagram, the connection details are represented by green, yellow, and red lines.

- The green line indicates that the node is connected to all the Interconnects properly and all functions are normal.
- The yellow line indicates that the node is only partially connected to the Interconnect. That is, the node might be connected to some of the interconnects, but not all.
- The red line indicates that the node is not connected to any of the interconnects.

The following details are displayed for each device, if the details are configured on the device:

- Name of the device provided while configuring the device. The device name is displayed as a label.
- Platform of the device. Platform can be QFX3500, QFX3600, or QFX5100 switches.
- Node group name. The name of the node group to which this device belongs to.

- Node group type. The server name of the node group to which this device belongs to.
- Source Port: source port of the device.
- Destination port: The port number on the destination device to which the source device is connected to.

RELATED DOCUMENTATION

[Understanding the Monitor Mode Tasks Pane | 1277](#)

[Setting Up QFabric Systems | 780](#)

[Network Director Documentation home page](#)

Monitoring QFabric Devices and Components

This topic describes how to monitor QFabric devices and their components, and which QFabric-specific monitors are available.

To monitor a QFabric device or its components:

1. Click **Monitor** in the Network Director banner.
2. Select a QFabric device node or a node within a QFabric device node that you want to monitor in the View pane.

The tabs and monitors that are available depend on the type of node you select.

3. To get information about a monitor, click the Help button in its title bar, or refer to [Table 341](#) for information about the QFabric-specific monitors that are available for each node type.

Table 341: Monitors Available for QFabric Devices and Components

QFabric Node Type Selected in the View Pane	Monitoring Tab	Available QFabric-Specific Monitors
QFabric device	Summary	<ul style="list-style-type: none"> • Status Monitor for QFabric Systems on page 1424 • Access vs. Uplink Port Utilization Trend Monitor on page 1385 • Node Device Summary Monitor on page 1395 • QFabric Interconnect Status Summary Monitor on page 1402
	Equipment	“Status Monitor for QFabric Systems” on page 1424
	Fabric Analysis	For information about fabric analysis, see “Analyzing QFabric Devices” on page 1361.
Directors folder	Summary	“QFabric Director Status Monitor” on page 1401
Director device	Summary	<ul style="list-style-type: none"> • Status Monitor for QFabric Directors on page 1423 • QFabric VM Status Summary Monitor on page 1403
Interconnects folder	Summary	“QFabric Interconnect Status Summary Monitor” on page 1402
Interconnect device	Equipment	“Status Monitor for QFabric Interconnects” on page 1425
Node Group folder (applies to Server, Redundant Server, and Network Node Groups)	Summary	“Node Device Summary Monitor” on page 1395
Node Group (applies to Server, Redundant Server, and Network Node Groups)	Summary	“Node Device Summary Monitor” on page 1395
Node device (applies to Server, Redundant Server, and Network Nodes)	Summary	“Access vs. Uplink Port Utilization Trend Monitor” on page 1385
	Equipment	“Status Monitor for QFabric Nodes” on page 1425

RELATED DOCUMENTATION

[Understanding Monitor Mode in Network Director](#) | 1269

[Network Director Documentation home page](#)

Analyzing Virtual Chassis Fabrics

IN THIS SECTION

- [Procedure for Analyzing a Virtual Chassis Fabric | 1367](#)
- [Using the Fabric Health Check Tab | 1367](#)
- [Using the Topology Tab | 1368](#)

This topic describes how to analyze Virtual Chassis Fabrics (VCFs). The Run Fabric Analyzer task analyzes a VCF and provides information about its health, connectivity, and topology. The Fabric Analyzer works on VCFs configured in spine-and-leaf mode.

This topic describes:

Procedure for Analyzing a Virtual Chassis Fabric

Fabric analysis is automatically performed on managed VCFs. You do not need to manually run the analyzer.

To see the results of the analysis of a VCF:

1. Click **Monitor** in the Network Director banner.
2. Select the VCF to analyze in the View pane.
3. Select the **Fabric Analysis** tab.

For information about using the tabs within the Fabric Analysis tab, see the following sections:

- [Using the Fabric Health Check Tab on page 1367](#)
- [Using the Topology Tab on page 1368](#)

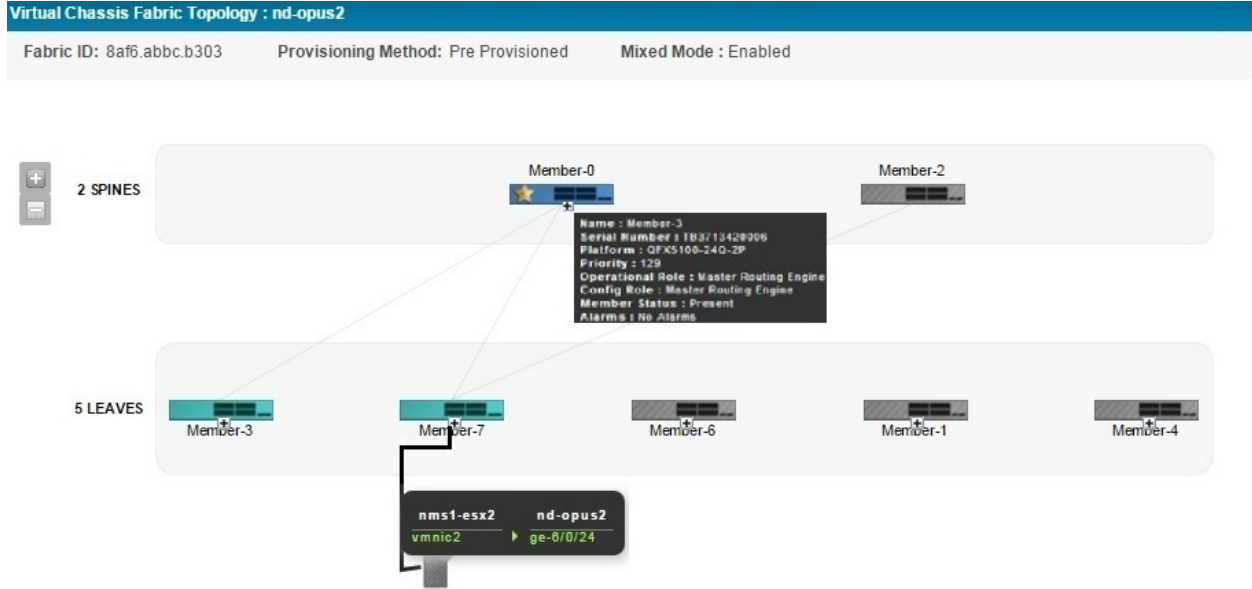
Using the Fabric Health Check Tab

To check the health of a VCF, select the **Fabric Health Check** tab on the Fabric Analysis tab in Monitor mode. The Fabric Health Check tab shows connectivity information about the VCF. Each leaf device is listed in the Leaves column, and each spine device is a table column. Each cell indicates the connectivity status between the leaf device and the spine device.

Using the Topology Tab

The Topology tab displays a topology diagram of the VCF as shown in [Figure 57](#). Mouse over a member to view details of that member.

Figure 57: Virtual Chassis Fabric Topology



The details that are displayed in the top panel of the Virtual Chassis Topology view are described in [Table 342](#).

Table 342: Common Details for the Virtual Chassis and Virtual Chassis Fabric

Details	Description
Fabric ID VC ID	All members of a Virtual Chassis configuration share one Virtual Chassis identifier (VCID). This identifier is derived from internal parameters.

Table 342: Common Details for the Virtual Chassis and Virtual Chassis Fabric (continued)

Details	Description
Provisioning Method	<p>The provisioning mode of the member. Provision mode can be <i>autoprovisioned</i>, <i>preprovisioned</i> or <i>not preprovisioned</i>.</p> <p>Autoprovisioning a Virtual Chassis Fabric (VCF) enables you to “plug and play” devices into your VCF after minimal initial configuration.</p> <p>In a VCF, you can have two to four members configured in the Routing Engine role. Of this, one member acts as the master Routing Engine and another member acts as the backup Routing Engine. In a preprovisioned configuration, the selection of which member functions as the master Routing Engine and which as the backup Routing Engine is determined by the software based on the master-election algorithm.</p> <p>In a configuration that is not preprovisioned, the selection of the master and backup is determined by the mastership priority value and secondary factors in the master-election algorithm.</p>
VC Mode	Indicates whether the Virtual Chassis is mixed or not.

The following details are displayed in the Virtual Chassis Topology view for each member depending on the role of the member as shown in [Table 343](#).

Table 343: Details of VC/VCF Members

Details	Description	Role
Name	Name of the member switch provided while configuring the device. The device name is displayed as a label.	Master Backup Line Card
Serial Number	Serial number of the member switch.	Master Backup Line Card
Platform	Platform of the device. Platform can be QFX3500, QFX3600, QFX5100 or QFX5110.	Master Backup Line Card
Priority	The mastership priority value. This is the most important factor in determining the role of the member switch within the Virtual Chassis configuration.	Master Backup Line Card

Table 343: Details of VC/VCF Members (continued)

Details	Description	Role
Operational Role	Operational role of the device. A device might be configured for a particular role, but can operate in the same or a different role. For example, a spine device configured with a Routing Engine role might operate as a line card. Therefore, the operational role of this device is Line Card. Operational role can be Routing Engine or Line Card.	Master Backup Line Card
Config Role	The configured role of the device. This can be Routing Engine or Line card.	Master Backup Line Card
Member Status	Displays the status of each member device: <ul style="list-style-type: none"> • Present—The device is connected and working fine. • Not Present—The device is not connected to the VC or VCF. • Inactive—The device is connected, but not running. • Non Provisioned—A configuration in which the roles of the members are assigned automatically; not configured statically (preprovisioned). • Pre Provisioned—A configuration that allows you to deterministically control the member ID and role assigned to a member by associating the member with its serial number. 	Master Backup Line Card

RELATED DOCUMENTATION

[Understanding the Monitor Mode Tasks Pane | 1277](#)

[Network Director Documentation home page](#)

Monitoring Virtual Networks

IN THIS CHAPTER

- [Using Monitor Mode for Virtual Devices | 1371](#)
- [Viewing vMotion History in Network Director | 1376](#)

Using Monitor Mode for Virtual Devices

IN THIS SECTION

- [Current Active Alarms Monitor | 1372](#)
- [Status Monitor | 1373](#)
- [Hosts By % Bandwidth Utilization | 1374](#)
- [Top VMs By Bandwidth Utilization | 1374](#)
- [Host NIC Bandwidth Utilization | 1375](#)
- [Virtual Switch Summary By Version | 1375](#)
- [Virtual Machine Bandwidth Utilization Trend | 1375](#)

The Monitor mode for virtual devices in your network enables you to view details about your virtual network using the following tabs:

- **Summary**—Displays the status of the virtual network, virtual machine, or virtual switch, active alarms, and the number of hosts and the version of VMware ESXi that is running on each host.
- **vMotion History**—vSphere vMotion is a feature that enables live migration of running virtual machines from one host to another with zero downtime and continuous network availability. You can view the status of the history of all the vMotions for your virtual network in the vMotion History tab. For more details, see [“Viewing vMotion History in Network Director” on page 1376](#).

Your current scope—that is, your view and node selection in the View pane—affects which Monitor widgets are available. For example, if you select a virtual switch, Network Director displays the status and the active alarms for the selected virtual switch.

This topic describes:

Current Active Alarms Monitor

The Current Active Alarms monitor shows any active alarm that has not yet been cleared. It is one of the four standard monitors available in Alarm mode. Current Active Alarms is a table that has four fields and appear by default. However, nine fields are available for selection. See [Table 344](#) for a description of the table.

Table 344: Current Active Alarms Monitor

Table Column	Description	Shown in Summary by Default	Shown in Detailed View by Default
Name	The alarm name.	Yes	Yes
ID	A system and sequentially generated identification number.	No	No
Assigned To	If assigned to an individual, it shows the name of the person assigned; otherwise, it shows System to mark that the alarm is still unassigned.	No	Yes
Severity	The severity of the alarm. Severity levels are: <ul style="list-style-type: none"> • Critical—A critical condition exists; immediate action is necessary. • Major—A major error has occurred; escalate or notify as necessary. • Minor—A minor error has occurred; notify or monitor the condition. • Info—An informational message; no action is necessary. 	Yes	Yes
Entity ID	The identification of the entity responsible for causing this alarm. The Entity ID is the key for correlation of events into an alarm. The Entity ID can be a MAC address of a radio or an IP address of the device.	Yes	Yes

Table 344: Current Active Alarms Monitor (continued)

Table Column	Description	Shown in Summary by Default	Shown in Detailed View by Default
Reporting Device IP	The IP address of the device or network element that generated the alarm. The SNMP agent is located at the source IP. In most cases, the source IP is the IP address of the switch or controller.	Yes	Yes
Reporting Device	The hostname or IP address of the reporting device.	Yes	Yes
Creation Date	The date and time the alarm was first reported.	No	No
Last Updated	The date and time that the information for the alarm was last modified.	Yes	Yes
Updated By	Either the system or the last user who modified the alarm.	No	No

Clicking the Details icon opens Alarm Details where you can sort and disposition alarms by state (Acknowledged, Clear, or Active).

Status Monitor

This monitor provides key information about the status for the virtual device selected in the Datacenter View. This monitor is on the Summary tab in Monitor mode.

[Table 345](#) describes the fields in this monitor.

Table 345: Status Monitor Fields

Field	Function	Scope
Name	The hostname of the virtual device.	All
Hardware	The hardware platform that the host uses.	Host
Connection Status	The connection status of the host and the virtual machines.	Host
Hypervisor	The version of VMware ESXi that is installed on the host.	Host
Number of Virtual Switches	The number of virtual switches available for the selected scope.	Virtual Network, Host

Table 345: Status Monitor Fields (continued)

Field	Function	Scope
Number of Virtual Machines	The number of virtual machines available for the selected scope.	All
Sync Status	Indicates whether the virtual network configuration is synchronized with Network Director.	Virtual Network
Orchestration Status	Displays the orchestration status of the virtual network.	Virtual Network
MTU	The maximum size of a protocol data unit that can be transmitted using the virtual switch.	Virtual Switch
Number of Port Groups	The number of port groups that are defined for the selected virtual switch. Port group is a template that stores a set of configuration that is used to create virtual switch ports on a virtual switch.	Virtual Switch

Hosts By % Bandwidth Utilization

The Hosts by % Bandwidth Utilization widget displays the hosts that use the maximum and minimum bandwidth in terms of percentage, for the selected context in the View pane. Click Top 5 in the filter list to see which hosts have the maximum bandwidth utilization in your virtual network. Likewise, you can click Bottom 5 in the filter list to see the hosts that have the minimum bandwidth utilization.

This widget is visible when you select the virtual network, a vCenter server, or the Hosts container from the View pane.

Click any host to view the virtual connectivity for the host. You can use the virtual connectivity view to monitor devices that the host is connected to, view connection between devices, and to view the bandwidth utilization at each device-level.

Top VMs By Bandwidth Utilization

The Top VMs By Bandwidth Utilization widget displays the virtual machines that are using up the maximum bandwidth for the selected context in the View pane. For instance, if you select the virtual network, this widget displays the top five virtual machines, from the entire virtual network, that are using the maximum bandwidth whereas if you select a host in the View pane, this widget displays the top five virtual machines that are part of the selected host.

This widget is visible when you select the virtual network, a vCenter server, or the Hosts container from the View pane.

Click any virtual machine to view the virtual connectivity for the virtual machine.

Host NIC Bandwidth Utilization

The Host NIC Bandwidth Utilization widget displays the percentage of bandwidth utilization for each physical NIC in a given host. Each bar in this bar chart represents a virtual NIC. You can mouse over each bar to see the name of the virtual NIC, percentage of bandwidth utilization, bandwidth that is currently used, and the bandwidth capacity of the virtual NIC.

This widget is visible when you select a host from the View pane.

Virtual Switch Summary By Version

The Virtual Switch Summary By Version displays the standard and distributed switch, and the software version that is installed on the switch, by using a pie chart. Mouse over a pie segment to view the actual number of devices and the percentage represented by that pie segment.

This widget is visible when you select the Virtual Switches container from the View pane.

Virtual Machine Bandwidth Utilization Trend

The Virtual Machine Bandwidth Utilization Trend page lists all the virtual machines that are part of the selected virtual network. For each virtual machine, you can view the graphical representation of the bandwidth utilization against time. You can use the Virtual Machine Bandwidth Utilization Trend to monitor traffic patterns, identify over utilized or under utilized virtual machines, and to optimize your virtual network.

This widget is visible when you select the Virtual Networks container or a Virtual Network from the View pane. If you access this widget after selecting the virtual networks container from the View pane, Network Director lists all the virtual machines in your virtual network spanning multiple vCenter servers, whereas selecting a virtual network from the View pane lists only those virtual machines that are part of the selected virtual network.

RELATED DOCUMENTATION

[Viewing vMotion History in Network Director | 1376](#)

[Network Director Documentation home page](#)

Viewing vMotion History in Network Director

vSphere vMotion is a feature that enables live migration of running virtual machines from one host to another with zero downtime and continuous network availability. vMotion is a key feature that enables the creation of a dynamic, automated and self-optimizing data center.

If a vMotion happens in any of the virtual machines that are under the management of Network Director, then Network Director initiates a job to track the vMotion and the corresponding changes to orchestration. You can view the status of the history of all the vMotions for your virtual network in the vMotion History page. You can also view the status of the orchestration job that is initiated because of this vMotion by clicking the Orchestration Job ID field.

After the orchestration job is completed successfully, you must manually resynchronize the physical switch's configuration by using Network Director. If the system of record (SOR) mode set for the Junos Space Network Management Platform is:

- Network as system of record (NSOR), then performing a resynchronization ensures that Junos Space automatically resynchronizes its configuration record to match the device configuration and sets the device configuration state to In Sync when the synchronization completes. For more details, see [“Resynchronizing Devices When Junos Space Is in NSOR Mode” on page 1222](#).
- Junos Space as system of record (SSOR), then you must perform a resynchronization and accept the out-of-band changes. Both the Junos Space configuration record and the Network Director Build mode configuration are resynchronized to reflect the out-of-band configuration changes. For more details, see [“Resynchronizing Devices When Junos Space Is in SSOR Mode” on page 1223](#).

To view the vMotion history:

1. While in the Monitor mode, select a Virtual Network.
2. Select the **vMotion History** tab.

The vMotion History page appears. You can view the details shown in [Table 346](#) in the vMotion History page.

Table 346: View vMotion History fields

Field	Description
VM Name	Name of the virtual machine that had undergone a vMotion.
vNetwork	Name of the virtual network.
Source Host	The host from which the virtual machine moved.
Destination Host	The host to which the virtual machine moved.

Table 346: View vMotion History fields (continued)

Field	Description
Started On	Time when the vMotion started.
Completed On	Time when the vMotion completed.
Status	Indicates the status of the vMotion.
Source Switches	Host name of the physical switch to which the host was connected before the vMotion.
Source Switch Port	Port on the source physical switch to which the host was connected.
Destination Switches	Host name of the physical switch to which the host is connected after the vMotion.
Destination Switch Port	Port on the destination physical switch to which the host is connected.
Orchestration Job IDs	The ID of orchestration job that was initiated as a result of the given vMotion. Click a job ID to view details about the orchestration job that got initiated as a result of the vMotion.
MAC Address	MAC address of the virtual machine.

3. You can check the status of the orchestration job corresponding to a given vMotion by clicking the orchestration job ID. Network Director opens the vMotion Orchestration window displaying job details such as the name, percentage complete, status, start and end time, and the summary of the job.

RELATED DOCUMENTATION

[Understanding Virtual Network Management | 785](#)

[Network Director Documentation home page](#)

General Monitoring

IN THIS CHAPTER

- Selecting Monitors To Display on the Summary Tab | 1378
- Changing Monitor Polling Interval and Data Collection | 1379
- Pinging Host Devices | 1379
- Troubleshooting Network Connections Using Traceroute | 1381

Selecting Monitors To Display on the Summary Tab

When you select the My Network node in the View pane, the Summary tab in Monitor mode enables you to select which monitors to display. If you select more than four monitors, a scroll bar appears to allow you to scroll to the additional monitors.

To select monitors to display on the Summary tab:

1. Click **Monitor** in the Network Director banner.
2. Select the **My Network** node in the View pane (the top node in the tree).
3. To select which monitors to display on the Summary tab:
 - a. Click **Select Monitors to Display** in the Tasks pane.

The Select Monitors window opens. The monitors that are already selected to display are listed in the Selected list. The other available monitors are listed in the Available list.

- b. To move a monitor from one list to the other list, click the monitor name, and then click the right or left arrow button, as appropriate.

- c. To change the order in which the selected monitors appear in the tab, select a monitor name and move it in the list using the up and down arrow buttons. The arrow buttons at the top and bottom of the stack of buttons move the selected monitor to the top or bottom of the list, respectively.
 - d. Click **Save** to save your changes, or click **Cancel** to cancel your changes.
4. To get information about a monitor, click the Help button in its title bar.

RELATED DOCUMENTATION

[Understanding Monitor Mode in Network Director | 1269](#)

[Network Director Documentation home page](#)

Changing Monitor Polling Interval and Data Collection

Network Administrators can change the default polling interval for monitors. The default polling period varies by monitor category. You can change these values in Preferences, found in the Network Director banner. You can also enable or disable the data collection processes used by monitors in Preferences.

RELATED DOCUMENTATION

[Setting Up User and System Preferences | 108](#)

[Network Director Documentation home page](#)

Pinging Host Devices

Use the Ping Host task in Monitor mode to determine whether an EX Series host can be reached over the network from the device selected in the network tree. Entering a hostname or an address creates a periodic ping task that sends a series of Internet Control Message Protocol (ICMP) echo (ping) requests to the specified host. The output of the task displays in the Response Console.

The Ping from Device to a Host task is available only for EX Series switches and QFX Series switches in your network.

1. Select either IP or HostName in the Remote Host Details box.

2. Type the IP address or hostname for the device that you want to reach.
3. Click **Ping** to use the default settings and start the requests or select the plus (+) symbol to use the Advanced Search Criteria. The fields in Advanced Search Criteria are described in [Table 347](#).

Table 347: Ping Host Advanced Search Criteria Field Descriptions

Field	Description	Default
Count	Indicates the number of ping requests to send. Valid values are 1 through 24.	5
Type of Service	Sets the type-of-service (ToS) field in the IP header of the ping packets. The range of values is 0 through 255. If the routing platform does not support ToS, the field is ignored.	0
Time To Live	Indicates the time-to-live hop count for the ping request packet. Valid values are 0 through 255.	0
Wait Interval	Indicates the amount of time in seconds between ping requests. Valid values are 0 through 24; a 0 value sends the request immediately.	1
Packet Size	Indicates the size of the ping request packet in bytes. The routing platform adds 8 bytes of ICMP header to this size before sending the request packet.	56
Interface	Sends the ping requests on the interface you specify. If you do not specify this option, ping requests are sent on all interfaces.	All
Source	Uses the source address that you specify in the ping request packet.	None

RELATED DOCUMENTATION

[Understanding Monitor Mode in Network Director | 1269](#)

[Network Director Documentation home page](#)

Troubleshooting Network Connections Using Traceroute

Traceroute is a diagnostic tool that enables you to display the route that a packet takes to reach the destination and measure transit delays of packets across an Internet Protocol (IP) network. You can use traceroute to troubleshoot and identify points of failure in your switching network. In traceroute, the source device sends three Internet Control Message Protocol (ICMP) echo request packets to the destination device. This is done sequentially till the source receives an ICMP echo reply message from the destination device. The time-to-live (TTL) value is used in determining the number of intermediate devices that the packets traverse before reaching the destination device.

You can use traceroute for EX Series switches (with or without ELS), QFX Series switches, and QFabric systems.

To start a traceroute from the selected device to another device in your network:

1. Select either IP or HostName in the Remote Host Details box.
2. Type the IP address or hostname for the device to which you want to start a traceroute.
3. Click **Trace** to use the default settings and start the traceroute or select the plus (+) symbol to use the Advanced Options. The fields in Advanced Options are described in [Table 348](#).

Table 348: Traceroute Advanced Options Field Descriptions

Field	Description	Default
Interface	Sends the Internet Control Message Protocol (ICMP) echo request packets on the interface you specify. If you do not specify this option, ICMP packets are sent on all interfaces.	Select a value from the list.
Time To Live	Indicates the time-to-live hop count for the ICMP echo request packets. Default value is 30. Valid values are 1 through 255.	30
Wait Interval	Indicates the amount of time in seconds between echo requests. Default value is 5. Valid values are 1 through 24.	5
Type of Service	Sets the type-of-service (ToS) field in the IP header of the echo packets. The range of values is 0 through 255. If the routing platform does not support ToS, the field is ignored.	0

RELATED DOCUMENTATION

[Understanding Monitor Mode in Network Director](#) | 1269

[Network Director Documentation home page](#)

Monitor Reference

IN THIS CHAPTER

- 802.11 Packet Errors Monitor | **1384**
- Access vs. Uplink Port Utilization Trend Monitor | **1385**
- AP Status Monitor | **1386**
- Current Sessions Monitor | **1388**
- Current Sessions by Type Monitor | **1388**
- Current SSID Statistics Monitor | **1389**
- Error Trend Monitor | **1391**
- Equipment Status Summary Monitor | **1393**
- Equipment Summary By Type Monitor | **1394**
- Node Device Summary Monitor | **1395**
- Percentage of Packets Retransmitted Monitor | **1396**
- Port Status Monitor | **1397**
- Port Status for IP Fabric Monitor | **1399**
- Port Utilization Monitor | **1399**
- Power Supply and Fan Status Monitor | **1400**
- QFabric Director Status Monitor | **1401**
- QFabric Interconnect Status Summary Monitor | **1402**
- QFabric VM Status Summary Monitor | **1403**
- Radio Status Monitor | **1403**
- Radio Technology Type Statistics Monitor | **1405**
- Resource Monitor For Wireless LAN Controllers | **1407**
- Resource Utilization Monitor for Switches, Routers, Virtual Chassis, Virtual Chassis Fabrics, and QFabric Systems | **1408**
- RF Interference Sources Monitor for Wireless Devices | **1409**
- RF Interference Sources Monitor For an Access Point | **1412**
- RF Throughput or Packet Throughput Level Monitor | **1413**
- Session Trends Monitor | **1415**
- Signal-to-Noise Ratio Monitor | **1418**
- SNR SSID Statistics Monitor | **1420**

- Status Monitor for Junos Fusion Systems | 1421
- Status Monitor for Layer 3 Fabrics | 1422
- Status Monitor for QFabric Directors | 1423
- Status Monitor for QFabric Systems | 1424
- Status Monitor for QFabric Interconnects | 1425
- Status Monitor for QFabric Nodes | 1425
- Status Monitor for Switches and Routers | 1426
- Status Monitor for Virtual Chassis | 1427
- Status Monitor for Virtual Chassis Members | 1428
- Status Monitor for Wireless Access Points | 1429
- Status Monitor for Wireless LAN Controllers | 1430
- Top APs by Session Monitor | 1431
- Top APs by Traffic Monitor | 1432
- Top Talker - Wireless Devices Monitor | 1433
- Top Talker - Wired Devices Monitor | 1435
- Top Users Monitor | 1436
- Top Sessions by MAC Address Monitor | 1438
- Traffic Trend Monitor | 1439
- Unicast vs Broadcast/Multicast Monitor | 1440
- Unicast vs Broadcast/Multicast Trend Monitor | 1441
- User Session Details Window | 1442
- Virtual Chassis Topology Monitor | 1443
- VC Equipment Summary By Type Monitor | 1445

802.11 Packet Errors Monitor

The 802.11 Packet Errors monitor displays the number of packet errors experienced by the object selected in the View pane. The object is polled and plotted at the standard polling rate.

You can perform the following actions on this graph:

- Change the time period over which to display the percentage of retransmitted packets by selecting a time period from the list in the upper right corner.
- Display a numeric value by mousing the cursor where a vertical grid line bisects a data line.

Packet error data is available when you select a radio, an access point, a floor, a building, or a site in any view. Radios and access points are displayed automatically in the View pane—you must configure floors, buildings, and sites. See [“Creating a Site” on page 228](#), [“Configuring Buildings” on page 229](#), and [“Configuring Floors” on page 231](#) for details.

If your packet error rate is too high, you can try to lower it by:

- Locating and eliminating noise that could be causing causes spurious packets. For more information about noise, see [“Monitoring RF Interference Sources on Wireless Devices” on page 1329](#), [“Monitoring RF Interference Sources on One Radio” on page 1325](#), and [“Monitoring RF Interference Sources For Radios on One Access Point” on page 1328](#).
- Checking for weak signals. If automatic power tuning is not enabled, try enabling it. For more information, see [“Understanding Auto Tune Power Policy for Wireless Radios” on page 866](#).
- Assigning access points to different channels. When channels are too congested, packet collision and corruption can occur. If automatic channel tuning is not enabled, try enabling it. For more information, see [“Understanding Wireless Radio Channels” on page 856](#) and [“Understanding Adaptive Channel Planner” on page 861](#).

RELATED DOCUMENTATION

[Monitoring RF 802.11 Packet Errors | 1323](#)

[Creating a Site | 228](#)

[Configuring Buildings | 229](#)

[Configuring Floors | 231](#)

[Monitoring RF Interference Sources on Wireless Devices | 1329](#)

[Monitoring RF Interference Sources on One Radio | 1325](#)

[Monitoring RF Interference Sources For Radios on One Access Point | 1328](#)

[Understanding Auto Tune Power Policy for Wireless Radios | 866](#)

[Understanding Wireless Radio Channels | 856](#)

[Understanding Adaptive Channel Planner | 861](#)

[Network Director Documentation home page](#)

Access vs. Uplink Port Utilization Trend Monitor

The Access vs. Uplink Port Utilization Trend monitor shows trends in the bandwidth utilization of access and uplink ports within the selected QFabric device or node device. It is available on the Summary tab in Monitor mode.

NOTE: After a device is discovered, trend data does not appear immediately—it appears only after three polling periods have occurred.

The information is shown in a line graph. The vertical axis shows bandwidth utilization percentage. The horizontal axis shows the times when data was polled. At each poll, the bandwidth utilization percentage of each port type (access and uplink) is indicated by a dot. The dots are connected by lines to show the trend over time.

You can perform the following actions on the line graph:

- Change the time period over which to display the traffic trends by selecting a time period from the list in the upper right corner.
- Highlight a line in the graph by mousing over the line's legend.
- Remove or restore a line by clicking its legend.
- Display a numeric value by mousing over where a vertical grid line crosses a data line.

RELATED DOCUMENTATION

[Understanding Monitor Mode in Network Director | 1269](#)

[Network Director Documentation home page](#)

AP Status Monitor

The AP Status monitor displays status information about all access points that belong to the selected node. It is found on the Equipment tab in Monitor mode.

The default view of the summary shows four of the eight available fields that you can configure to be shown or hidden. The details page shows an expanded version with all of the fields, which can also be tailored.

This monitor currently displays for the top wireless node and for wireless LAN controller nodes. If, for example, you select Wireless Network node, it shows global status information for all access points. If you select a wireless LAN controller, it shows status information for the access points belonging to the controller.

[Table 349](#) describes the fields that are available in both the summary and detail view of the AP Status monitor. Fields that are available, but hidden, are also displayed.

Table 349: AP Status Monitor Fields

Field	Description	Summary, Detailed, or Hidden View
AP Name	Name of the access point.	Summary
Serial Number	Serial number of the access point.	Summary
Model	The model number of the access point.	Summary (hidden) Detailed
IP Address	The IP address assigned to the AP.	Summary (hidden) Detailed
Status	Operational status of the access point: <ul style="list-style-type: none"> • Down—The access point is offline. • Up—The access point is online and enabled. • Up Redundant—The access point is online, reporting to this controller as redundant and another controller as primary. 	Summary
Uptime	The length of time since the access point last booted.	Summary
Version	The version of the Mobility System Software (MSS) running on the access point.	Summary (hidden) Detailed
Primary Controller	The primary controller for the access point.	Detailed
Secondary Controller	The secondary controller for the access point.	Detailed
Location	Location of the access point.	Summary (hidden) Detailed

RELATED DOCUMENTATION

[Monitoring the Status of Aggregated Access Points and Radios | 1354](#)

[Monitoring the Status of Wireless Controllers, Access Points, and Radios | 1356](#)

[Device Inventory Report | 1511](#)

[Network Director Documentation home page](#)

Current Sessions Monitor

Current Sessions monitor displays the number of active sessions at any given point of time. This monitor is available at the My Network level from the Summary tab.

This monitor displays the number of current active sessions on a pie chart.

RELATED DOCUMENTATION

[Current Sessions by Type Monitor | 1388](#)

[Network Director Documentation home page](#)

Current Sessions by Type Monitor

IN THIS SECTION

- [Current Sessions by Type | 1389](#)
- [Current Session Details | 1389](#)

The Current Sessions by Type monitor provides summary and detailed information about the active sessions within the node selected in the View pane. This monitor is available in the Client tab.

NOTE: If the selected scope is a single switch, this monitor is named Current Sessions by VLAN, and shows the distribution of current sessions by VLAN.

Current Sessions by Type

The summary view of the Current Sessions by Type monitor shows a pie chart of the active sessions within the node selected in the View pane. The chart shows the distribution of sessions by the session type. To change the session type shown in the monitor, select from the **Choose Sessions By Type** list.

Current Session Details

To see detailed information about the sessions in the Current Sessions monitor, click the **Details** button in the monitor title bar. The User Session Details window opens. For information about this window, see [“User Session Details Window” on page 1442](#).

RELATED DOCUMENTATION

[Monitoring Client Sessions | 1321](#)

[Network Director Documentation home page](#)

Current SSID Statistics Monitor

IN THIS SECTION

- [Current SSID Statistics Summary | 1389](#)
- [SSID Statistics Details | 1390](#)

The Current SSID Statistics monitor provides summary and detailed information about SSID statistics within the node selected in the View pane. This monitor is available on the Client tab.

Current SSID Statistics Summary

The summary view of the Current SSID Statistics monitor displays summary information about SSID statistics within the node selected in the View pane. Select the statistic to show from the Choose Statistic by SSID list box at the bottom of the monitor. [Table 350](#) describes the information shown for each statistic.

Table 350: Current SSID Statistics Summary Monitor

Statistic	Description
Number of Users	Shows a pie chart of the number of users per SSID. Mouse over the slices of the pie chart to see additional information. The total number of users on all SSIDs is shown below the chart.
Session Airtime (Last 10 min)	Shows a pie chart of the session airtime of each SSID. An SSID's session airtime is the total airtime of all sessions using the SSID during the polling period.
Bandwidth Usage	Shows a pie chart of the wireless network bandwidth usage per SSID. Mouse over the slices of the pie chart to see additional information. The total bandwidth is shown below the chart.
Number of Sessions	Shows a pie chart of the number of sessions per SSID. Mouse over the slices of the pie chart to see additional information. The total number of sessions on all SSIDs is shown below the chart.

SSID Statistics Details

To see detailed information about the current SSID statistics, click the **Details** button in the monitor title bar. [Table 351](#) describes the information in the SSID Statistics Detail window.

Table 351: SSID Statistics Detail Window

Column	Description
SSID Name	Name of the SSID.
Number of Users	Number of users on the SSID.
% of Users	Percentage of total wireless users that are using the SSID.
Session Airtime	Session airtime of the SSID. An SSID's session airtime is the total airtime of all sessions using the SSID during the polling period.
% of Session Airtime	Percentage of total session airtime used by the SSID.
B/W Usage	Bandwidth used by that SSID.
% of B/W Usage	Percentage of total wireless bandwidth used by the SSID.
Number of Sessions	Number of sessions on the SSID.

RELATED DOCUMENTATION

[Monitoring Client Sessions | 1321](#)

[Network Director Documentation home page](#)

Error Trend Monitor

IN THIS SECTION

- [Error Trend | 1391](#)
- [Error Trend Details | 1392](#)

The Error Trend monitor displays inbound and outbound error trends on the node you selected in the View pane. This monitor is available in the Traffic tab.

This topic describes:

Error Trend

A line graph shows the rate inbound and outbound errors over time. The horizontal axis shows the times when samples were taken. The vertical axis shows errors per second.

NOTE: After a device is discovered, trend data does not appear immediately—it appears only after three polling periods have elapsed.

You can perform the following actions on the line graph:

- Change the time period over which to display the traffic trends by selecting a time period from the list in the upper right corner.
- Highlight a line in the graph by mousing over the line's legend.
- Remove or restore a line by clicking its legend.
- Display a numeric value by mousing over a data point.

Error Trend Details

The Error Trend details window displays detailed information about errors on the node you selected in the View pane. It contains the following elements:

- A line graph shows the rate of errors over time. The horizontal axis shows the times when samples were taken. The vertical axis shows errors.

NOTE: After a device is discovered, trend data does not appear immediately—it appears only after three polling periods have elapsed.

You can perform the following actions on the line graph:

- Change the time period over which to display the traffic trends by selecting a time period from the list in the upper right corner.
- Highlight a line in the graph by mousing over the line's legend.
- Remove or restore a line by clicking its legend.
- Display a numeric value by mousing over a data point.
- Error Trend Details table—Shows detailed information about the data gathered at each sample. For information about this table, see [Table 352](#)
- Error Trend Additional Details table—Shows additional error trend details and enables you to display them on the graph. For information about this table, see [Table 353](#).

Table 352: Error Trend Details Table

Column	Description
Time	Time when a data sample was taken from devices.
Errors In	Number of inbound errors reported in the sample.
Errors Out	Number of outbound errors reported in the sample.
CRC Errors In	Number of inbound cyclic redundancy check (CRC) errors reported in the sample.
CRC Errors Out	Number of outbound CRC errors reported in the sample.

Table 353: Error Trend Additional Details Table

Column	Description
Series Name	Name of the data series.
Series Value	Value of the data series.
Show	Select the check box to display the series on the graph. Clear the check box to remove the series from the graph.

RELATED DOCUMENTATION

[Monitoring Traffic on Devices | 1283](#)

[Network Director Documentation home page](#)

Equipment Status Summary Monitor

The Equipment Status Summary monitor provides status highlights for the wireless controller ports, access points, and radios in the current scope. Both the summary and details show up to five available fields.

[Table 354](#) describes the fields in this monitor.

Table 354: Equipment Status Summary Fields

Field	Function	Default View
Device	Indicates the type of device being run by the wireless controller: access points, radios, and ports.	Summary Details
Up	Indicates how many of the devices are up.	Summary Details
Down	Indicates how many of the devices are down.	Summary Details
Unknown	Indicates if the controller cannot identify the device.	Summary Details
Disabled	Indicates if the device is disabled.	Summary Details

RELATED DOCUMENTATION

[Monitoring the Status of Wireless Controllers, Access Points, and Radios | 1356](#)

[AP Status Monitor | 1386](#)

[Radio Status Monitor | 1403](#)

[Network Director Documentation home page](#)

Equipment Summary By Type Monitor

IN THIS SECTION

- [Equipment Summary By Type | 1394](#)
- [Equipment Summary By Type Details | 1394](#)

The Equipment Summary By Type monitor provides summary and detailed information about the type and number of devices in the scope selected in the View pane. This monitor is available on the Summary tab in Monitor mode.

Equipment Summary By Type

The summary view of the Equipment Summary By Type monitor shows the distribution of device types in the selected scope. Switches in a Virtual Chassis are counted separately from standalone switches. Similarly, the count of satellite devices and aggregation devices in a Junos Fusion system are displayed separately in a pie chart.

Mouse over a segment of the pie chart to see the actual number of devices of that type. Click the details icon to open the Equipment Summary By Type Detail View window.

Equipment Summary By Type Details

The Equipment Summary By Type Detail View window provides details about the distribution of device types in the selected scope. Each table row represents a device type. Device types are defined by the combination of a device family, platform, and operating system version (for some device types). See [Table 355](#) for a description of the table columns.

Table 355: Equipment Summary By Type Detail View

Table Column	Description
Device Family	Device family.
Platform	Device platform.
OS Version	Operating system version running on the device.
Device Type	Device type.
Count	Number of devices of this platform in the selected scope.

RELATED DOCUMENTATION

[Selecting Monitors To Display on the Summary Tab | 1378](#)

[Network Director Documentation home page](#)

Node Device Summary Monitor

The Node Device Summary monitor displays information about the port utilization of the nodes within the selected QFabric fabric or container within a fabric. It is on the Summary tab in Monitor mode. The information is presented in a bar chart. The vertical axis shows node names. The horizontal axis shows the number of ports. Ports are categorized based on the percentage of allocated bandwidth they use: over 80%, between 50-80%, and below 50%. The bar color codes for the categories are shown in the legend below the chart. The five nodes that are using the highest percentage of their bandwidth are shown on the monitor.

RELATED DOCUMENTATION

[Understanding Monitor Mode in Network Director | 1269](#)

[Network Director Documentation home page](#)

Percentage of Packets Retransmitted Monitor

The Percentage of Packets Retransmitted monitor displays the percentage of wireless data packet retransmissions experienced by the access point or radio selected in the Network Director View pane.

You can perform the following actions on this line graph:

- Change the time period over which to display the percentage of retransmitted packets by selecting a time period from the list in the upper right corner.
- Display a numeric value by placing the cursor where a vertical grid line bisects a data line.

Ideally, packet retransmission does not exceed 10% of the total number of packets sent. If your retransmission percentage is higher, you can try to lower it by:

- Locating and eliminating noise that could be causing causes spurious packets. For more information on noise, see [“Monitoring RF Interference Sources on Wireless Devices” on page 1329](#), [“Monitoring RF Interference Sources on One Radio” on page 1325](#), and [“Monitoring RF Interference Sources For Radios on One Access Point” on page 1328](#).
- Checking for weak signals. Automatic power tuning will help improve weak signals. For more information, see [“Understanding Auto Tune Power Policy for Wireless Radios” on page 866](#).
- Assigning access points to different channels. When channels are too congested, packet collision and corruption can occur. Enable automatic channel tuning enabled if it is not already enabled. For more information, see [“Understanding Wireless Radio Channels” on page 856](#) and [“Understanding Adaptive Channel Planner” on page 861](#).

Retransmitted packet data is available when you either select a radio or an access point.

RELATED DOCUMENTATION

[Monitoring RF Interference Sources on Wireless Devices | 1329](#)

[Monitoring RF Interference Sources on One Radio | 1325](#)

[Monitoring RF Interference Sources For Radios on One Access Point | 1328](#)

[Understanding Auto Tune Power Policy for Wireless Radios | 866](#)

[Understanding Wireless Radio Channels | 856](#)

[Understanding Adaptive Channel Planner | 861](#)

[Network Director Documentation home page](#)

Port Status Monitor

IN THIS SECTION

- [Port Status Summary | 1397](#)
- [Port Status Details | 1397](#)

The Port Status monitor provides summary and detailed information about the status of the physical network interfaces for the selected node in the View pane.

If the selected node represents an individual device, the monitor displays data specific to the ports on the device. If the selected node contains multiple devices, the monitor displays data aggregated from all the ports on all the devices.

This topic describes:

Port Status Summary

The summary view of the Port Status monitor displays two pie charts:

- Admin Status—Of the interfaces on the selected node, shows the proportion of interfaces that are administratively enabled and that are administratively disabled.
- Free vs Used—Of the network interfaces that are administratively enabled, shows the proportion of interfaces that are in use (operationally up) and that are not in use (operationally down).





Mouse over a pie segment to view the actual number of ports. Click the details icon to open the Port Status Details window.

Port Status Details

The Port Status Details table provides details about the physical network interfaces for the selected node, as shown in [Table 356](#).

NOTE: You must have a transceiver installed in an SFP, SFP+, or XFP port for information about the port to appear.

Table 356: Port Status Details Table

Field	Description
Port Name	The name of the physical interface.
MAC Address	<p>For standalone EX Series switches, the first five groups of hexadecimal digits are determined when the switch is manufactured. The switch then assigns a unique MAC address to each interface by assigning a unique identifier as the last group of hexadecimal digits.</p> <p>For Virtual Chassis members, the first four groups of hexadecimal digits are determined when the switch is manufactured. The fifth group of hexadecimal digits reflects the role of the member in the chassis, such as master or linecard.</p>
Serial Number	The hardware serial number of the device.
Host Name	The hostname of the device.
Description	A text description of the physical interface.
Current Negotiated Speed (Mbps)	The actual operating speed of the port, in megabits per second (Mbps). Depending on the results of autonegotiation, this speed might be less than the maximum speed supported by the port as indicated by port type.
Configured Speed	The speed configured for the port. If the speed is configured to be determined by autonegotiation, the configured speed is shown as Auto.
Duplex Mode	The duplex mode: full (full-duplex), half (half-duplex), or auto (autonegotiation).
Port Type	The port type (for example, 1 Gigabit Ethernet or 10 Gigabit Ethernet interface).
Admin Status	Indicates the administrative state of the port as  UP or  DOWN.
Operational Status	Indicates the operational status of the port as  UP or  DOWN.
PoE	<p>Indicates whether the PoE traps are enabled for the port. Possible values are:</p> <ul style="list-style-type: none"> • Enabled—PoE traps are generated for the port • Disabled—PoE traps are not generated for the port • N/A—PoE traps are not applicable for the port
Last Flap Time	Date and time at which the advertised link became unavailable, and then, available again.

RELATED DOCUMENTATION

[Monitoring the Status of Virtual Chassis Members | 1358](#)

[Network Director Documentation home page](#)

Port Status for IP Fabric Monitor

The Port Status for IP Fabric monitor provides summary information about the status of the physical network interfaces in the Layer 3 Fabric selected in the View pane. The status information the monitor provides is organized by the device role in the fabric (spine or leaf) and by port role (access or uplink).

To use the monitor, first select either **Spines** or **Leaves** and then **Access Ports** or **Uplink Ports**. The information displayed is based on your selections and consists of the following:

- A table that lists the number of each type of port—for example, the number of 10 Gigabit Ethernet ports.
- Two pie charts that display:
 - The administrative status of the physical interfaces—that is the proportion of interfaces that are administratively enabled versus administratively disabled.
 - The operational status of the physical interfaces—that is, of the interfaces that are administratively enabled, the proportion of interfaces that are in use (operationally up) versus not in use (operationally down).

RELATED DOCUMENTATION

[Status Monitor for Layer 3 Fabrics | 1422](#)

[Network Director Documentation home page](#)

Port Utilization Monitor

The Port Utilization Monitor displays a bar chart with information about the port traffic utilization on the node selected in the View pane. Each bar in the chart represents the port traffic utilization data gathered at a polling interval. The vertical axis shows the number of ports polled. The horizontal axis shows the time when each poll was taken. The data shown in the graph is aggregated from all the ports contained in the node selected in the View pane.

Each bar is divided into the following colored sections to indicate the distribution of port traffic utilization at the polling interval:

- Green indicates ports that operated at less than 50% of negotiated speed.
- Yellow indicates ports that operated at between 50% and 80% of negotiated speed.
- Red indicates ports that operated at more than 80% of negotiated speed.

You can perform the following actions on the bar chart:

- Change the time period over which to display the data by selecting a time period from the list in the upper right corner.
- Remove or restore a utilization category (bar color) by clicking its legend.
- Display a numeric value by mousing over a bar.

RELATED DOCUMENTATION

[Monitoring Port Utilization | 1300](#)

[Network Director Documentation home page](#)

Power Supply and Fan Status Monitor

IN THIS SECTION

- [Power Supply and Fan Status | 1401](#)
- [Power Supply and Fan Status Details | 1401](#)

The Power Supply and Fan Status monitor provides information about the availability and status of power supplies and cooling fans for the node you select in the View pane.

This monitor is available when you select a switch, router, or a Virtual Chassis member in any view. It appears on the Equipment tab when in Monitor mode.

This topic describes:

Power Supply and Fan Status

The summary view of the Power Supply and Fan Status monitor displays two pie charts:

- **Power Supply Units**—On the node selected, shows the proportion of power supplies that are detected as absent against those that are present in the bay. Those units that are present indicate their operating status as OK, Check, or Failed. The totals shown below the title indicate the total number of power supplies that the device is capable of holding.
- **Fans**—On the node selected, shows the proportion of fans that are detected against those that are present in the bay. Those units that are present indicate their operating status as OK, Check, or Failed. The totals shown below the title indicate the total number of fans that the device is capable of holding.

Mouse over the pie segments to view the number of power supplies or fans in each segment. The total number of units is shown at the bottom of the graph. Click the details icon to open the Power Supply and Fan Status Details window.

Power Supply and Fan Status Details

The Power Supply and Fan Status Details window provides a tabular status view of each power supply and fan in the device.

The top table lists all of the power supplies available in the device. The chart shows the individual status of each power supply as OK, Absent, Check, or Failed.

The lower table lists the fans in the device. The chart shows the status of each fan unit as OK, Absent, Check, or Failed.

RELATED DOCUMENTATION

[Monitoring the Status of a Virtual Chassis | 1357](#)

[Monitoring the Status of Virtual Chassis Members | 1358](#)

[Network Director Documentation home page](#)

QFabric Director Status Monitor

The QFabric Director Status monitor shows the status of the QFabric director devices within the selected Directors folder in the View pane, in a table. It is on the Summary tab in Monitor mode. [Table 357](#) describes the table columns.

Table 357: Status Monitor for QFabric Directors Table

Column	Description
Name	Device name.
Member ID	Member ID.
Status	Device status.
Role	Device role.
Uptime	Length of time the device has been up.

RELATED DOCUMENTATION

[Understanding Monitor Mode in Network Director | 1269](#)

[Network Director Documentation home page](#)

QFabric Interconnect Status Summary Monitor

The QFabric Interconnect Status Summary monitor shows the status of the selected QFabric fabric's interconnect devices in a table. It is on the Summary tab in Monitor mode. [Table 358](#) describes the table columns.

Table 358: QFabric Interconnect Status Summary Monitor Table Description

Column	Description
Name	Device name.
Status	Device status.
Port Utilization %	Percentage of device's allocated bandwidth that is being used.
Uptime	Length of time the device has been up.

RELATED DOCUMENTATION

[Understanding Monitor Mode in Network Director | 1269](#)

[Network Director Documentation home page](#)

QFabric VM Status Summary Monitor

QFabric VM Status Summary Monitor shows the status of the virtual machines (VMs) within the QFabric director device selected in the View pane, in a table. It is on the Summary tab in Monitor mode. [Table 359](#) describes the table columns.

Table 359: QFabric VM Status Summary Monitor Table

Column	Description
VM Name	Name of the VM.
CPU Utilization %	Percentage of CPU the VM is using.

RELATED DOCUMENTATION

[Understanding Monitor Mode in Network Director | 1269](#)

[Network Director Documentation home page](#)

Radio Status Monitor

The Radio Status monitor, on the Equipment tab in Monitor mode, provides information about the radios and access points being controlled at this node. For example if you select Wireless Network, information about all of the radios is displayed; if you select a wireless controller, then only information for the radios for that controller is shown. The monitor has a summary view and a detailed view. The default view is the summary, showing four of the six available fields. These fields in the Radio Status monitor are described in [Table 360](#).

Table 360: Radio Status Monitor Fields

Field	Description	Summary, Detailed, or Hidden View
Radio Identifier	The system identification for the radio, composed of: <ul style="list-style-type: none"> • The access point name • A colon (:) • The radio number 	Summary Detailed
MAC Address	The radio MAC address.	Summary (Hidden) Detailed
Status	Up (enabled), Down (disabled), or NA (unable to determine status).	Summary Detailed
Radio Type	The type of wireless clients that can connect to the access point.	Summary Detailed
Channel	The configured operating channel for AP communication.	Summary Detailed
Tx Power	The transmit power level for a radio.	Summary (Hidden) Detailed

RELATED DOCUMENTATION

[Monitoring the Status of Aggregated Access Points and Radios | 1354](#)

[Monitoring the Status of Wireless Controllers, Access Points, and Radios | 1356](#)

[Device Inventory Report | 1511](#)

[Network Neighborhood Report | 1516](#)

[Network Director Documentation home page](#)

Radio Technology Type Statistics Monitor

IN THIS SECTION

- [Radio Technology Type Statistics Summary | 1405](#)
- [Radio Technology Type Statistics Details | 1406](#)

The Radio Technology Type Statistics monitor provides summary and detailed information about the usage of radio technology types within the node selected in the View pane. This monitor is available on the Summary tab in Monitor mode.

Radio Technology Type Statistics Summary

The summary view of the Radio Technology Type Statistics monitor shows summary information about the usage of radio technology types within the node selected in the View pane. Select the category of information to show from the Tech Type Params By list box at the bottom of the monitor. [Table 361](#) describes the information shown for each category.

Table 361: Radio Technology Type Statistics Summary Categories

Category	Description
Sessions by Technology Type	Shows a pie chart of the distribution of radio technologies among the active sessions. You can mouse over the slices of the pie chart to see additional information.
Users by Technology Type	Shows a pie chart of the distribution of radio technologies among the active users. A user is an account that authenticates to get network access. A user can have multiple active sessions by using multiple network devices with the same account. You can mouse over the slices of the pie chart to see additional information.
Average SNR by Technology Type	Shows a bar chart of the signal-to-noise ratio (SNR) of the current sessions. Select the radio technology type to show from the list to the left of the chart. The chart shows the number of sessions on the vertical axis and the SNR on the horizontal axis. The average SNR of all sessions on the chart is shown below the chart.
Bandwidth Usage by Technology Type	Shows a pie chart of the percentage of wireless bandwidth used by each radio technology. Mouse over the slices of the pie chart to see additional information.
Amount of Time (last 10 min)	Shows a bar chart of the amount of time each radio technology type has been continuously active on the network.

Radio Technology Type Statistics Details

To see detailed information about the radio technology type statistics, click the **Details** button in the monitor title bar. [Table 362](#) describes the information in the The Radio Type Statistics window.

Table 362: Radio Type Statistics Window

Column	Description
Tech Type	Radio technology type.
No. of Sessions	Number of sessions using the radio technology type.
No. of Users	Number of users using the radio technology type.
Users Percentage	Percentage of users using the radio technology type.
Bandwidth Usage	Bandwidth used by the radio technology type.
Bandwidth Percentage	Percentage of bandwidth used by the radio technology type.
Time Amount	Amount of time during the polling period when at least one session using that radio type was active. The maximum amount is the length of the polling period.
Time Percentage	Percentage of time during the polling period when at least one session using that radio type was active.
Average SNR	Average SNR for the radio technology type.

RELATED DOCUMENTATION

[Monitoring Client Sessions | 1321](#)

[Network Director Documentation home page](#)

Resource Monitor For Wireless LAN Controllers

IN THIS SECTION

- [Resource Utilization Summary | 1407](#)
- [CPU and Memory Utilization Charts | 1407](#)

The Resource Utilization monitor shows wireless LAN controller CPU and memory use for the last hour using two needle gauges. The gauges display usage from 1-100 percent as resources are consumed, making it easy to see if these resources are being under or overused.

This monitor is available when you select a wireless LAN controller. It appears on the Equipment tab when in Monitor mode.

This topic describes:

Resource Utilization Summary

The summary view of the Resource Utilization monitor displays two needle gauges:

- **CPU Usage**—For the selected node, it displays the device's CPU usage within the last hour marked in tenths from 0 to 100 percent. Under the gauge, the total percent usage is also shown.
- **Memory Usage**—For the selected node, it displays the device's memory consumption within the last hour marked in tenths from 0 to 100 percent. Under the gauge, the total percent usage is also shown.

To view the resource utilization over a period of 24 hours, a month, or a year, click the details icon. The CPU Utilization and Memory Utilization charts provide a more granular presentation of the data.

CPU and Memory Utilization Charts

The CPU and Memory Utilization charts allow you to view consumption rates for over different periods of time. You can select the time period from a list or specify a custom value. If you select Custom, an additional dialog box opens enabling you to select a starting and ending date and time.

The data is presented in two line charts or graphs, one for CPU utilization and the other for memory utilization. Select a time-frame for analysis from the list to update both graphs. Depending on the time-frame selected, the charts refresh to display time increments proportional to the time-frame. For example, when 1 Hour is selected, the time increments are 5 minutes apart; when 1 day is selected, the time increments are 1 hour apart. Mouse over the data intersections to view the precise value at that point.

RELATED DOCUMENTATION

[Monitoring the Status of Wireless Controllers, Access Points, and Radios | 1356](#)

[Network Director Documentation home page](#)

Resource Utilization Monitor for Switches, Routers, Virtual Chassis, Virtual Chassis Fabrics, and QFabric Systems

IN THIS SECTION

- [Resource Utilization Summary | 1408](#)
- [Resource Utilization Details | 1409](#)

The Resource Utilization monitor shows a line chart for CPU and memory use for a switch, router, Virtual Chassis, Virtual Chassis Fabric, and QFabric. The vertical axis shows the percentage of the resource being consumed. The horizontal axis shows the times when samples were taken. The time period that the chart represents can be selected from a list.

This monitor appears on the Equipment tab in Monitor mode.

This topic describes:

Resource Utilization Summary

The summary view of the Resource Utilization monitor shows a line chart representing memory and CPU use. There are six possible categories shown on the chart, depending on which device type is selected:

CPU User—(Orange) the percentage of time that the CPU uses to run user processes, such as the database.

CPU System—(Green) the percentage of time that the CPU uses on all processes for the system.

CPU Background—(Light Blue) the percentage of time that the CPU uses on background processes.

CPU Interrupt—(Red) the percentage of time that the CPU uses for interrupt handling.

CPU Idle—(Purple) the percentage of time that the CPU is available for work.

5 min Mem avg—(Dark Blue) the amount of memory being used over a 5-minute average.

You can interact with the chart to manipulate the data being displayed by:

- Mouse over a line's legend to highlight the line.
- Removing or restoring a line by clicking the legend item.
- Displaying specific chart values by mousing over a datapoint.
- Changing the time period that the chart covers.

If you select Custom, an additional dialog box opens, enabling you to select a starting and ending date and time.

Resource Utilization Details

In Resource Utilization Details, you can view utilization rates for memory and CPU over different periods of time. You can select the time period from a list or specify a custom value. If you select Custom, an additional dialog box, enabling you to select a starting and ending date and time.

The data is presented in two line charts or graphs, one for memory utilization and the other for CPU utilization. Select a timeframe for analysis from the list to update both graphs. Depending on the timeframe selected, the charts refresh to display time increments proportional to the timeframe. For example, if you select 1 Hour, the time increments are 5 minutes apart; if you select 1 Day, the time increments are 1 hour apart. Mouse over a data point to view the precise value at that point.

RELATED DOCUMENTATION

[Resource Monitor For Wireless LAN Controllers | 1407](#)

[Network Director Documentation home page](#)

RF Interference Sources Monitor for Wireless Devices

The RF Interference Sources monitor for wireless devices consists of a summary pie chart that reflects all wireless traffic experienced by the object selected in the View pane. You can select any one of the objects listed in [Table 363](#) in the view pane:

Table 363: Wireless Objects With Interference Tracking










Icon	Object
	Entire Wireless Network in any view.
	Wireless Mobility Domain in any view.

Table 363: Wireless Objects With Interference Tracking (continued)

Icon	Object
	Controller Cluster in any view. NOTE: You cannot see interference for a single controller.
	Individual access point in any view.
	Individual radio in any view.
	Selecting a floor in logical view displays all access points on that floor—to create a floor, see “Configuring Floors” on page 231 .
	Selecting a building in logical view displays all access points in that building—to create a building, see “Configuring Buildings” on page 229 .
	Selecting a site from the logical view displays all access points in that building—to create a site, see “Creating a Site” on page 228 .
	Wiring closet—to create a wiring closet, see “Setting Up Closets” on page 233 .

Network Director tracks and monitors interference from these sources:

- Microwave ovens—Most domestic microwave ovens use 2.45 GHz, and can interfere with Wi-Fi channels from 8 to 10 (or even 7 to 11). Interference varies depending on the model of the oven—for example, commercial restaurant microwave ovens sweep over a wider spectrum and have a higher duty cycle.
- Continuous wave devices continuously transmit at a particular frequency without attempting to share the radio frequency medium with other devices. Devices that use continuous wave technology in the same frequency bands as wireless LAN networks will interfere with wireless communications, reducing performance or totally preventing communication. Several examples of devices that may use continuous wave transmission that interferes with Wi-Fi are video surveillance cameras and baby monitors.
- Bluetooth devices
- Phone FHSS from cordless phones
- Unknown devices

To track these devices, Network Director polls the controllers at the standard interval. The categories with the largest sections of the pie chart cause the most interference.

You can perform the following actions on the pie chart:

- Change the time period over which to display interference by selecting a time period from the list in the upper right corner.

- Display a numeric value for interference occurrences by mousing over a section of the chart.
- Click the monitor's title to see a list of interfering objects along with the information listed in [Table 364](#).

Table 364: Information on RF Interference Sources for a Radio

Information	Description
Last Seen	Date and time the interference was last detected.
Transmitter ID	If the interference is caused by an object with a MAC address, the MAC address is displayed. If the object has no MAC address, MSS calculates a MAC address, using the characteristics of the object. This way, you can correlate interference events over time.
Listener MAC	MAC address of the access point that detected the interference.
AP	Name of the access point that detected the interference.
Controller	Name of the controller that reported the interference.
Channel	Channel the interference affected.
RSSI	Received signal strength indication (RSSI), in decibels referred to 1 milliwatt (dBm). A higher value indicates a stronger signal.
Duty Cycle	Reported fraction of time that the source is emitting RF.
Source Type	Possible sources of interference include Bluetooth, Continuous Wave, Microwave Oven, Unknown, and Phone FHSS.
CIM (%)	Estimated severity of interference on this channel caused by the source.

Interference is frequently not a problem on wireless networks with light traffic, but as traffic becomes heavier, throughput and capacity decrease and other problems become apparent. RF interference can cause packet retransmission (see [“Monitoring the Percentage of RF Packet Retransmissions” on page 1338](#)). Interference is also a security concern because jamming can bring down the network .

Ideally, interference retransmission does not cause more than 10% of the total number of packets sent. If your retransmission percentage is higher, you can try to lower it by:

- Locating and eliminating offending devices. If the item cannot be removed, you can add electromagnetic interference (EMI) shielding such as grounded mesh, foils, insulating foams, or insulating paint. This will limit the interference to a small area.
- Moving clients to channels with less interference. Keep in mind, however, that Bluetooth devices, cordless phones, 802.11FH devices, and jamming emissions are broadband, so it's not possible to change channels

away from them—they are everywhere in the band. For more information, see [“Understanding Wireless Radio Channels”](#) on page 856 and [“Understanding Adaptive Channel Planner”](#) on page 861.

For more information about wireless interference, see [“Understanding Wireless Interference”](#) on page 914.

RELATED DOCUMENTATION

[Monitoring RF Interference Sources on One Radio | 1325](#)

[Monitoring RF Interference Sources For Radios on One Access Point | 1328](#)

[Understanding Wireless Interference | 914](#)

[Network Director Documentation home page](#)

RF Interference Sources Monitor For an Access Point

The RF Interference Sources monitor for single access points consists of a bar chart that reflects interference experienced with the traffic of one or both radios on the access point selected in the View pane. Some access points have two radios and some access points have one radio. Network Director tracks and monitors radio interference from these sources:

- Microwave ovens—Most domestic microwave ovens use 2.45 GHz, and can interfere with Wi-Fi channels from 8 to 10 (or even 7 to 11). Interference varies depending on the model of the oven—for example, commercial restaurant microwave ovens sweep over a wider spectrum and have a higher duty cycle.
- Continuous wave devices continuously transmit at a particular frequency without attempting to share the radio frequency medium with other devices. Devices that use continuous wave technology in the same frequency bands as wireless LAN networks will interfere with wireless communications, reducing performance or totally preventing communication. Several examples of devices that may use continuous wave transmission that interferes with WiFi are video surveillance cameras and baby monitors.
- Bluetooth devices
- Phone FHSS from cordless phones
- Unknown devices

To track these interference devices, Network Director polls the access point’s controller at the standard interval. The categories with the largest bars in the chart cause the most interference.

You can perform the following actions on the bar chart:

- Change the time period over which to display interference by selecting a time period from the list in the upper right corner.
- Display a numeric value by mousing over a bar in the chart.

- Add or remove one or both radio's data from the chart by clicking **Radio 1** or **Radio 2** in the legend.

Interference is frequently not a problem on wireless networks with light traffic, but as traffic becomes heavier, throughput and capacity decrease and other problems become apparent. RF interference can cause packet retransmission (see [“Monitoring the Percentage of RF Packet Retransmissions” on page 1338](#)). Interference is also a security concern because jamming can bring down the network.

Ideally, interference retransmission does not cause more than 10% of the total number of packets sent. If your retransmission percentage is higher, you can try to lower it by:

- Locating and eliminating offending devices. If the item cannot be removed, you can add electromagnetic interference (EMI) shielding such as grounded mesh, foils, insulating foams, or insulating paint. This will limit the interference to a small area.
- Moving the affected access point.
- Moving clients to channels with less interference. Keep in mind, however, that Bluetooth devices, cordless phones, 802.11FH devices, and jamming emissions are broadband, so it's not possible to change channels away from them—they are everywhere in the band. For more information, see [“Understanding Wireless Radio Channels” on page 856](#) and [“Understanding Adaptive Channel Planner” on page 861](#).

For more information about wireless interference, see [“Understanding Wireless Interference” on page 914](#).

RELATED DOCUMENTATION

[Monitoring RF Interference Sources For Radios on One Access Point | 1328](#)

[Monitoring RF Interference Sources on Wireless Devices | 1329](#)

[Monitoring RF Interference Sources on One Radio | 1325](#)

[Understanding Wireless Interference | 914](#)

[Understanding Wireless Radio Channels | 856](#)

[Understanding Adaptive Channel Planner | 861](#)

[Network Director Documentation home page](#)

RF Throughput or Packet Throughput Level Monitor

The RF Throughput or Packet Throughput Level monitor displays the amount of data throughput in kilobytes per second (KBps) experienced in the last hour by the object selected in the View pane. Total network throughput for each radio or for each access point (either one can be selected) is measured in KBps at the configured polling interval and plotted on this line chart. The throughput rates are reflected on the left side of the chart.

If you are viewing data for a time period longer than one hour, each data point on the graph represents data consolidated from more than one polling period. In this case, the graph shows multiple lines, which are labeled in the legend:

- maxThroughput—The largest value sampled during the consolidated polling periods.
- avgThroughput—The average of the values sampled during the consolidated polling periods.
- minThroughput—The smallest value sampled during the consolidated polling periods.

The area between the maxThroughput and minThroughput lines is shaded to indicate the range of values.

You can perform the following actions on this line graph:

- Change the time period over which to display throughput by selecting a time period from the list in the upper right corner.
- Display a numeric value by placing the cursor where a vertical grid line bisects a data line.
- Highlight a line in the graph by mousing over the line's legend.
- Remove or restore a line by clicking its legend.
- Select which SSID to show in the graph by selecting from the **Choose SSID** list.

Throughput data is available when you select either a radio or access point in any view of the View pane.

Throughput is decreased by Layer 2 retransmissions, increased numbers of clients, and the overhead associated with 802.11 protocols. You can try to increase throughput by:

- Adding equipment such as controllers and access points to cope with over-subscription.
- Using optimal configuration, such as WPA2 encryption, for 802.11n devices.
- Configuring separate WLAN Service profiles for voice and data—for data, see [“Creating and Managing a WLAN Service Profile” on page 1090](#). For directions on creating a voice-specific WLAN Service profile, see [“Configuring a Voice SSID with Network Director” on page 1124](#) and [“Creating and Managing a WLAN Service Profile” on page 1090](#).
- Creating separate Radio profiles for transmissions using long and short guard intervals—see [“Creating and Managing a Radio Profile” on page 932](#).
- Locating and eliminating noise that could be causing interference. For more information, see [“Monitoring RF Interference Sources on Wireless Devices” on page 1329](#), [“Monitoring RF Interference Sources on One Radio” on page 1325](#), [“Monitoring RF Interference Sources For Radios on One Access Point” on page 1328](#), and [“Monitoring RF Signal-to-Noise Ratio” on page 1334](#).
- Checking for weak signals. If automatic power tuning is not enabled, try enabling it. For more information, see [“Understanding Auto Tune Power Policy for Wireless Radios” on page 866](#).
- Assigning access points to different channels. When channels are too congested, packet collision and corruption can occur. If automatic channel tuning is not enabled, try enabling it. For more information,

see “Understanding Wireless Radio Channels” on page 856 and “Understanding Adaptive Channel Planner” on page 861.

- Correcting conditions that trigger alarms - for a list of alarms, see the “Current Active Alarms Monitor” on page 1464.

RELATED DOCUMENTATION

[Creating and Managing a WLAN Service Profile | 1090](#)

[Configuring a Voice SSID with Network Director | 1124](#)

[Creating and Managing a WLAN Service Profile | 1090](#)

[Creating and Managing a Radio Profile | 932](#)

[Monitoring RF Interference Sources on Wireless Devices | 1329](#)

[Monitoring RF Interference Sources on One Radio | 1325](#)

[Monitoring RF Interference Sources For Radios on One Access Point | 1328](#)

[Monitoring RF Signal-to-Noise Ratio | 1334](#)

[Understanding Auto Tune Power Policy for Wireless Radios | 866](#)

[Understanding Wireless Radio Channels | 856](#)

[Understanding Adaptive Channel Planner | 861](#)

[Current Active Alarms Monitor | 1464](#)

[Network Director Documentation home page](#)

Session Trends Monitor

IN THIS SECTION

● [Session Trends | 1416](#)

● [Session Details | 1416](#)

The Session Trends monitor provides summary and detailed trend information about the number of active sessions and users within the node selected in the View pane.

NOTE: After a device is discovered, trend data does not appear immediately—it appears only after three polling periods have occurred.

Session Trends

The summary view of the Session Trends monitor displays a line graph of the number of active sessions and users over time within the node selected in the View pane. The vertical axis is the number of active sessions or users. The horizontal axis shows the polling interval times.

You can perform the following actions on the line graph:

- Change the time period over which to display the trend data by selecting a time period from the list in the upper right corner.
- Select which SSID to monitor from the **Choose SSID** list.
- Select which VLAN to monitor from the **Choose VLAN** list. This option appears only when you have selected a switch in the View pane.
- Display the number of sessions or users at a polling interval by mousing over the plotted data point.
- Highlight a line in the graph by mousing over the line's legend.
- Remove or restore a line by clicking its legend.

Because data points plotted against the x-axis can represent data consolidated from multiple polling periods, three lines are plotted for session count and for user count: the maximum, minimum, and average counts that occurred during the consolidated polling periods.

Session Details

The Session Details window provides detailed trend information about the number of active sessions and users within the current node selected in the View pane. It contains these panes:

- The top pane contains a line graph of the number of active sessions and users over time within the node selected in the View pane.

You can perform the following actions on the line graph:

- Change the time period over which to display the trend data by selecting a time period from the list in the upper right corner.
- Display the number of sessions or users at a polling interval by mousing over the plotted data point.
- Highlight a line in the graph by mousing over the line's legend.
- Remove or restore a line by clicking its legend.

Because data points plotted against the x-axis can represent data consolidated from multiple polling periods, three lines are plotted for both session count and user count: the maximum, minimum, and average session and user counts that occurred during the consolidated polling periods.

- The bottom pane contains a table with detailed information about the active sessions.

The following table describes the columns that appear in current session details tables.

Table 365: User Session Details Table

Table Column	Description
User Name	Client's user name
MAC Address	Client's MAC address.
Device Type	Client's device type.
Device Group	Client's device group.
Device Profile	Client's device profile.
Controller IP	IP address of the controller to which the client is connected.
AP ID	ID of the wireless access point to which the client is connected.
AP NAME	Name of the wireless access point to which the client is connected.
SSID	SSID to which the wireless client is connected.
VLAN	VLAN to which the client is connected.
Client IP	Client's IP address.
Auth Type	Authorization type used for the client.
B/w[KBps]	Bandwidth used by the client.
Data Usage (KBytes)	Data transmitted and received by the client, in kilobytes.
Elapsed Time	Length of time the session has been active.
Sample Time	Time when the most recent sample was taken.
RSSI	Received signal strength indication (RSSI). Specified in decibels referred to 1 milliwatt (dBm). A higher value indicates a stronger signal.

Table 365: User Session Details Table (continued)

Table Column	Description
Roam In Time	The time when the session roamed in to the wireless access point.

TIP: Some table columns are hidden by default. To select which columns to display, mouse over any column heading, click the arrow that appears, mouse over **Columns** in the drop-down menu, and then select the columns to display from the list.

RELATED DOCUMENTATION

[Monitoring Client Sessions | 1321](#)

[Network Director Documentation home page](#)

Signal-to-Noise Ratio Monitor

IN THIS SECTION

- [Monitoring Signal-to-Noise Ratio | 1418](#)
- [Signal-to-Noise Ratio Details | 1419](#)

Monitoring Signal-to-Noise Ratio

Signal-to-noise ratio (SNR) is a measure of the level of a desired signal against the level of background noise, measured in decibels (dB). You can imagine this as a person trying to be heard in a noisy restaurant, where his voice is the signal and the background chatter blocks his voice. The SNR charts display the ratio between signal and background noise, the individual signal level, and the individual background noise level.

If you are viewing data for a time period longer than one hour, each data point on the graph represents data consolidated from more than one polling period. In this case, the graph shows multiple lines, which are labeled in the legend:

- maxSnr—The largest value sampled during the consolidated polling periods.

- snr—The average of the values sampled during the consolidated polling periods.
- minSnr—The smallest value sampled during the consolidated polling periods.

The area between the maxSnr and minSnr lines is shaded to indicate the range of values.

You can perform the following actions on the SNR chart:

- Change the time period over which to display the SNR by selecting a time period from the list in the upper right corner.
- Display a numeric value by placing the cursor where a vertical grid line bisects a data line.
- Highlight a line in the graph by mousing over the line's legend.
- Remove or restore a line by clicking its legend.
- Click **Details** to view these three charts for a radio, one above the other:
 - Received signal strength indicator (RSSI), which measures the power of a radio signal
 - Signal to noise ratio (SNR) which measures a signal against the current RF background noise level
 - Noise floor which is the sum of all the noise interference for the radio—the current RF background noise level

Higher numbers on this chart indicate that a radio has more signal than noise, which is desirable. If the chart has more noise than signal (indicated by values less than 40 on the chart), the signal becomes more unreadable, because the noise level severely competes with it. A reading of 0-20 on this chart would indicate an unacceptable level of noise or a really low signal. This can cause a reduction in data speed because of frequent errors that require the source transmitter to resend data packets—see [“Monitoring the Percentage of RF Packet Retransmissions” on page 1338](#).

SNR is computed only for individual radios.

Signal-to-Noise Ratio Details

There are three charts in the Signal-to-Noise Ratio (SNR) Details window: RSSI, SNR, and Noise Floor. [Table 366](#) briefly describes how these charts can be interpreted.

Table 366: Interpreting Signal-to-Noise Ratio Values

	SNR	RSSI	Noise Floor
Definition	Signal-to-noise ratio is the ratio of a signal's strength to the sum of all interference. (Signal-to-noise Ratio = RSSI/Noise Floor).	RSSI is signal strength, the first value used in the signal-to-noise ratio.	Noise is any signal (interference) that is not Wi-Fi traffic such as cordless phones, microwaves, radar, etc. This is the second value in the signal-to-noise ratio.

Table 366: Interpreting Signal-to-Noise Ratio Values (continued)

	SNR	RSSI	Noise Floor
How is it measured?	SNR is the ratio of signal to background noise, measured as a positive value between 0 dB and 80 dB. You want the signal to be high and the background noise to be low. This produces a higher ratio, which is better.	RSSI is measured in decibels from -20 through -100.	Noise floor is measured in decibels from -90 through -120.
What does the chart mean?	If the chart has more noise than a signal (indicated by values less than 40 on the chart), the signal becomes more unreadable, because the noise level severely competes with it. A reading of 0-20 on this chart would indicate an unacceptable level of noise.	A louder signal is better, so the higher the RSSI is, the better. Typically voice networks require a better signal level than a data network does. Normal signal strength in a network would be around -45 dB through -87 dB.	A quiet noise floor is better, so the closer to -120 the noise floor is, the better because that means there is little to no interference. Typical environment noise floors are about 95 dB.

Deal with a low SNR reading by either increasing the signal (RSSI) or reducing the background noise. To get an idea what an acceptable SNR reading is for your network, check the values when the network is operating optimally—you might be able to do this by changing the time period on the chart.

RELATED DOCUMENTATION

[Monitoring RF Signal-to-Noise Ratio | 1334](#)

[Monitoring the Percentage of RF Packet Retransmissions | 1338](#)

[Network Director Documentation home page](#)

SNR SSID Statistics Monitor

IN THIS SECTION

- [SNR SSID Statistics Summary | 1421](#)
- [SNR SSID Statistics Details | 1421](#)

The SNR SSID Statistics monitor provides summary and detailed information about signal-to-noise ratio (SNR) statistics within the node selected in the View pane. This monitor is available in the Client tab.

SNR SSID Statistics Summary

The summary view of the Current SSID Statistics monitor displays a bar chart of SNR statistics within the node selected in the View pane. The number of sessions appears on the vertical axis. The SNR appears on the horizontal axis. Each horizontal axis unit represents a range of SNR values. Select an SSID to display from the **Choose SSID** list below the chart, or select **All** to see all SSIDs.

SNR SSID Statistics Details

To see detailed information about SNR SSID statistics, click the **Details** button in the monitor title bar. [Table 367](#) describes the information in the SNR SSID Details window.

Table 367: SNR SSID Details Window

Column	Description
SSID Name	Name of the SSID.
SNR	SNR for the session.
MAC Address	MAC address of the client device.
User Name	User name of the session.
AP Name	Name of the wireless access point to which the session is connected.

RELATED DOCUMENTATION

[Monitoring Client Sessions | 1321](#)

[Network Director Documentation home page](#)

Status Monitor for Junos Fusion Systems

The Status monitor for Junos Fusion systems provides key information about the status of equipment in a Junos Fusion system and is available on the Equipment tab in Monitor mode.

[Table 368](#) describes the fields in this monitor.

Table 368: Status Monitor for Junos Fusion System Fields

Field	Function
Power Supply Status	Indicates the aggregated power supply status for devices in the Junos Fusion system. Power supplies are categorized as absent , OK , check , or failed and the number of power supplies in each category are given.
FAN Status	Indicates the aggregated fan status for devices in the Junos Fusion system. Fans are categorized as absent , OK , check , or failed and the number of fans in each are given.
Used MAC Addresses	Indicates the number of MAC addresses in use in the Junos Fusion system.
Used VLANs	Indicates the number of VLANs in use in the Junos Fusion system.
Status	Indicates the number of devices in the Junos Fusion system that are up and that are down.
Alarm Severity Status	Displays the highest severity of any alarms active on any device in the Junos Fusion system and the number of alarms at that severity level.

RELATED DOCUMENTATION

[Network Director Documentation home page](#)

Status Monitor for Layer 3 Fabrics

The Status monitor for Layer 3 Fabrics provides key information about the status of equipment in a Layer 3 Fabric and is available on the Equipment tab in Monitor mode.

[Table 369](#) describes the fields in this monitor.

Table 369: Status Monitor for Layer 3 Fabric Fields

Field	Function
Power Supply Status	Indicates the aggregated power supply status for devices in the Layer 3 Fabric. Power supplies are categorized as absent, OK, check, or failed and the number of power supplies in each category are given.

Table 369: Status Monitor for Layer 3 Fabric Fields (continued)

Field	Function
FAN Status	Indicates the aggregated fan status for devices in the Layer 3 Fabric. Fans are categorized as absent, OK, check, or failed and the number of fans in each are given.
Used MAC Addresses	Indicates the number of MAC addresses in use in the Layer 3 Fabric.
Used Vlans	Indicates the number of VLANs in use in the Layer 3 Fabric.
Status	Indicates the number of devices in the Layer 3 Fabric that are up and that are down.
Alarm Severity Status	Displays the highest severity of any alarms active on any device in the Layer 3 Fabric and the number of alarms at that severity level.

RELATED DOCUMENTATION

[Port Status for IP Fabric Monitor | 1399](#)

[Network Director Documentation home page](#)

Status Monitor for QFabric Directors

The Status monitor for QFabric Directors shows the status of the selected QFabric Director in a table. It is on the Summary tab in Monitor mode. [Table 370](#) describes the fields in this monitor.

Table 370: Status Monitor for QFabric Systems Directors Fields

Field	Function
Name	Device name.
Member ID	Member ID.
Status	Device status.
Uptime	Length of time device has been running.
Role	Device role.

Table 370: Status Monitor for QFabric Systems Directors Fields (continued)

Field	Function
IP Address	Device IP address.
VMs	Number of virtual machines (VMs) running on the device.

RELATED DOCUMENTATION

[Understanding Monitor Mode in Network Director | 1269](#)

[Network Director Documentation home page](#)

Status Monitor for QFabric Systems

The Status monitor for QFabric Systems shows the status of the selected QFabric fabric in a table. It is on the Summary tab in Monitor mode. [Table 371](#) describes the fields in this monitor.

Table 371: Status Monitor for QFabric Systems Fields

Field	Function
Serial Number	Indicates the hardware serial number of the fabric.
IP Address	Indicates the IP address of the fabric.
Uptime	Indicates the amount of time since the last boot of the fabric in days, hours, minutes, and seconds.
Status	Indicates whether the fabric is up or down.
Used MAC Addresses	Indicates the number of MAC addresses in use on the fabric.
Used VLANs	Indicates the VLAN memberships for this fabric.
Temperature Range (°C)	Indicates the ambient temperature of the coldest and hottest devices in the fabric (in degrees Celsius).
Junos Version	Indicates the version and release level of Junos OS running on the fabric.

RELATED DOCUMENTATION

[Understanding Monitor Mode in Network Director | 1269](#)

[Network Director Documentation home page](#)

Status Monitor for QFabric Interconnects

The Status monitor for QFabric Interconnects shows the status of the selected QFabric Interconnect in a table. It is on the Equipment tab in Monitor mode. [Table 372](#) describes the fields in this monitor.

Table 372: Status Monitor for QFabric Systems Interconnects Fields

Field	Function
Serial Number	Device serial number.
Model	Device model.
Status	Device status.
Uptime	Length of time device has been running.
Temperature	Device temperature, in degrees Celsius.

RELATED DOCUMENTATION

[Understanding Monitor Mode in Network Director | 1269](#)

[Network Director Documentation home page](#)

Status Monitor for QFabric Nodes

The Status monitor for QFabric Nodes shows the status of the selected QFabric Node in a table. It is on the Equipment tab in Monitor mode. [Table 373](#) describes the fields in this monitor.

Table 373: Status Monitor for QFabric Systems Nodes Fields

Field	Function
Serial Number	Device serial number.

Table 373: Status Monitor for QFabric Systems Nodes Fields (*continued*)

Field	Function
Model	Device model
Status	Device status.
Uptime	Length of time device has been running.
Temperature	Device temperature, in degrees Celsius.

RELATED DOCUMENTATION

[Understanding Monitor Mode in Network Director | 1269](#)

[Network Director Documentation home page](#)

Status Monitor for Switches and Routers

This monitor provides key information about the status for a standalone switch or a router when the device is selected in any of the views. This monitor is on the Equipment tab in Monitor mode.

Table 374 describes the fields in this monitor.

Table 374: Status Monitor Fields

Field	Function
Serial Number	Indicates the hardware serial number of the device.
IP Address	Indicates the IP address of the device.
Uptime	Indicates the amount of time since the last boot of the unit in days, hours, minutes, and seconds.
Status	Indicates whether the device is up or down.
Used MAC Addresses	Indicates the number of MAC addresses in use on the device.
Used VLANs	Indicates the number of VLAN memberships for this device.

Table 374: Status Monitor Fields (continued)

Field	Function
Last Configured Time	Indicates the date and time when the device was last configured.
Temperature (°C)	Indicates the ambient temperature (in degrees Celsius).
Junos Version	Indicates the version and release level of Junos OS running on the device.

RELATED DOCUMENTATION

[Network Director Documentation home page](#)

Status Monitor for Virtual Chassis

This monitor provides status information, including power supply and fan information, for a Virtual Chassis. It is on the Equipment tab in Monitor mode.

The Summary view shows key status fields in a table format. Power supply and fan data is represented as small bar chart entries in the table. The Details view also shows the same status information, but expands the power supply and fan information. [Table 375](#) displays these fields and their location in the monitor.

Table 375: Virtual Chassis Status Monitor Fields

Field	Function	Location
Serial Number	Indicates the hardware serial number of the master member.	Summary Detailed
IP Address	Indicates the IP address of the master member.	Summary Detailed
Uptime	Indicates the amount of time since the last boot of the system in days, hours, minutes, and seconds.	Summary Detailed
Status	Indicates whether the Virtual Chassis is up or down.	Summary Detailed
Used MAC Addresses	Indicates the number of MAC addresses in use on the Virtual Chassis.	Summary Detailed

Table 375: Virtual Chassis Status Monitor Fields (continued)

Field	Function	Location
Used VLANs	Indicates the VLAN memberships for the Virtual Chassis.	Summary Detailed
Last Configured Time	Indicates the date and time since the Virtual Chassis was last configured.	Summary Detailed
Temperature Range (°C)	Indicates the temperature of the coolest and hottest devices in the Virtual Chassis (in degrees Celsius).	Summary Detailed
Junos Version	Indicates the version and release level of Junos OS running on the device.	Summary Detailed
Power Supply Status	Indicates the number of power supplies that are detected as absent or present in the bay. The graphic bar and total count for missing and present power supplies is shown as OK, Check, or Failed.	Summary
Power Supply Status	In the Power Supply and Fan Status table, there are separate table entries for each power supply state with the totals for that state.	Detailed
Fan Status	Indicates the number of cooling fans that are detected as absent or present in the bay. The graphic bar and total count for missing and present fans is shown as OK, Check, or Failed.	Summary
Fan Status	In the Power Supply and Fan Status table, there are separate table entries for each power supply state with the totals for that state.	Detailed

RELATED DOCUMENTATION

[Monitoring the Status of a Virtual Chassis | 1357](#)

[Monitoring the Status of Virtual Chassis Members | 1358](#)

[Network Director Documentation home page](#)

Status Monitor for Virtual Chassis Members

Use the Member Status monitor to view key information about the status of Virtual Chassis members. It is displayed on the Equipment tab in Monitor mode when you select a Virtual Chassis member.

[Table 376](#), describes the fields in this monitor.

Table 376: Status Monitor for Members Fields

Field	Description
Serial Number	Indicates the hardware serial number of the member.
Member ID	Identifies by number a member switch in a Virtual Chassis.
Member Serial Number	Indicates the hardware serial number of the member.
FPC Slot	Identifies the Flexible PIC Concentrator (FPC) slot number for the member: same as Member Slot.
Member Model	The model number of the member.
Member Mixed Mode	Indicates whether the switch is configured to run in mixed member mode. Valid fields are true or false.
Member Role	Indicates the function and responsibility of the switch in the Virtual Chassis. Possible values are master, backup, and linecard.

RELATED DOCUMENTATION

[Monitoring the Status of Virtual Chassis Members | 1358](#)

[Network Director Documentation home page](#)

Status Monitor for Wireless Access Points

This monitor provides key information about the status for the wireless access point selected in any of the views. This monitor is on the Summary tab in Monitor mode.

[Table 377](#) describes the fields in this monitor.

Table 377: Status Monitor Fields

Field	Description
AP Name	Name of the access point.
AP Number	Number of the access point.

Table 377: Status Monitor Fields (continued)

Field	Description
Model	The model number of the access point.
Serial Number	Serial number of the access point.
IP Address	The IP address assigned to the access point.
Uptime	The length of time since the access point last booted.
status	Operational status of the access point: <ul style="list-style-type: none"> • Down—The access point is offline. • Up—The access point is online and enabled. • Up Redundant—The access point is online, reporting to this controller as redundant and another controller as primary.
Version	The version of the Mobility System Software (MSS) running on the access point.
Primary Controller	The primary controller for the access point.
Secondary Controller	The secondary controller for the access point.
Location	Location of the access point.

RELATED DOCUMENTATION

[Monitoring the Status of Wireless Controllers, Access Points, and Radios | 1356](#)

[Network Director Documentation home page](#)

Status Monitor for Wireless LAN Controllers

The Status monitor for wireless LAN controllers provides status highlights for the wireless controller. View [Table 378](#) for a description of the fields in the monitor.

This monitor is available on the Equipment tab when you select a wireless controller from any view while in Monitor mode.

Table 378: Wireless Controller Status Fields

Field	Function
Serial Number	Indicates the hardware serial number of the master member.
IP Address	Indicates the IP address of the master member.
Uptime	Indicates the amount of time since the last boot of the system in days, hours, minutes, and seconds.
Status	Indicates whether the controller is up or down.
MSS Version	Indicates the version and release level of Mobility System Software running on the device.

RELATED DOCUMENTATION

[Monitoring the Status of Wireless Controllers, Access Points, and Radios | 1356](#)

[Network Director Documentation home page](#)

Top APs by Session Monitor

IN THIS SECTION

- [Top APs by Session Summary | 1431](#)
- [Top APs by Session Details | 1432](#)

The Top APs by Session monitor provides summary and detailed information about the wireless access points with the most active sessions within the node selected in the View pane. This monitor is available in the Client tab.

Top APs by Session Summary

The summary view of the Top APs by Session monitor displays a bar chart of the wireless access points with the most active sessions within the node selected in the View pane. The wireless access points are shown on the vertical axis. The number of sessions is shown on the horizontal axis.

Top APs by Session Details

To see detailed information about the top wireless access points by sessions, click the **Details** button in the monitor title bar. [Table 379](#) describes the information in the Top APs by Sessions window.

Table 379: Top APs by Sessions Window

Column	Description
AP Name	Wireless access point name.
Serial Number	Wireless access point serial number.
WLC Controller	Wireless controller that controls the wireless access point.
Location	Location of the wireless access point.
Number of Sessions	Number of active sessions on the wireless access point.
Bandwidth (KBytes)	Wireless access point bandwidth.

RELATED DOCUMENTATION

[Monitoring Client Sessions | 1321](#)

[Network Director Documentation home page](#)

Top APs by Traffic Monitor

IN THIS SECTION

- [Top APs by Traffic Summary | 1433](#)
- [Top APs by Traffic Details | 1433](#)

The Top APs by Traffic monitor provides summary and detailed information about the wireless access points with the most traffic within the node selected in the View pane. This monitor is available on the Summary tab.

Top APs by Traffic Summary

The summary view of the Top APs by Traffic monitor displays a bar chart of the wireless access points with the most traffic within the node selected in the View pane. The wireless access points are shown on the vertical axis. The traffic is shown on the horizontal axis.

Top APs by Traffic Details

To see detailed information about the top wireless access points by traffic, click the **Details** button in the monitor title bar. [Table 380](#) describes the information in the Top APs by Traffic Details window.

Table 380: Top APs by Traffic Details Window

Column	Description
AP Name	Wireless access point name.
Serial Number	Wireless access point serial number.
WLC Controller	Wireless controller that controls the wireless access point.
Location	Location of the wireless access point.
Number of Sessions	Number of active sessions on the wireless access point.
Bandwidth (KBytes)	Wireless access point bandwidth.

RELATED DOCUMENTATION

[Monitoring Client Sessions | 1321](#)

[Network Director Documentation home page](#)

Top Talker - Wireless Devices Monitor

IN THIS SECTION

- [Top Talker-Wireless Devices Summary | 1434](#)
- [Top Talker-Wireless Devices Details | 1434](#)

The Top Talker-Wireless Devices widget provides summary and detailed information about the top client device types that are generating wireless network traffic. Device types correspond to the platform or operating system of the device, for example, Windows or Android.

This topic describes:

Top Talker-Wireless Devices Summary

The summary view of the Top Talker-Wireless Devices monitor has a bar chart showing summary information about the top client device types that are generating wireless network traffic. Device types are listed on the vertical axis. Data usage in kilobytes is shown on the horizontal axis. You can mouse over a bar to see more information about it, including the number of devices of that type.

Top Talker-Wireless Devices Details

To see detailed information about the Top Talker-Wireless Devices, click the **Details** button in the monitor title bar. The Top Talker-Wireless Devices monitor details window has a table containing detailed information about the top client device types that are generating wireless network traffic. [Table 381](#) describes the columns in the table.

Table 381: Top Talker-Wireless Devices Details Window

Column	Description
Number of Device(s)	Number of devices of the device type.
Device Type	Device type.
Data Usage (KBytes)	Data used by devices of the device type, in kilobytes.

RELATED DOCUMENTATION

[Monitoring Client Sessions | 1321](#)

[Network Director Documentation home page](#)

Top Talker - Wired Devices Monitor

IN THIS SECTION

- [Top Talker - Wired Devices Summary | 1435](#)
- [Top Talker - Wired Devices Details | 1435](#)

The Top Talker - Wired Devices monitor provides summary and detailed information about the wired devices that are using the most bandwidth.

Top Talker - Wired Devices Summary

The summary view of the Top Talker - Wired Devices monitor has a bar chart that shows summary information about the wired devices that are using the most bandwidth. Device names or addresses are listed on the vertical axis. Data usage in kilobytes is shown on the horizontal axis. You can mouse over a bar to see more information about that device.

Top Talker - Wired Devices Details

To see detailed information about the top talkers, click the **Details** button in the monitor title bar. The Top Talker - Wired Devices monitor details window has a table containing detailed information about the devices that are using the most bandwidth. [Table 382](#) describes the columns in the table. To close the details page, click the **Minimize** button in the title bar.

Table 382: Top Hosts Monitor Details

Column	Description
Host Name	Host's host name.
MAC Address	Host's MAC address
Data Usage (KBytes)	Data used by the host, in kilobytes.
Device Serial Number	Device's serial number.

RELATED DOCUMENTATION

Top Users Monitor

IN THIS SECTION

- [Top Users | 1436](#)
- [Top Session By User Details | 1436](#)

The Top Users monitor provides summary and detailed information about the users within the node you selected in the View pane that use the most bandwidth. This monitor is available on the Client tab.

This monitor includes only wireless network users, not users on wired connections. If the node you select in the View pane contains only wired users, this monitor does not appear. If the node contains both wired and wireless users, only the wireless users appear in the monitor.

This topic describes:

Top Users

The summary view of the Top Users monitor displays a bar chart of the top bandwidth users within the node you selected in the View pane. The vertical axis shows the user names. The horizontal axis shows different information depending on the time period you select in the list in the title bar:

- If you select the **Current** time period, the horizontal axis shows the user's incremental data usage.
- If you select any time period other than **Current**, the horizontal axis shows the user's total data usage.

You can mouse over a bar to see more information about that user. You can select which SSID to monitor from the **Choose SSID** list.

Top Session By User Details

The Top Session By User Details window displays detailed information about the top users within the node you selected in the View pane.

To change the number of top users displayed, select a number from the Top *N* list in the upper right corner.

To change the time period for which data is shown, select a time period from the time period list. By default, the Current time period is selected. When Current is selected, the data from the most recent polling period is shown. When any time period other than Current is selected, the data for the entire selected time period is shown.

The following table describes the columns that appear in Top Session By Users Details table.

Table 383: Top Session Details Table

Table Column	Description
User Name	Client's user name To copy the text in this table cell: Click the cell, highlight the text, right-click the cell, and select Copy from the context menu.
MAC Address	Client's MAC address. To copy the text in this table cell: Click the cell, highlight the text, right-click the cell, and select Copy from the context menu.
Number of Sessions	Number of sessions.
AP Name	Name of the wireless access point to which the client is connected.
AP ID	ID of the wireless access point client is connected to.
Incremental Data Usage (KBytes)	The session's current incremental data usage. Appears only when the Current time period is selected.
Total Data Usage (KBytes)	The session's total data usage. Appears when any time period other than Current is selected.

RELATED DOCUMENTATION

[Monitoring Client Sessions | 1321](#)

[Network Director Documentation home page](#)

Top Sessions by MAC Address Monitor

IN THIS SECTION

- [Top Sessions | 1438](#)
- [Top Session by MAC Details | 1438](#)

The Top Sessions by MAC Address monitor provides summary and detailed information about the sessions within the node you selected in the View pane that use the most bandwidth. This monitor is available in the Client tab.

This monitor includes only wireless network sessions, not sessions on wired connections. If the node you selected in the View pane contains only wired sessions, this monitor does not appear. If the node contains both wired and wireless sessions, only the wireless sessions appear in the monitor.

This topic describes:

Top Sessions

The summary view of the Top Sessions by MAC Address monitor displays a bar chart of the sessions within the node you selected in the View pane that consume the most bandwidth. The vertical axis shows the session MAC addresses. The horizontal axis shows different information depending on the time period you select in the list in the title bar:

- If you select the **Current** time period, the horizontal axis shows the session's incremental data usage.
- If you select any time period other than **Current**, the horizontal axis shows the session's total data usage.

You can mouse over a bar to see more information about that session. You can select which SSID to monitor from the **Choose SSID** list.

Top Session by MAC Details

The Top Session by MAC Details window displays detailed information about the top sessions within the node you selected in the View pane.

To change the number of top sessions displayed, select a number from the Top N list in the upper right corner.

To change the time period for which data is shown, select a time period from the time period list. By default, the Current time period is selected. When Current is selected, the data from the most recent polling period

is shown. When any time period other than Current is selected, the data for the entire selected time period is shown.

The following table describes the columns that appear in Top Sessions by MAC Details table.

Table 384: Top Session Details Table

Table Column	Description
User Name	Client's user name To copy the text in this table cell: Click the cell, highlight the text, right-click the cell, and select Copy from the context menu.
MAC Address	Client's MAC address. To copy the text in this table cell: Click the cell, highlight the text, right-click the cell, and select Copy from the context menu.
Number of Sessions	Number of sessions.
AP Name	Name of the wireless access point to which the client is connected.
AP ID	ID of the wireless access point client is connected to.
Incremental Data Usage (KBytes)	The session's current incremental data usage. Appears only when the Current time period is selected.
Total Data Usage (KBytes)	The session's total data usage. Appears when any time period other than Current is selected.

RELATED DOCUMENTATION

[Monitoring Client Sessions | 1321](#)

[Network Director Documentation home page](#)

Traffic Trend Monitor

The Traffic Trend monitor displays inbound and outbound traffic trends on the node you selected in the View pane. This monitor is available in the Traffic tab. A line graph shows the rate of each type of traffic over time. The horizontal axis shows the times when samples were taken. The vertical axis shows the data rate, in packets per second.

NOTE: After a device is discovered, trend data does not appear immediately—it appears only after three polling periods have elapsed.

You can perform the following actions on the line graph:

- Change the time period over which to display the traffic trends by selecting a time period from the list in the upper right corner.
- Highlight a line in the graph by mousing over the line's legend.
- Remove or restore a line by clicking its legend.
- Display a numeric value by mousing over a data point.

RELATED DOCUMENTATION

[Monitoring Traffic on Devices | 1283](#)

[Network Director Documentation home page](#)

Unicast vs Broadcast/Multicast Monitor

The Unicast vs Broadcast/Multicast monitor displays a pie chart of the current distribution of unicast, broadcast, and multicast traffic types on the node you selected in the View pane. This monitor is available in the Traffic tab.

The traffic is divided into these categories:

- Unicast inbound
- Unicast outbound
- Broadcast inbound
- Broadcast outbound
- Multicast inbound
- Multicast outbound

Mouse over a pie segment to view the actual number of packets.

RELATED DOCUMENTATION

Unicast vs Broadcast/Multicast Trend Monitor

The Unicast vs Broadcast/Multicast Trend monitor displays trends in the data rates of unicast, broadcast, and multicast traffic on the node you selected in the View pane. This monitor is available on the Traffic tab. A line graph shows the rate of each type of traffic over time. The horizontal axis shows the times when samples were taken. The vertical axis shows the data rate, in packets per second.

NOTE: After a device is discovered, trend data does not appear immediately—it appears only after three polling periods have elapsed.

The traffic is divided into these categories:

- Unicast inbound
- Unicast outbound
- Broadcast inbound
- Broadcast outbound
- Multicast inbound
- Multicast outbound

You can perform the following actions on the line graph:

- Change the time period over which to display the traffic trends by selecting a time period from the list in the upper right corner.
- Highlight a line in the graph by mousing over the line's legend.
- Remove or restore a line by clicking its legend.
- Display a numeric value by mousing over a data point.

RELATED DOCUMENTATION

User Session Details Window

The User Session Details window provides information about the active sessions within the node selected in the View pane. To open this window, click the **Details** button in the Current Sessions or Current Sessions by Type monitors.

The following table describes the columns that appear in the user session details table:

Table 385: User Session Details Table

Table Column	Description
User Name	Client's user name
MAC Address	Client's MAC address.
Device Type	Client's device type.
Device Group	Client's device group.
Device Profile	Client's device profile.
Controller IP	IP address of the controller to which the client is connected.
AP ID	ID of the wireless access point to which the client is connected.
AP NAME	Name of the wireless access point to which the client is connected.
SSID	SSID to which the wireless client is connected.
VLAN	VLAN to which the client is connected.
Client IP	Client's IP address.
Auth Type	Authorization type used for the client.
B/w[KBps]	Bandwidth used by the client.
Data Usage (KBytes)	Data transmitted and received by the client, in kilobytes.
Elapsed Time	Length of time the session has been active.
Sample Time	Time when the most recent sample was taken.

Table 385: User Session Details Table (continued)

Table Column	Description
RSSI	Received signal strength indication (RSSI). Specified in decibels referred to 1 milliwatt (dBm). A higher value indicates a stronger signal.
Roam In Time	The time when the session roamed in to the wireless access point.

TIP: Some table columns are hidden by default. To select which columns to display, mouse over any column heading, click the arrow that appears, mouse over **Columns** in the drop-down menu, and then select the columns to display from the list.

RELATED DOCUMENTATION

[Monitoring Client Sessions | 1321](#)

[Network Director Documentation home page](#)

Virtual Chassis Topology Monitor

The Virtual Chassis Topology monitor provides a fast and simple way to view the members and their relationships. It is available on the Equipment tab in Monitor mode. View [Table 386](#) for a description of the fields in the monitor.

The summary shows up to five available fields that you can configure to be displayed or hidden. The details page shows an expanded version with up to eleven fields that can also be tailored.

Table 386: Virtual Chassis Topology Fields

Field	Function	Default in Topology Monitor
Member	Identifies by member ID a member switch in a Virtual Chassis.	Summary (hidden) Details (hidden)
Member Role	Indicates the function and responsibility of the switch in the Virtual Chassis. Possible values are master, backup, and linecard.	Summary Details

Table 386: Virtual Chassis Topology Fields (*continued*)

Field	Function	Default in Topology Monitor
Member ID	Same as Member.	Summary Details
FPC Slot	Identifies the Flexible PIC Concentrator (FPC) slot number for the member.	Summary Details
Member Status	Identifies whether the member is present in the Virtual Chassis or Not Present.	Details
Member Serial Number	Identifies the hardware serial number of the switch.	Details
Member Model	Specifies the Juniper model number of the switch.	Details
Member Location	Identifies the wiring closet for the switch.	Details (hidden)
Member Mixed Mode	Indicates whether the switch is configured to run in mixed member mode. Valid fields are true or false.	Details
Neighbor ID	Identifies the neighbors by the member ID.	Summary Details
Neighbor Interface	The Virtual Chassis Port of the neighbor.	Details

RELATED DOCUMENTATION

[Virtual Chassis Topology Monitor | 1443](#)

[Network Director Documentation home page](#)

VC Equipment Summary By Type Monitor

IN THIS SECTION

- [VC Equipment Summary By Type | 1445](#)
- [VC Equipment Member Status By Type Details | 1445](#)

The VC Equipment Summary By Type displays the operational status of the virtual chassis members associated with Logical View, Location View, Device View and Custom Group View.

VC Equipment Summary By Type

The summary view of the VC Equipment Summary By Type monitor shows the total number of virtual chassis members and the connection status of the virtual chassis members in the selected scope.

NOTE: The VC Equipment Summary tab is enabled only for members that are at the virtual chassis container level, virtual chassis level or virtual chassis device level in the applicable view panes.

Mouse over a segment of the pie chart to see the actual number of VC devices of that type. The connection states are as follows:

- UP (green)—Device is connected to Network Director
- DOWN (red)—Device is not connected to Network Director

Click the details icon to open the VC Equipment Member Status By Type Detail View window.

VC Equipment Member Status By Type Details

The VC Equipment Member Status By Type Detail View window provides details about the VC members in the selected scope. See [Table 387](#) for the description of the table columns.

Table 387: VC Equipment Member Status By Type Detail View

Table Column	Description
Hostname	Configured name of the device or IP address if no hostname is configured.

Table 387: VC Equipment Member Status By Type Detail View (continued)

Table Column	Description
IP Address	IP Address of the device.
Serial Number	Serial number of device chassis.
Platform	Model number of the device.
OS Version	Operating system version running on the device.
Device Family	Device family of the device: <ul style="list-style-type: none"> • JUNOS-EX for EX Series switches • JUNOS for Campus Switching ELS
Device Type	Type of the device: <ul style="list-style-type: none"> • VC—Virtual Chassis master • VC Member—Virtual Chassis member switch
Connection State	Connection status of the device in Network Director. <ul style="list-style-type: none"> • UP (green up arrow)—Device is connected to Network Director • DOWN (red down arrow)—Device is not connected to Network Director
Configuration State	Displays the configuration status of the device: <ul style="list-style-type: none"> • In Sync—The configuration on the device is in sync with the Network Director configuration for the device. • Out Of Sync—The configuration on the device does not match the Network Director configuration for the device. This state is usually the result of the device configuration being altered outside of Network Director. You cannot deploy configuration on a device from Network Director when the device is Out Of Sync. To resolve this state, use the Resynchronize Device Configuration task in Deploy mode • Sync failed—An attempt to resynchronize an Out Of Sync device failed. • Synchronizing—The device configuration is in the process of being resynchronized • N/A—The device is down or is an access point.

RELATED DOCUMENTATION

[Selecting Monitors To Display on the Summary Tab | 1378](#)

[Network Director Documentation home page](#)



Using Fault Mode

[About Fault Mode | 1448](#)

[Using Fault Mode | 1454](#)

[Fault Reference | 1459](#)

About Fault Mode

IN THIS CHAPTER

- [Understanding Fault Mode in Network Director | 1448](#)
- [Understanding the Fault Mode Tasks Pane | 1452](#)

Understanding Fault Mode in Network Director

IN THIS SECTION

- [What Are Events and Alarms? | 1448](#)
- [Alarm Severity | 1449](#)
- [Alarm Classification | 1449](#)
- [Alarm State | 1451](#)
- [Alarm Notifications | 1451](#)
- [Threshold Alarms | 1451](#)

The Fault mode shows you information about the health of your network and changing conditions of your equipment. Use Fault mode to find problems with equipment, pinpoint security attacks, or to analyze trends and categories of errors.

This topic describes:

What Are Events and Alarms?

Activity on a network device consists of a series of *events*. A software component on the network device, called an *entity*, is responsible for running the Simple Network Management Protocol (SNMP) to log and monitor these events. When certain types of events are persistent, or when the condition causing the event crosses a threshold, SNMP sends a notification, also called a *trap* to Network Director. Network

Director correlates traps, describing a condition, into an *alarm*. For example, multiple power supply traps coming from a device are correlated into a single power supply alarm for the device.

There are many types of alarms. An alarm can be as routine as when the device changes state or as serious as when a power supply has failed. When an alarm is sent, or *raised*, it stays raised until the triggering condition is resolved or *cleared*. The system can clear the alarm when the state changes again or an administrator can clear it manually, which indicates that the condition is now resolved.

SNMP also plays another role in Network Director. Enabling devices for SNMP with the appropriate read-only V1/V2/V3 credentials, can speed up device discovery.

Alarm Severity

Alarms are ranked by their impact to the network. The following list shows the ranking of alarms in Network Director from alarms that have the most impact to alarms that have the least impact on the network. It also shows the color scheme associated with each level of severity that is reflected in related graphs.

Critical (Red)—A critical condition exists; immediate action is necessary.

Major (Orange)—A major error has occurred; escalate or notify as necessary.

Minor (Yellow)—A minor error has occurred; notify or monitor the condition.

Info (Blue)—An informational message; no action is necessary. Informational alarms do not necessarily indicate an error. It could indicate that a device or entity has changed state.

Administrators can override the default severity of an alarm and set the severity to match their inhouse guidelines. Changing the severity level for an alarm is done on the Fault tab of System Preferences.

Alarm Classification

Network Director organizes alarms into categories so you can view trends in the types of errors occurring on a network. These categories, shown in [Table 388](#) are derived from the SNMP Management Information Base (MIB) that is the information database or module containing the trap information for the event.

Table 388: Network Director Alarm Classifications

Category	Description
AP/Radio	Indicates alarms for access points and their radios. These alarms are generated from access points.
BFD	Indicates alarms for Bidirectional Forwarding Detection sessions. These alarms are generated from EX Series switches.
BGP	Indicates alarms for BGP4.

Table 388: Network Director Alarm Classifications (*continued*)

Category	Description
Chassis	Indicates alarms for switch hardware, in this case, EX Series switches.
ClientAndUserSession	Indicates alarms for wireless clients.
Cluster/Modo	Indicates alarms about wireless network clusters and mobility domains.
Configuration	Indicates alarms for configuration management.
Controllers	Indicate device alarms.
CoS	Indicates class of service alarms.
DHCP	Indicates local server DHCP alarms.
DOM	Indicates Digital Optical Monitoring alarms that are generated from optical interfaces.
FlowCollection	Indicates alarms generated when collecting and exporting traffic flows.
General	Indicates alarms that are common to all network devices, such as link up/down or authentication.
GenericEvent	Indicates an alarm that is generated from an Op script or event policies.
L2ALD	Indicates MAC address alarms generated from the Layer 2 Address Learning Daemon (L2ALD).
L2CP	Indicates alarms generated by Layer 2 Control Protocol features.
MACFDB	Indicates an alarm for when MAC addresses are learned or removed from the forwarding database of the monitored device.
Misc	Indicates alarms that do not fit into the other categories.
PassiveMonitoring	Indicates alarms that occur on a passive monitoring interface.
Ping	Indicates alarms that a generated during a Ping request.
RFDetect	Indicates alarms from radio frequency conditions. These alarms are generated from wireless controllers.
RMon	Indicates RMON alarms

Table 388: Network Director Alarm Classifications (*continued*)

Category	Description
SONET	Indicates a SONET or SDH alarm on an interface.
SONET APS	Indicates alarms generated on a SONET interface that participates in Automatic Protection Switching (APS).
VirtualChassis	Indicates alarms generated from Virtual Chassis members regarding member or port status.
VNetwork	Indicates virtual networking alarms.

Alarm State

Once an alarm is active, it has one of these states:

- Active—Alarms that are current and not yet acknowledged or cleared.
- Cleared—Alarms that are resolved and the device or entity has returned to normal operation.

Some alarm states go directly from active to cleared state and require little to no administrative effort. However, other alarms with a high severity should be acknowledged and investigated.

In addition to acknowledging and clearing an alarm, you can assign an alarm to someone and you can append a note or annotation to an alarm. Annotations are helpful for documenting the resolution of an alarm or time estimates for a fix. Changes to an alarm's state are made through the Alarm State monitor in Fault mode.

Alarm Notifications

Alarms can be enabled for email notification. When an alarm with notification enabled is generated, an email is sent to a set of specified addresses. There is a list of global email addresses that receive notifications from all alarms with notification enabled. Each alarm type can also have a list of addresses that receive notification when that alarm type is generated. Administrators can enable notification for alarm types and specify addresses to receive email notifications. These tasks are done on the Fault tab of System Preferences.

Threshold Alarms

Threshold alarms are alarms that are generated when a monitored value crosses the configured threshold. They provide enhanced visibility into potential issues on the network. Administrators configure and manage threshold alarms the same way as other alarms, and can set the threshold level of individual threshold alarms on the Fault tab of System Preferences.

RELATED DOCUMENTATION

[Setting Up User and System Preferences | 108](#)

[Alarms by Severity Monitor | 1466](#)

[Alarms by Category Monitor | 1466](#)

[Current Active Alarms Monitor | 1464](#)

[Alarms by State Monitor | 1467](#)

[Network Director Documentation home page](#)

Understanding the Fault Mode Tasks Pane

The Tasks pane in Fault mode provides you with a set of tools for effectively managing alarms on your system.

From the Tasks pane, you can:

- Filter known alarms to locate a specific alarm or error condition by clicking Search Alarms. Use this task to isolate alarms that occurred during a known time-frame or that have annotations associated with them. Although each of the Fault mode monitors can sort the alarms, Search Alarms enable you to submit multiple search and sort arguments as part of your search query.
- View a list of Aruba devices that are connected to a Juniper Networks switch port that is down. You can also launch the Aruba Airwave application to view the alerts for these devices. The [Table 389](#) lists the device management task available in Fault mode.

Table 389: Device Management Tasks in Fault Mode

Task	Description
View Aruba Wireless Device Inventory	Displays the list of Aruba devices connected to Juniper Networks switch port that is down.
Launch Aruba Airwave	Launches the System > Alerts page of the Aruba Airwave application.

In addition, Network Director enables you to group the tasks that you perform frequently and create a list of key tasks. You can add any task from the Tasks pane to the Key Tasks list by selecting a task and clicking the plus (+) sign that appears adjacent to the task. For some modes, you can see that Network Director has predefined some key tasks for you. You can modify this set of tasks to suit your requirements. This feature is available in Task pane irrespective of your current mode, scope, or view.

RELATED DOCUMENTATION

[Searching Alarms | 1454](#)

[Changing Alarm State | 1457](#)

[Understanding Fault Mode in Network Director | 1448](#)

[Understanding Aruba Airwave Integration with Network Director | 1548](#)

[Network Director Documentation home page](#)

Using Fault Mode

IN THIS CHAPTER

- [Customizing Alarms | 1454](#)
- [Searching Alarms | 1454](#)
- [Changing Alarm State | 1457](#)

Customizing Alarms

Ensure that all devices are enabled for SNMP trap forwarding. This task, [Set SNMP Trap Configuration](#), is found in Deploy mode.

Network Director enables you to tailor alarms by:

- Enabling or disabling individual alarms.
- Setting the amount of time alarms are retained in the system.

You can customize alarms using Preferences in the Network Director banner.

RELATED DOCUMENTATION

[Setting Up User and System Preferences | 108](#)

[Network Director Documentation home page](#)

Searching Alarms

Use [Search Alarms](#), available from the Tasks pane, to filter and isolate information about a specific alarm. Use this page to specify complex sorting and filtering criteria for all alarms.

Each field in the Search Alarm window helps narrow the current list of alarms. The more search items you specify, the more specific your results. All fields are optional.

1. Select or type the known descriptors for the alarm. These fields are described in [Table 390](#).
2. Click **Search** to run the query. The Alarms Details page opens with the results of your search.
3. Review the alarm. From this page you can change the state of the alarm, annotate, or assign the alarm to personnel. For more information about changing the state of an alarm, view [“Changing Alarm State” on page 1457](#).

Table 390: Alarm Search Fields

Search Criteria	Description
State	<p>Use the list to select which alarm states to search for:</p> <ul style="list-style-type: none">● All—Alarms of all states.● Active—Alarms that are current and not yet acknowledged or cleared.● Clear—Alarms that are resolved and the device or entity has returned to normal operation.

Table 390: Alarm Search Fields (continued)

Search Criteria	Description
Category	<p>Fill in one of the available alarm categories:</p> <ul style="list-style-type: none"> • AP/Radio • BFD • BGP • Chassis • ClientAndUserSession • Cluster/Modo • Configuration • Controllers • CoS • DHCP • DOM • FlowCollection • GENERAL • GenericEvent • L2ALD • L2CP • MACFDB • Misc • PassiveMonitoring • Ping • RFDetect • RMon • SONET • SONETAPS • VirtualChassis • VNetwork
Severity	<p>Pull down the list to select the severity level. Not all possible alarm severities are listed. Only the severity levels of your current active alarms are shown. Possible selections are:</p> <ul style="list-style-type: none"> • Critical • Major • Minor • Info

Table 390: Alarm Search Fields (*continued*)

Search Criteria	Description
Advanced Search Criteria	
(from) Date	Pull down the calendar and select the starting date of the search.
(from) Time	Pull down the list to select the starting time of the search. Search times are in military (24-hour) clock format in 30 minute intervals.
(to) Date	Pull down the calendar and select the ending date of the search.
(to) Time	Pull down the list to select the ending time of the search. Search times are in military (24-hour) clock format in 30 minute intervals.
Notes	Enter any keywords or phases that were listed in an existing annotation.

RELATED DOCUMENTATION

[Understanding Fault Mode in Network Director | 1448](#)

[Network Director Documentation home page](#)

Changing Alarm State

When an alarm becomes active, it remains active until either the system determines that the condition is resolved or system personnel change the status. Critical alarms always need immediate attention and seldom resolve on their own, but informational messages are often expected actions and results. When a condition is severe or persistent and needs attention, follow these steps:

1. Locate the alarm.
 - a. Click **Fault** in the Network Director banner to enter Fault mode.
 - b. Click the Alarm Details icon on any of the monitors to open the Alarm Details page. Scroll or sort the alarms to find the alarm in question. As an alternate method, click **Search Alarms** in the Tasks pane and filter the active alarm list.

- c. Select the alarm.
2. Review the Event Details that triggered the trap for the alarm. These events provide insight into the cause or location of the problem.
3. Click **Acknowledge** to indicate that the problem is now known. You should receive a message saying the alarm is acknowledged.
4. Depending whether you can resolve the alarm with the information at hand or not, either assign the alarm to a member of your staff or clear the alarm. Click **Clear** to clear the alarm or click **Assign** and fill in the assignee's name.

At any time in the life cycle of an alarm, you can attach information about the alarm to the alarm record by clicking **Annotate**. Fill in your name in the **Notes By** field and add the note description in the **Notes** field. Click **Add** to record the annotation.

RELATED DOCUMENTATION

[Alarm Detail Monitor | 1459](#)

[Network Director Documentation home page](#)

Fault Reference

IN THIS CHAPTER

- Alarm Detail Monitor | 1459
- Current Active Alarms Monitor | 1464
- Alarms by Category Monitor | 1466
- Alarms by Severity Monitor | 1466
- Alarms by State Monitor | 1467
- Alarm Trend Monitor | 1468

Alarm Detail Monitor

IN THIS SECTION

- Finding Specific Alarms | 1460
- Sorting Alarms | 1461
- Reading Events | 1463
- Investigating Event Attributes | 1464
- Changing the Alarm State | 1464

Use the Alarm Detail monitor to sort alarms, view an alarm in depth, and to assign a disposition to an alarm.

By clicking the **Details** icon, you can access the Alarm Detail monitor from any of the four alarm monitors available on the main page in Fault mode (Severity, Category, Current, or State). It is also available from the Current Active Monitors available from the Summary tab in Monitor mode.

This topic describes:

Finding Specific Alarms

Use the Alarm Detail monitor to locate a specific alarm, research the events causing the alarm, and to assign a disposition to the alarm. When an alarm is highlighted in the sorting sequence, the events contributing to the alarm are listed in Event Details and the variable settings are shown in Event Attribute Detail.

To locate an alarm and to assign a disposition to the alarm:

1. Sort the list using the Display list. Sorting choices vary depending on how you arrived here. View [“Sorting Alarms” on page 1461](#) for details on sorting options.
2. Review the sorted list. Each entry shows a minimum of one to a maximum of nine fields. These fields are described in [Table 391](#).
3. Examine the events and event attributes that contributed to sending the alarm. Events and event attributes are discussed in [“Reading Events” on page 1463](#) and [“Investigating Event Attributes” on page 1464](#).

Table 391: Alarm Detail Fields

Field	Value	Shown in Detailed View by Default
Name	The alarm name.	Yes
ID	A system and sequentially-generated identification number.	No
Assigned To	If assigned to an individual, it shows the name of the person assigned; otherwise, it shows System to mark that the alarm is still unassigned.	Yes
Severity	The severity of the alarm. Severity levels are: <ul style="list-style-type: none"> ● Critical—A critical condition exists; immediate action is necessary. ● Major—A major error has occurred; escalate or notify as necessary. ● Minor—A minor error has occurred; notify or monitor the condition. ● Info—An informational message; no action is necessary. 	Yes
Acknowledged	Indicates if the alarm has been acknowledged.	Yes
Entity ID	The identification of the entity responsible for causing this alarm. The Entity ID is the key for correlation of events into an alarm. The Entity ID could be a MAC address of a radio or an IP address of the device.	Yes

Table 391: Alarm Detail Fields (continued)

Field	Value	Shown in Detailed View by Default
Reporting Device IP	The IP address of the reporting device that generated the alarm.	Yes
Reporting Device Name	The hostname of the reporting device.	Yes
Creation Date	The date and time the alarm was first reported.	No
Last Updated	The date and time that the information for the alarm was last modified.	Yes
Updated By	Either the system or the last user who modified the alarm.	No

NOTE: You can enter the alarm name, entity ID, or reporting device IP in the search field and press enter to perform searches and display only those searched alarms in the Alarm Detail table.

You can also use Searching Alarms in the Tasks pane to perform searches using multiple arguments. With multiple arguments, you can isolate a single alarm from a long alarm list. For more information, see [“Searching Alarms” on page 1454](#).

Sorting Alarms

Depending on the monitor you chose to access Alarm Detail, your sorting options change to reflect the summary monitor. The different sort options are listed in [Table 392](#).

Table 392: Sort Options for Alarms

Alarms by Severity Sort	Alarms by Category Sort	Alarms by State and Current Active Alarms Sort
All	AP/Radio	Active
Info	BFD	Clear
Minor	BGP	
Major	Chassis	
Critical	ClientandUserSession	

Table 392: Sort Options for Alarms (continued)

Alarms by Severity Sort	Alarms by Category Sort	Alarms by State and Current Active Alarms Sort
	ClusterAndModDo	
	Config	
	Controllers	
	CoS	
	DHCP	
	DOM	
	FlowCollection	
	GENERAL	
	GenericEvent	
	L2ALD	
	L2CP	
	MACFDB	
	Misc.	
	PassiveMonitoring	
	Ping	
	RFDetect	
	RMon	
	SONET	
	SONETAPS	
	VirtualChassis	

Table 392: Sort Options for Alarms (continued)

Alarms by Severity Sort	Alarms by Category Sort	Alarms by State and Current Active Alarms Sort
	VNetworkS	

Reading Events

When you select an alarm in Alarm Detail, the Event Detail table updates with information about the events that are associated with the alarm. [Table 393](#) lists the fields in Event Detail.

Table 393: Event Detail Fields

Field	Value
Name	The event name; also known as the SNMP trap name.
ID	A system-generated, hexadecimal code that uniquely identifies the event.
Description	If the event is an SNMP event, it is shown as a system-generated event.
Type	The type of event, either fault or system alert.
Category	<p>The category of the event message. The category corresponds to the alarm categories shown in the Alarms by Category monitor and the Alarm Settings window. These categories are:</p> <ul style="list-style-type: none"> • RFDetect • General • Chassis • AP/Radio • BFD • CoreandControllers • Misc.
Source	The identification of the entity that is the cause of this event ; it is not necessarily the ID of the event that generated the event.
Originator	The identification of the entity that generated this event, for example, the switch IP or controller IP address.
Time Updated	The date and time of the last update to the event.

Investigating Event Attributes

The Event Attribute Detail window reflects the variables set during the event. In SNMP terminology, these attributes are known as variable bindings or varbinds. These attributes can provide key information about triggers. For example, if a fan fails, the attribute field could indicate the location of the fan in the chassis.

Changing the Alarm State

When an alarm is first reported, it is considered an active alarm. To change the alarm state, to assign the alarm to a person, or simply to record notes about the alarm, use the buttons on Alarm Details. These buttons are:

- Acknowledge—Use this button to acknowledge or record that the alarm is known and is being addressed.
- Clear—Use this button to clear or remove the alarm. The clear state says that the issue sending the alarm has been resolved and no long requires attention.
- Annotate—Use this button to record actions taken to resolve the alarm.
- Assign—Use this button to assign active or acknowledged alarms to staff.

RELATED DOCUMENTATION

[Alarms by Category Monitor | 1466](#)

[Alarms by Severity Monitor | 1466](#)

[Alarms by State Monitor | 1467](#)

[Current Active Alarms Monitor | 1464](#)

[Network Director Documentation home page](#)

Current Active Alarms Monitor

The Current Active Alarms monitor shows any active alarm that has not yet been cleared. It is one of the four standard monitors available in Alarm mode. Current Active Alarms is a table that has four fields and appear by default. However, nine fields are available for selection. View [Table 394](#) for a description of the table.

Table 394: Current Active Alarms Monitor

Table Column	Description	Shown in Summary by Default	Shown in Detailed View by Default
Name	The alarm name.	Yes	Yes
ID	A system and sequentially-generated identification number.	No	No
Assigned To	If assigned to an individual, it shows the name of the person assigned; otherwise, it shows System to mark that the alarm is still unassigned.	No	Yes
Severity	The severity of the alarm. Severity levels are: <ul style="list-style-type: none"> • Critical—A critical condition exists; immediate action is necessary. • Major—A major error has occurred; escalate or notify as necessary. • Minor—A minor error has occurred; notify or monitor the condition. • Info—An informational message; no action is necessary. 	Yes	Yes
Entity ID	The identification of the entity responsible for causing this alarm. The Entity ID is the key for correlation of events into an alarm. The Entity ID could be a MAC address of a radio or an IP address of the device.	Yes	Yes
Reporting Device IP	The IP address of the device or network element that generated the alarm. The SNMP agent is located at the source IP. In most cases, the source IP is the IP address of the switch or controller.	Yes	Yes
Reporting Device	The hostname or IP address of the reporting device.	Yes	Yes
Creation Date	The date and time the alarm was first reported.	No	No
Last Updated	The date and time that the information for the alarm was last modified.	Yes	Yes
Updated By	Either the system or the last user who modified the alarm.	No	No

Clicking the Details icon opens Alarm Details where you can sort and disposition alarms by state (Acknowledged, Clear, Active).

RELATED DOCUMENTATION

[Alarm Detail Monitor | 1459](#)

[Understanding Fault Mode in Network Director | 1448](#)

[Network Director Documentation home page](#)

Alarms by Category Monitor

Alarms by Category is a table of all active alarms sorted by category. Use this monitor to view where errors are trending. These categories are the same categories shown in the Alarm Settings page.

This monitor is available in all views in the main window when in Fault mode.

The table shows the active categories and the number of alarms per category. Clicking the Details icon on Alarms by Category opens Alarm Details where you can sort these categories and change the state of the alarms.

To create a similar report for a specific period of time, use the Alarm Summary report in Report mode.

RELATED DOCUMENTATION

[Alarm Detail Monitor | 1459](#)

[Understanding the Fault Mode Tasks Pane | 1452](#)

[Setting Up User and System Preferences | 108](#)

[Alarm Summary Report | 1503](#)

[Network Director Documentation home page](#)

Alarms by Severity Monitor

Alarms by Severity is a pie-chart that shows the breakdown of all alarms since the last system restart. It is available on the main page when in Fault mode.

If you mouse over each segment, the total number of alerts for those alarms is shown. Alarm severity levels are:

- Critical (Red)—A critical condition exists; immediate action is necessary.
- Major (Orange)—A major error has occurred; escalate or notify as necessary.
- Minor (Yellow)—A minor error has occurred; notify or monitor the condition.
- Info (Wedgewood Blue)—An informational message; no action is necessary. Informational alarms do not necessarily indicate an error. It could indicate that a device or entity has changed state.

Clicking the Details icon on Alarms by Severity opens Alarm Details where you can sort and disposition individual.

RELATED DOCUMENTATION

[Alarm Detail Monitor | 1459](#)

[Understanding the Fault Mode Tasks Pane | 1452](#)

[Setting Up User and System Preferences | 108](#)

[Network Director Documentation home page](#)

Alarms by State Monitor

The Alarms by State monitor is a pie-chart representation of the states of an alarm: active and cleared. Use this graph to get an overall perspective of the amount of alarms that are active compare to those that are cleared. The Alarms by State monitor is on the main pane when in Fault mode.

Mouse over each segment of the pie-chart shows the number of alarms in these states:

- Active—Alarms that are current and not yet cleared.
- Cleared—Alarms that are resolved and the device or entity has returned to normal operation.

You can create an Alarms by State report for a specified node or a period of time using the Alarms Summary Report in Report mode.

Changing the state of an alarm using Network Director is performed on the Alarm Detail page. Clicking the Details icon on Alarms by State opens Alarm Details where you can sort and set the disposition of the alarms.

RELATED DOCUMENTATION

[Alarm Detail Monitor | 1459](#)

[Understanding the Fault Mode Tasks Pane | 1452](#)

[Setting Up User and System Preferences | 108](#)

[Alarm Summary Report | 1503](#)

[Network Director Documentation home page](#)

Alarm Trend Monitor

The Alarm Trend monitor provides trend information about alarms. The trend information is shown on a line chart, where each alarm severity is shown as a colored line. The legend for the line colors is displayed below the chart. The alarm count is shown on the vertical axis. The time of the data samples is shown on the horizontal axis. This monitor includes tabs that show alarm trend information for active alarms and for new alarms. You can select the time period to display from the list in the title bar.

RELATED DOCUMENTATION

[Understanding Fault Mode in Network Director | 1448](#)

[Network Director Documentation home page](#)

7

PART

Working in Report Mode

[About Report Mode | 1470](#)

[Creating and Managing Reports | 1475](#)

[Report Reference | 1498](#)

About Report Mode

IN THIS CHAPTER

- Understanding Report Mode in Network Director | 1470
- Understanding the Report Mode Tasks Pane | 1472
- Understanding the Scope of Reports You Can Create | 1474

Understanding Report Mode in Network Director

In Report mode in Junos Space Network Director, you can create standardized reports from the monitoring and fault data collected by Network Director. An essential part of the network management lifecycle, reporting provides administrators and management insight into the network for maintenance, troubleshooting, trend and capacity analysis, and provides records that can be archived for compliance requirements.

Network Director provides reports in PDF and HTML formats that use graphs and tables to clearly convey data. Reports are also available in CSV format for importing into spreadsheets. [Figure 58](#) shows some examples of PDF reports.

Figure 58: Examples of Network Director Reports

Network Director Report: Top Users by Data Usage Report
 Generated On: 17 March 2013 06:26 PM UTC
 Generated By: super
 Report Name: Top Users for Alphanet
 Scope: Perspective: Logical Node
 No. of Users: 10
 From: 17 March 2013 06:26 PM UTC
 To: 17 March 2013 06:26 PM UTC

Network Director Report: Alarm Summary Report
 Generated On: 04 February 2013 02:47 PM UTC
 Generated By: super
 Report Name: My Network Alarms and Traffic
 Scope: Perspective: Logical Node
 From: 03 February 2013 02:47 PM UTC
 To: 04 February 2013 02:47 PM UTC

Network Director Report: Device Inventory Report
 Generated On: 13 March 2013 12:30 PM UTC
 Generated By: super
 Report Name: Monthly Device Inventory
 Scope: Perspective: Logical Node
 Report Filters: Device Type: ALL
 Connection State: ALL

Network Director Report: My Network Alarms and Traffic
 Generated On: 17 March 2013 06:26 PM UTC
 Generated By: super
 Report Name: Top Users for Alphanet
 Scope: Perspective: Logical Node
 No. of Users: 10
 From: 17 March 2013 06:26 PM UTC
 To: 17 March 2013 06:26 PM UTC

Device Type Summary

Device Type	Count
EX Switches	6
EX VC Members [VC = 2]	6
Wireless Controllers	5
Wireless Access Points	20
Total	46

EX Switches

HostName	IP Address	Model	Software Version	Serial No	Device Type	Connection State	Configuration State
ip-agnd	10.83.200.194	EX3300-24T	12.2R2.4	B4C208375304	SWITCH	UP	IN_SYNC
ip-tunam01	10.83.200.199	EX4500-40F	12.3R1.7	D65210211234	SWITCH	UP	IN_SYNC
10.83.200.129	10.83.200.129	EX3308	12.2R2.4	CA1710120295	SWITCH	UP	IN_SYNC
stjwa1037-ACC-3	10.83.202.80	EX3200-48P	12.3R1.7	BL0300231181	SWITCH	DOWN	SYNCHRONIZING
stjwa1024-JACC-1	10.83.202.75	EX4200-24T	11.4R7.4	8M0208469895	SWITCH	UP	IN_SYNC
stjrandr17	10.83.212.131	EX3308	12.2R2.4	CA0808430076	SWITCH	UP	IN_SYNC
stjdragn-14	10.83.213.34	EX3300-24T	12.2R2.4	GD0211096483	VC	UP	IN_SYNC
Member-0	N/A	EX3300-24T	N/A	GD0211096503	VC Member	N/A	N/A

Number of Sessions

Number of Sessions	Total Data Used
1	2.89 Mb
1	758 Kb
1	258 Kb
1	109 Kb
1	36 Kb
1	33 Kb
1	24 Kb
1	16 Kb
1	0 Bytes

Active Alarms by Severity

Severity	Count
Informational	17
Warning	14
Major	0
Critical	0
Emergency	0

In addition to choosing the formats for your reports, you can:

- Run reports on-demand or schedule them to run at a specific time or on a recurring schedule.
- Select the portion of network you want the report to cover by selecting a scope in the View pane when you create a report definition. For example, you can run a Device Inventory report on your entire network, on all devices in a wiring closet, or on all EX4200 switches.
- Select the report options—for example, the historical time frame you want an Audit Trail report to cover or the type of devices you want to include in a Device Inventory report.
- Have reports sent to an e-mail address or automatically archived on a file server.

The process for generating reports is simple. Select a scope in the View pane and then create a report definition by using the Create Report Definition wizard. When you complete the report definition, the reports are immediately scheduled to run according to the scheduling choices you have made.

RELATED DOCUMENTATION

[Understanding the Scope of Reports You Can Create | 1474](#)

[Understanding the Report Mode Tasks Pane | 1472](#)

[Network Director Documentation home page](#)

Understanding the Report Mode Tasks Pane

Network Director has built-in reporting features to create standardized reports from your network data. You can schedule these reports to run either in real time or in batch to gain insight into the network for ensuring compliance, performing maintenance, or troubleshooting.

The Report mode analyzes data from different perspectives and filters the data based on the node selected in the network tree.

For example, if you want to view inventory reports on only your wireless controllers, you can select the Device view and the Wireless LAN Controllers node in the network tree to provide granular information on just those devices. After selecting the view and node in the network tree, create a report definition. In this definition file you select from a number of preconfigured reports and set the time frame, schedule, and output options.

From the Reports Tasks pane, you can:

- Set up a new report or change how an existing report is run by clicking Report Definition. From this page, you can launch a wizard that guide you through the process of defining a report or changing a report definition file. The report definition file is based on the report content on the view and the node you select in the network tree. The Filter option in the View pane does not affect the report content.
- View the summary details of the last run of a report, export a report, or to delete a report output by clicking Manage Generated Reports. This page is also the default Reports page. After a report definition is created and a report is generated from that definition, it is shown in the Generated Reports page.

Reports are stored on the application server on which Network Director is running. However, because reports can be large, the report is delivered in a compressed or *zipped* format. and can be stored offline or on a Secure Copy Protocol (SCP) server.

- Set or change the path to an SCP server for report storage. You can also test the connectivity to the server by clicking Test Connection on the Manage SCP Servers page.
- Set or change the path to an SMTP server for e-mail notifications of alerts or for mailing reports to administrators. You can test the connectivity to the server by clicking Test Connection on the Manage SMTP Servers page.

- Create or change a schedule that is used by one or more reports by clicking Manage Schedules. Unless you want to run the report immediately, you need to create a schedule and associate it with the report definition file. Create the schedule before you create the report definition file.

For example, you might want to run several reports that run on the weekend and are available first thing on Monday morning. You could create a single schedule that runs at midnight on Saturday and is delivered to you through e-mail.

- Launch the Aruba Airwave application Reports page that displays the reports generated for Aruba wireless devices connected to Juniper Networks switches.
- Add frequently performed tasks to Key tasks list. You can add any task from the Tasks pane to the Key Tasks list by selecting a task and clicking the plus (+) sign that appears adjacent to the task. For some modes, you can see that Network Director has predefined some key tasks for you. You can modify this set of tasks to suit your requirements. This feature is available in Task pane irrespective of your current mode, scope, or view.

RELATED DOCUMENTATION

[Managing Reports in Network Director | 1475](#)

[Managing Generated Reports | 1489](#)

[Managing Reports on SCP Servers | 1492](#)

[Mailing Reports | 1494](#)

[Scheduling Reports | 1484](#)

[Understanding the Scope of Reports You Can Create | 1474](#)

[Launching the Aruba Airwave Application | 1554](#)

[Network Director Documentation home page](#)

Understanding the Scope of Reports You Can Create

The Report mode enables you to create standard reports from your network information. Reports are based on a report definition that can either be global or granular. You control this global or granular scope of the report definition by your selections in the View pane (the selected view and network tree node).

For example, if you want to run your reports against all switches, you could select the Logical view and the Switching Network node in the network tree. Or if you wanted to run reports on all the devices on a floor of a building, you would select the Location view and navigate to the floor node of a building in the network tree. To pinpoint the performance on a particular switch, you would select the Device view and the individual switch node in the network tree.

TIP: When naming your report definition, include the scope in the name. You cannot tell the scope from the report definition after you have created the definition; you can, however, determine the scope from the generated report.

The reports generated from the report definition file are either formatted and sent to you through e-mail or sent using Secure Copy Protocol (SCP) to a designated repository.

RELATED DOCUMENTATION

[Creating Reports | 1478](#)

[Understanding Report Mode in Network Director | 1470](#)

[Understanding the Report Mode Tasks Pane | 1472](#)

[Network Director Documentation home page](#)

Creating and Managing Reports

IN THIS CHAPTER

- [Managing Reports in Network Director | 1475](#)
- [Creating Reports | 1478](#)
- [Scheduling Reports | 1484](#)
- [Managing Generated Reports | 1489](#)
- [Retaining Reports | 1491](#)
- [Managing Reports on SCP Servers | 1492](#)
- [Mailing Reports | 1494](#)

Managing Reports in Network Director

IN THIS SECTION

- [How to Locate and Manage Reports | 1476](#)
- [Managing Report Definitions | 1476](#)

Reports are generated from a report definition. These definitions establish the type of report, when it is run, and how the report output is presented and preserved. You create, modify, or delete these report definitions from the Manage Report Definition page.

This topic describes:

How to Locate and Manage Reports

The Manage Report Definition page is available from the Report Tasks pane while the Report mode is selected. To locate this page:

1. Select **Report** in the Network Director banner. The Report Tasks pane opens displaying the tasks available in the Report mode.
2. Select **Manage Report Definitions** in the Tasks pane. The Manage Report Definition page opens displaying all existing report definition files.
3. Use the Manage Report Definition page to review existing report definitions, create new definitions, or change a definition.
 - Create a new report definition by clicking **Add**. See [“Creating Reports” on page 1478](#) for help using the Report wizard.
 - Modify an existing report definition by selecting a report definition in the table and clicking **Edit**.
 - Delete an existing report definition by selecting a report type in the table and clicking **Delete**.
 - View details of the report composition, the scope, and perspective of the report definition by clicking **Details**.

Managing Report Definitions

Use the Manage Report Definition page to review existing report definitions, or follow the Report wizard to create new report definitions, delete definitions, or see report details.

Existing report definitions are listed on the page in the format discussed in [Table 395](#). The reports created from these definitions are found under the Manage Generated Reports task.

Table 395: Manage Report Definition Fields

Field	Description
Report Definition	The name of the report definition. Specify a name that indicates the purpose of the report.

Table 395: Manage Report Definition Fields (continued)

Field	Description
Format	<p>The format or file extension of the report output; the final rendering of the output. Valid values are:</p> <ul style="list-style-type: none"> • PDF—(Portable Definition Format) is used for output that is either viewed in a reader or printed. • CSV—(Comma Separated Format) is used for output that is exported into a spreadsheet. • HTML—(Hypertext Markup Language) is used for output that is viewed in a Web browser.
Reporting Mode	<p>(Optional) Where the generated report is sent. Valid values are:</p> <ul style="list-style-type: none"> • Email—Sends a zipped file of the report to an e-mail address. • SCP—Sends a zipped file to a secure server.
Schedule	(Optional) When the report is scheduled to run.
Last Updated By	The userid of the last person to modify the report definition.
Last Updated Time	Time when the report definition was last updated.
Execute Report	Click Run Now to run the report.

RELATED DOCUMENTATION

[Creating Reports | 1478](#)

[Managing Generated Reports | 1489](#)

[Understanding the Scope of Reports You Can Create | 1474](#)

[Understanding the Report Mode Tasks Pane | 1472](#)

[Retaining Reports | 1491](#)

[Network Director Documentation home page](#)

Creating Reports

IN THIS SECTION

- [How to Create a Report Definition | 1478](#)
- [Creating a Report Definition | 1480](#)
- [Setting Report Options | 1482](#)
- [Reviewing the Report Definition | 1483](#)
- [Changing a Report Definition | 1484](#)

Network Director has built-in reporting features to create standardized reports from your network data. You can schedule these reports either to run in real time or in batch to provide insight into the network for compliance, maintenance, or troubleshooting. To define a new report, you select from a number of preconfigured report types and set the scheduling and output options.

This topic describes:

How to Create a Report Definition

You create new reports from the Report Definition page while in the Report mode. To locate this page:

1. Select **Report** in the Network Director banner. The Report Tasks pane opens, displaying the tasks available in the Report mode.
2. (Optional) Select the node on which to run the report in the View pane. Some reports are designed to run at a specific scope in the network tree. For example, if you select an EX Series switch node and attempt to run a Network Neighborhood report that reports on RF strength, the report runs, but is empty.
3. Select **Manage Report Definitions** in the Tasks pane.

[Table 396](#) describes the information provided about report definitions on the Manage Report Definition page.

Table 396: Manage Report Definition Fields

Field	Description
Report Definition	The name of the report definition. Specify a name that indicates the purpose of the report.
Format	<p>The format or file extension of the report output; the final rendering of the output. Valid values are:</p> <ul style="list-style-type: none"> ● PDF—(Portable Definition Format) is used for output that is either viewed in a reader or printed. ● CSV—(Comma Separated Format) is used for output that is exported into a spreadsheet. ● HTML—(Hypertext Markup Language) is used for output that is viewed in a Web browser.
Reporting Mode	<p>(Optional) Where the generated report is sent. Valid values are:</p> <ul style="list-style-type: none"> ● Email—Sends a zipped file of the report to an e-mail address. ● SCP—Sends a zipped file to a secure server.
Schedule	(Optional) When the report is scheduled to run.
Last Updated By	The user ID of the last person to modify the report definition.
Last Updated Time	Time when the report definition was last updated.
Execute Report	Click Run Now to run the report.

Creating a Report Definition

A report definition defines the properties that are used to generate one or more reports. It includes these properties:

- Name of the report definition
- Report type(s)
- Reporting filters
- Scheduling options
- Output format

To create a report definition:

1. Click **Add** on the Manage Report Definition main page to open the Create Report Definition wizard Basic Settings page.
2. Type a name for the report in the Report Definition Name field. After the report runs, you can find a report by this name in the Generated Reports list. Names can contain letters, numbers, spaces, dashes (-), and underscores (_).
3. Select the report types for the report definition in the Select Report Type area:
 - To add one or more report types:
 - a. Click **Add**. The Assign Report Types window opens.
 - b. Select one or more report types from the list in the Assign Report Types window.
 - c. Click **OK**.

The report types you added appear in the Select Report Type list. [Table 397](#) describes the information about report types that is available in the Select Report Type table.

TIP: When adding multiple reports types, be sure all of the reports you select are supported for the node type selected in the view pane.

- To delete one or more report types, select their check boxes in the Select Report Type list, then click **Delete**.
4. (Optional) Edit the report type options for the added report types by clicking **Edit Report Options** in the Customize Report Options column. Configure report type options in the Filter Options window, then click **OK**. [Table 398](#) describes the available report type options.

NOTE: For wireless clients, you can specify the search options such as the user name, MAC address, and IP address of the clients to generate the reports. While you are performing the search using the IP address, you can specify the sub-string of the IP address to generate reports of all the sessions having the sub-string as the prefix of the IP address. Similarly, you can fetch complete session details without filters, by specifying '.' r '.'. These parameters match all the IPs or MACs and provides complete session details.

5. Click **Next** or **Report Options** to set up the report options. You can also click **Cancel** to exit the wizard. For details on report options, see ["Setting Report Options" on page 1482](#).

Table 397: Select Report Type Table Columns

Column Heading	Description
Type	The Report name.
Category	The general classification of the report.
Scope	Shows the scopes that are applicable for the report type. (Appears only in the Assign Report Types window that opens when you click the Add button.)
Description	A description of the use or purpose of the report.
Report Option	Lists the applied report type options.
Customize Report Options	Click the link to change the report type options for that report type.

Table 398: Report Type Options for Data Filtration

Filter Option	Description
Classification Reason	For the Rogue Summary report, filters the rogue devices included in the report to only those that are classified as rogue for the selected reason.
Connection State	Limits the report to devices in this state.
Device Types	Limits the report to this type of device.
Number of Users	Customizes the report to the specified number of users.
Percentage Utilization Exceeding	Specifies the utilization percentage threshold for the report. Only results that exceed the threshold will appear in the report.

Table 398: Report Type Options for Data Filtration (continued)

Filter Option	Description
Search Parameter	Specifies search parameters. Only results that match the parameters will appear in the report. The search parameters are compared to these properties of results to filter the results that appear in the report: IP address, MAC address, username. Separate multiple search parameters with commas (,).
Time Interval	Limits the report to the indicated time period. If you select Custom, the From and To fields become available, enabling you to set a specific reporting period.
Top N Count	Sets the number of items to include in reports that show a fixed number of items. For example, the Traffic and Congestion Summary report includes the top N number of devices that have the highest port utilization and latency. If the scope is a single device, the top N number of ports on the device are included in the report.

Setting Report Options

This page establishes the report schedule and the output format of the report.

1. Choose from the following scheduling options:

- Run the report now
- Select or create a schedule for the report
- Select to both run the report now and to run the report by a schedule

Options for report scheduling are shown in [Table 399](#).

Table 399: Schedule Options for Reports

Field	Action
Run Now	Select this option to immediately run the report one time.
Select Schedule	Select this option to either create a schedule so that it is run at regular intervals, or to select an already established schedule. <ul style="list-style-type: none"> • The Add Schedule link enables you to create a new schedule. • The Select button opens Choose Schedule window that displays the currently configured schedules. Select the check box to choose a schedule to use for the report. To associate the schedule to your report, click OK.

2. Establish the report output format and destination.

Field	Options
Select Format	<p>A report is available in these formats:</p> <ul style="list-style-type: none"> ● PDF—Choose this format if you want to print the report. Portable Definition Format (PDF) enables the report to be printed from any operating system with the same formatting results. ● CSV—Choose this format if you want to export the report data to a spreadsheet or other business application. The Comma-Separated Values (CSV) format takes the raw data from the report and delineates the fields with commas so that it imports into popular spreadsheet programs. ● HTML—Choose this format if you want to view the report in a browser. <p>NOTE: Because reports can be quite large, they are initially delivered as a zipped (compressed) file.</p>
Mode	<p>Reports can be sent to your e-mail address, to a secure server, or to both.</p> <ul style="list-style-type: none"> ● Select EMAIL and type the e-mail address to have the report sent through e-mail. Network Director uses SMTP server settings for e-mail routing. You can configure an SMTP server from the Tasks pane. ● Select SCP to send the report to the secure server that is marked as active, using Secure Copy Protocol. The settings for secure servers are available in Tasks > Manage SCP Servers.

3. Click **Next** or **Summary** to review the report definition.

Reviewing the Report Definition

The Report wizard guides you to the Summary Page where you can review your report configuration and make any changes before you run the report.

1. Review your Report Name and Report Type in basic settings. If you want to change either of these settings, click **Edit** to return to the Basic Settings page.
2. Review your Report Options. If you want to change these settings, click **Edit** to return to the Report Options page.
3. Click **Finish** when you are done with the report configuration and to exit the wizard.

Changing a Report Definition

You can change an existing report definition file from the Manage Report Definition page.

To change a report definition:

1. Select the check box for the report definition.
2. Click **Edit** to reopen the report definition in the Report wizard. The system returns you to the Summary page, where you can make changes to the report definition.
3. Click **Details** to review the details of the report definition or click **Delete** to remove the report definition. To remove all of the report definitions, select the check box in the header next to Report Definition to select all of the report definitions and click **Delete**.

RELATED DOCUMENTATION

[Managing Generated Reports | 1489](#)

[Understanding the Scope of Reports You Can Create | 1474](#)

[Managing Reports on SCP Servers | 1492](#)

[Mailing Reports | 1494](#)

[Scheduling Reports | 1484](#)

[Understanding Report Mode in Network Director | 1470](#)

[Network Director Documentation home page](#)

Scheduling Reports

IN THIS SECTION

- [How to Create or Manage Schedules | 1485](#)
- [Managing Schedules | 1485](#)
- [Creating New Schedules | 1486](#)
- [Editing Schedules | 1488](#)
- [Deleting Schedules | 1488](#)

You can run Network Director reports as needed or you can automate reports to run in batch by creating a schedule. You can associate a single schedule with one or more reports when you create the report definition. Although you can create a schedule during the report definition process, it is helpful to have the schedules configured before defining the report. To create a new schedule you name the schedule and set the time and frequency of the run.

This topic describes:

How to Create or Manage Schedules

You create a schedule from the Manage Schedules page. You can display this page from the Report mode Tasks pane.

- Select **Report** in the Network Director banner. The Report mode Tasks pane displays the tasks available in the Report mode. Reports run in any network view (Logical, Location, or Device).
- Select **Manage Schedules** in the Tasks pane. The Manage Schedules page opens, displaying all existing report schedules.

From the Manage Schedules page, you can:

- Create a new schedule
- Edit an existing schedule
- See the details of a schedule
- Delete a schedule

Managing Schedules

Use the Manage Schedules page to administer report schedules. From this page, you can create new schedules and view, modify, and delete existing schedules. On the Manage Schedules page, a table of existing report schedules appears. You can sort and customize this table to exclude fields that might not be relevant to your needs. The fields in the table are defined in [Table 400](#).

Table 400: Manage Report Schedules

Field	Description
Schedule Name	The name of the schedule. Indicate the purpose of the schedule in the name.
Schedule Type	Schedules are either One-Time or Recurring. <ul style="list-style-type: none"> • One-Time—These schedules are helpful for running a non-repeating report in batch-mode, such as running it at 3:30 a.m. • Recurring—These schedules are helpful for routine reports and trend analysis.

Table 400: Manage Report Schedules (*continued*)

Field	Description
Recurrence Pattern	How often the pattern repeats. The pattern applies only to recurring schedules. Valid values are: <ul style="list-style-type: none"> • Hourly • Daily • Weekly • Monthly
Description	Details of the reoccurrence pattern.
Status	Either Active or blank. Active indicates the schedule is running.

From this page you can:

- Display a summary of all the parameter settings of a schedule by selecting the schedule and clicking Details. The Report Schedule Summary opens.
- Create a new schedule by clicking Add. The Create Schedule window opens.
- Select a schedule and change the settings by clicking Edit. The Edit Schedule window opens.
- Select a schedule and click **Delete** to remove a schedule.



CAUTION: Take care when deleting a schedule. Network Director enables you to delete a schedule even if it is active.

Creating New Schedules

Use the Create Schedule window to create a one-time or a recurring schedule.

To create a new schedule:

1. Click **Add** on the Manage Schedules page. The Create Schedule window opens.
2. Enter the name for the schedule in **Schedule Name**. Indicate the purpose of the schedule in the name.
3. Choose a one-time run-option or a recurring schedule from the list.
 - One-time schedule options are described in [Table 401](#).
 - Recurring schedule options are described in [Table 402](#).

Depending on the schedule range selected, these settings change dynamically.

Table 401: One-Time Schedule Options

Field	Action
Execute Start Date	Select the date when the schedule is run. You can either fill in a date directly or click the calendar icon to pick a date from a traditional calendar.
Execute Start Time	Select the time the report is run. Time is shown in 24-hour clock format in increments of 15 minutes.

Table 402: Recurring Schedule Options

Field	Action
Hourly	Run the report associated with the schedule, hourly between these hours at intervals of x minutes. Start and end times are shown in 24-hour clock format, in increments of 15 minutes. For example, if you want to schedule a report to run from 1 am to 3 pm at 30 minute intervals, your settings would be: between 13:00 and 15:00 at 30 min(s).
Daily	Run the report associated with the schedule either every weekday or on the specified number of sequential days. Use the up and down arrows to set the number of sequential days or click Every weekday .
Weekly	Run the report associated with the schedule on one or more days of the week. Set the schedule to repeat the run in the specified number of weeks. Use the up and down arrows to set the weekly frequency of how often the schedule is repeated. Click one or more of the days when the report is run.
Monthly	Run the report associated with the schedule either on a certain day of the month or on a specified day and week for the specified number of months. For example, if your organization has a congestion spike at the end of the fiscal quarter, you might want to run reports on the last Friday every 4 months.

- Specify when to start and end implementation of this schedule as described in [Table 403](#).

Table 403: Range of Recurrence Fields

Field	Action
Start Time:	Select the time of day. The clock is in 24-hour format, in increments of 15 minutes, when the schedule begins to run.

Table 403: Range of Recurrence Fields (*continued*)

Field	Action
Start Date:	Select the date when the schedule is first implemented. Format is <i>yyyy-mm-dd</i> .
No end date	Select to indicate to continue to use this schedule until it is modified or deleted.
End After: x occurrence	Select to run the schedule for a specified number of times. Use the up and down arrow keys to specify the number of times to run the schedule.
End by:	Select a time and date to stop running the schedule. Date format is <i>yyyy-mm-dd</i> . Select the time from the list; clock times are in increments of 15 minutes.

5. Click **Add** to finish and to validate the schedule.

Editing Schedules

To change an existing schedule:

1. Select a schedule from the list in the Manage Schedules page.
2. Click **Edit** to reopen the schedule settings.
3. Change the settings based on the values described in [Table 401](#).
4. Click **Edit** to save the revised settings.

Deleting Schedules



CAUTION: Network Director enables you to delete an active schedule. Be aware that deleting a schedule that is active will cause reports not to run.

You can also permanently remove a schedule by selecting the schedule in the Manage Schedule page and clicking Delete.

RELATED DOCUMENTATION

| [Creating Reports](#) | 1478

Managing Generated Reports

IN THIS SECTION

- [Reviewing Generated Reports | 1489](#)
- [Viewing Report Details | 1490](#)
- [Exporting Reports | 1491](#)
- [Deleting Generated Reports | 1491](#)

After a report definition is created and the initial report is run, Network Director populates the Generated Reports page with summary information about the run of the report.

From the Generated Reports page you can view the report details, export the report to view or store in a new location, or delete the report.

This topic describes:

Reviewing Generated Reports

After a reports runs, information about the report is recorded on the Generated Reports page, as shown in [Table 404](#).

Table 404: Fields in the Generated Reports Page

Field	Description
Report Definition	The name assigned at report creation time.
Executed By	The userid of the report owner who ran the report.
Start Time	When the report began to run.
End Time	When the report ended.
Format	The chosen report format, possible values are PDF, CSV, or HTML.

Table 404: Fields in the Generated Reports Page (*continued*)

Field	Description
Generated Report	Links to view or download the report.

From the Generated Reports page, you can either select the check box in the heading to select all of the reports or select the check box for the individual reports and:

- View details about the running of the report by clicking Report Details.
- Export the report by clicking Export.
- Delete the report by clicking Delete.

Viewing Report Details

Use the Generated Report Details window to see information about the report composition. Viewing the Report Details is helpful when a report comprises many smaller reports. See [Table 405](#) for these field descriptions.

Table 405: Generated Report Details

Field	Description
Report Definition Name	The name assigned at report creation time.
Executed By	The userid of the report owner who ran the report.
Start Time	When the report began to run.
End Time	When the report ended.
Format	The chosen report format, possible values are PDF, CSV, or HTML.
Generated Report	This link opens the Report Details window where you can view or download the report.

Reports are kept on the application server until either you delete them or they are deleted by the system. The amount of time a report is saved on the system depends on the report retention settings for Network Director. Network Administrators can globally set the report retention period for reports in the system Preferences, located in the Network Director banner.

Exporting Reports

Use the Export Report window to store multiple reports to a file location. Because reports can be large, they are delivered as compressed *zipped* files for both viewing or storing. After you choose a report to export, you are prompted to select a program to view the report or to download the file.

You can either unzip the report for viewing or save the report to a file location.

TIP: If you choose to save the file, you might want to give a unique name to the unzipped file. After unzipping the report, the name of the report reverts to the type of report you selected. It does not retain the name of the report.

Deleting Generated Reports

Use Delete to discard unneeded or outdated reports. When you click **Delete**, you are prompted to confirm the file deletion. Because deleted reports cannot be recovered, save the report offline before deleting them from the system.

RELATED DOCUMENTATION

[Creating Reports | 1478](#)

[Managing Reports in Network Director | 1475](#)

[Retaining Reports | 1491](#)

[Network Director Documentation home page](#)

Retaining Reports

Reports are stored on the Junos Space server where Network Director is running; however, you set the report retention time in Network Director. The default for the report retention period is 30 days. Because accumulating old reports could eventually impact system performance, you might want to consider changing this setting. The report retention period is set in Preferences on the Report tab.

Another option is to store the report to another location such as a Secure Copy Protocol (SCP) server and then delete the report from the Network Director application server.

RELATED DOCUMENTATION

[Setting Up User and System Preferences | 108](#)

[Managing Generated Reports | 1489](#)

[Managing Reports on SCP Servers | 1492](#)

[Network Director Documentation home page](#)

Managing Reports on SCP Servers

IN THIS SECTION

- [How to Configure SCP Servers | 1492](#)
- [Managing SCP Servers | 1493](#)

If your organization requires reports be stored on a secure server using Secure Copy Protocol (SCP), you can set one or more servers as a reports repository. A reports repository enables you to keep reports long-term for compliance requirements or for your organizational needs.

This topic describes:

How to Configure SCP Servers

SCP servers are used as report repositories for Network Director reports. You can set up or manage a secure server for Network Director reports from the Manage SCP Servers page in the Report mode. To access this page:

1. Select **Report** in the Network Director banner. The Report Tasks pane opens, displaying the tasks available in the Report mode.
2. Select **Manage SCP Servers** in the Report Tasks pane. The Manage SCP Servers page opens, displaying all existing report schedules.

Use the Manage SCP Servers page to:

- View existing SCP server settings
- Set up new SCP servers for reports
- Edit SCP server settings

- Test the connection to an SCP server
- Make an SCP server active
- Delete an SCP server setting

Managing SCP Servers

Use the Manage SCP Servers page to view and manage SCP server settings.

- The Manage SCP Servers page lists any existing server settings. The fields in the Manage SCP Server page are described in [Table 406](#).

Table 406: Managing SCP Server Fields

Field	Description
Server Name	The name you are using to identify the SCP server.
IP Address	The IP address or hostname of the SCP server.
Port Number	The forwarding port number. Default port number for SCP is 22.
Active	Either yes or no to indicate whether it is the active server. Only one server can be active at a time.
Base Path	The path on the server where the reports are to be stored.

- Create new or edit existing server settings:
 1. Establish a new server definition by clicking **Add** or edit an existing definition by clicking **Edit**. Either an Add SCP Settings or Edit SCP Settings page opens.
 2. Fill in the settings described in [Table 407](#).

Table 407: Defining an SCP Server

Field	Action
Server Name	Type a name for this SCP server.
IP Address/Host Name	Type the IP address of SCP server.
Port Number	Type the forwarding port number. Default port number for SCP is 22.
User Name	Type the account name accessing the server.

Table 407: Defining an SCP Server (continued)

Field	Action
Password	Type the password twice for the account on the secure server.
Default Path	Type the file path to the server where the reports are to be stored.
Set Active	Select if you want this server to be the active server. While many servers can be set up as SCP servers for reports, only one server is marked as active.

3. Click **Done** to complete the process.
4. Click **Test Connection** to ensure your server is set up correctly. Network Director attempts to connect to the SCP server and tells you whether the connection could be established.
5. Select a server and click **Set Active** to make that server available for secure services.
You can also delete any SCP server definition from use by Network Director reports by clicking **Delete**.

RELATED DOCUMENTATION

[Understanding the Report Mode Tasks Pane | 1472](#)

[Managing Generated Reports | 1489](#)

[Network Director Documentation home page](#)

Mailing Reports

IN THIS SECTION

- [How to Configure SMTP Servers | 1495](#)
- [Managing SMTP Servers | 1495](#)
- [Adding or Editing SMTP Server Settings | 1497](#)

You can set up one or more electronic mail servers to send reports to e-mail addresses. These servers use the Simple Mail Transfer Protocol (SMTP) to forward the reports. While you can configure many servers as SMTP servers, you can only designate one as the primary mail server.

This topic describes:

How to Configure SMTP Servers

An SMTP server is responsible for sending e-mails. Network Director uses the SMTP server to send reports to users. Under most circumstances, you need only one SMTP server. However, you might want to configure more than one SMTP server if you need a server with a distinct SMTP server configuration. In this case, you would configure multiple SMTP servers and mark the server you want to use as Active.

You can set up or manage SMTP servers from the Manage SMTP Servers page in the Report mode. To access this page:

1. Select **Report** in the Network Director banner. The Report Tasks pane opens, displaying the tasks available in the Report mode. The Generated Reports page loads in the main window.
2. Select **Manage SMTP Servers** in the Report Tasks pane. The Manage SMTP Servers page opens in the main window, displaying all existing SMTP servers configured for Network Director.

Use the Manage SMTP Servers page to:

- View existing SMTP server settings
- Set up new SMTP servers
- Edit existing SMTP server settings
- Test the connection to a SMTP server
- Set an SMTP server as the active server
- Delete an SMTP configuration
- See details of the SMTP configuration

Managing SMTP Servers

Use the Manage SMTP Servers page to view and manage SMTP server settings. The Manage SMTP Servers page lists any existing SMTP server settings. The fields in the Manage SMTP Server page are described in [Table 408](#).

Table 408: Managing SMTP Server Fields

Field	Description	Hidden or Displayed by Default
Name	The name you are using to identify the SMTP server.	Displayed
Host Address	The IP address or hostname of the SMTP server.	Displayed
Port	The forwarding port number. Default port number for SMTP is 587.	Displayed
Active	Either yes or no to indicate whether it is the active server. Only one server can be active at a time.	Displayed
User Auth	Indicates whether SMTP authentication is required for the server. This field is either yes or no.	Displayed
Use TLS	Indicates whether Transport Layer Security (TLS) protocol is used to provide shared-secret encryption.	Displayed
User Name	Indicates the username when user credentials are required for SMTP authentication.	Hidden
From E-mail Address	The e-mail account that sends the report.	Hidden

1. Establish a new server definition by clicking **Add** or edit an existing definition by selecting the server and clicking **Edit**. Either an Add SMTP Settings or Edit SMTP Settings page opens. See [“Adding or Editing SMTP Server Settings” on page 1497](#) for details on setting up or changing server settings.
2. Click **Done** to complete the process.
3. Click **Test Connection** to ensure your server is set up correctly. Network Director tells you whether the attempted connection to the SMTP server could be established.
4. Select a server and click **Set Active** to make that server responsible for sending e-mail.

You can also delete any SMTP server definition from use by Network Director reports by clicking **Delete**.

Adding or Editing SMTP Server Settings

The process of establishing a new SMTP server or to changing the values on an existing server is straightforward. Simply enter or change the values in the fields in the Add SMTP Server or Edit SMTP Server page. These fields are described in [Table 409](#).

Table 409: Defining an SMTP Server

Field	Action
Server Name	Type a name for this SMTP server.
Host Address	Type the IP address of SMTP server.
Port Number	Type the forwarding port number. Default port number for SMTP is 587.
From Email Address	Type the e-mail address used to send the notification.
Set as Active Server	Checking this box sets the server as the Active server. If there is only one server, you cannot clear this box.
Use SMTP Authentication	Checking this box requires the server to use SMTP authentication. You must provide user credentials to use SMTP Authentication.
User Name	Type the account name accessing the server for SMTP authentication.
Password	Type the password twice that is used for authentication.
Use TLS	Select if you want this server to use TLS protocol on the SMTP server.

RELATED DOCUMENTATION

[Understanding the Report Mode Tasks Pane | 1472](#)

[Network Director Documentation home page](#)

Report Reference

IN THIS CHAPTER

- Active User Sessions Report | 1498
- Alarm History Report | 1500
- Alarm Summary Report | 1503
- Audit Trail Report | 1505
- Wireless Client Details Report | 1507
- Client Devices Report | 1510
- Device Inventory Report | 1511
- Fabric Analyzer Report | 1513
- Network Device Traffic Report | 1514
- Network Neighborhood Report | 1516
- Radio Traffic Report | 1518
- Port Bandwidth Utilization Report | 1519
- RF Interference Detail Report | 1521
- RF Interference Summary Report | 1522
- Rogue Summary Report | 1524
- Wireless Security Alarms Report | 1525
- Top Users by Data Usage Report | 1526
- Traffic and Congestion Summary Report | 1528

Active User Sessions Report

The Active User Sessions report is a standardized report generated in Network Director to show the activity level of current users on a specified node. There are two portions of the report: a report header and the report body. The contents of the report header are found in [Table 410](#).

Table 410: Active User Session Report Header

Field	Description
NETWORK DIRECTOR REPORT:	The type of report; In this case, the Active User Sessions report.
Generated On:	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2013 11:35 PM PST.
Generated By:	The username of the user that generated the report.
Report Name	The name of the report assigned when the report was created.
Scope	The report is generated from this view and node in the network. <ul style="list-style-type: none"> • <i>Perspective</i>—Can be a Logical view, Location view, or a Device view of the network. • <i>Node</i>—Represents the selected object on which the report is based.

The base report is a table with the fields described in [Table 411](#).

Table 411: Active User Session Report Fields

Field	Description
User Name	The name that identifies the user to the network.
Client MAC Address	The MAC address of the user.
Controller IP	The IP address of the wireless controller.
Total Bytes	The total number of bytes for the session. Bytes are shown in system international (SI) notation. For example, terabytes (TB), gigabytes (GB), megabytes (MB), and Kilobytes (KB).
AP Name	The network name for the access point.
SSID	The network identifier for the access point that the client is using.
Client IP	The IP address of the client.
Auth Type	The authentication type.
Bandwidth (KBps)	The transfer speed in Kilobytes per second.

Table 411: Active User Session Report Fields (continued)

Field	Description
VLAN	The VLAN name.
Session Elapsed Time	The amount of time after the user started the session.

RELATED DOCUMENTATION

[Understanding the Scope of Reports You Can Create | 1474](#)

[Creating Reports | 1478](#)

[Managing Generated Reports | 1489](#)

[Managing Reports on SCP Servers | 1492](#)

[Network Director Documentation home page](#)

Alarm History Report

IN THIS SECTION

[Alarm History Header | 1500](#)

[Alarm History Tables | 1501](#)

The Alarm History report is a standardized report generated in Network Director. It shows all active, acknowledged, and cleared alarms that occurred within a specified period of time for the indicated node. The report has two portions: a report header and the report body.

This topic describes:

Alarm History Header

The Alarm History report header provides file creation information about the report. The contents of the report header are described in [Table 412](#).

Table 412: Alarm History Report Header

Field	Description
NETWORK DIRECTOR REPORT	The type of report—in this case, the Alarm History report.
Generated On	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2012 11:35 PM PST.
Generated By	The username of the user that generated the report.
Report Name	The name of the report assigned when the report was created.
Scope	The report is generated from this view and node in the network. Perspective can be Logical view, Location view, or Device view of the network. Node represents the selected object upon which the report is based.
Report Filters	The period of time specified for data collection.

Alarm History Tables

The body of the Alarm History report is a set of tables: one table for each alarm state (active, acknowledged, and cleared). In each table the alarms are listed by severity level. The fields of the tables have a common format, which is described in [Table 413](#).

Table 413: Active Alarm History Fields

Field	Description
Alarm Name	The SNMP alarm name.
Severity	One of six levels that indicate the gravity of the alarm based on the impact on the system: <ul style="list-style-type: none"> ● Critical ● Major ● Minor ● Warning (customer defined) ● Alert (customer defined) ● Notice ● Info

Table 413: Active Alarm History Fields (continued)

Field	Description
Category	One of twenty-four functional areas: <ul style="list-style-type: none"> • AP/Radio • BFD • BGP • Chassis • ClientAndUserSession • Configuration • Controllers • CoS • DHCP • DOM • FlowCollection • General • GenericEvent • L2ALD • L2CP • MACFDB • Misc. • PassiveMonitoring • Ping • RFDetect • RMon • SONET • SonetAPS • VirtualChassis
Description	An indication of what caused the alarm.
Source	The Entity ID of the network device sending the trap.
Acknowledged	Whether or not the alarm has been acknowledged.
Updated On	The date when the alarm was last updated (assigned, annotated, or acknowledged) otherwise, the date the alarm was created.

RELATED DOCUMENTATION

[Understanding the Scope of Reports You Can Create | 1474](#)

[Creating Reports | 1478](#)

[Managing Generated Reports | 1489](#)

[Managing Reports on SCP Servers | 1492](#)

[Network Director Documentation home page](#)

Alarm Summary Report

IN THIS SECTION

● [Alarm Summary Header | 1503](#)

● [Alarm Summary Charts | 1504](#)

The Alarm Summary Report is a standardized report generated in Network Director. It shows a graphical summary of the alarms that occurred within a specified period of time, node, and network view. The report has two portions: a report header and a report body.

This topic describes:

Alarm Summary Header

The report header provides file creation information about the report. The contents of the report header are described in [Table 414](#).

Table 414: Alarm Summary Report Header

Field	Description
NETWORK DIRECTOR REPORT	The type of report—in this case, the Alarm Summary report.
Generated On	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2012 11:35 PM PST.
Generated By	The username of the user that generated the report.

Table 414: Alarm Summary Report Header (continued)

Field	Description
Report Name	The name of the report assigned when the report was created.
Scope	The report is generated from this view and node in the network. Perspective can be Logical view, Location view, or Device view of the network. Node represents the selected object upon which the report is based.
Report Filters	The period of time specified for data collection.

Alarm Summary Charts

The body of the Alarm Summary report contains a series of colored charts that provide insight into the trends or distribution of alarms in the network. The first chart summarizes the proportion of active and clear alarms. The rest of the charts are divided into two identical sets: one set for active alarms and one set for clear alarms. If there are no alarms in an active or in a clear state, the set of charts for that state are omitted.

The charts are:

- Alarms by State—Shows the proportion of alarms in active and clear states.
- Active Alarms by Severity (Clear Alarms by Severity)—Shows the proportion of alarms in each severity classification. The default severity levels are:
 - Critical (Red)**— A critical condition exists; immediate action is necessary.
 - Major (Orange)**— A major error has occurred; escalate or notify as necessary.
 - Minor (Yellow)**— A minor error has occurred; notify or monitor the condition.
 - Info (Light Blue)** —An informational message; no action is necessary.

In Preferences, administrators can override the default severity of an alarm and set the severity to match their inhouse guidelines.

- Active Alarm Summary Chart (Clear Alarm Summary Chart)—Shows alarms by category and severity.
- Active Alarms by Category (Clear Alarms by Category)—Shows the number of alarms in each alarm category and, within category, the number of alarms that are acknowledged or unacknowledged.
- Active Alarms by Type (Clear Alarms by Type)—Shows the number of alarms of each specific type. The types are color-coded by severity.

- Top 10 Sources of Active Alarms (Top 10 Sources of Clear Alarms)—Identifies the top 10 devices that are generating the most alarms. Shows the number of alarms each device is generating by severity.
- Active Alarms by Timestamp (Clear Alarms by Timestamp)—This section of the report contains two graphs that plot alarms by their created timestamp:
 - Alarms by timestamp per severity—Plots each alarm of given severity against the time the alarm was created. For example, if during the time period covered by the report there are 10 alarms of major severity, the graph shows 10 orange data points that are plotted against the time the alarms were created.
 - Active alarms by timestamp for top 10 sources (Clear alarms by timestamp for top 10 sources)—For each source, plots the alarms generated by the source against the time they were created.

RELATED DOCUMENTATION

[Understanding the Scope of Reports You Can Create | 1474](#)

[Creating Reports | 1478](#)

[Managing Generated Reports | 1489](#)

[Managing Reports on SCP Servers | 1492](#)

[Network Director Documentation home page](#)

Audit Trail Report

The Audit Trail report is a standardized report generated in Network Director. It shows a history of users accessing the system, modifications to applications, and network management activities for a specific period of time. The report is defined and generated from the Report mode in Network Director.

There are two portions of the report: a report header and the report body. The contents of the report header are shown in [Table 415](#).

Table 415: Audit Trail Report Header Information

Field	Description
NETWORK DIRECTOR REPORT:	The type of report; In this case, the Audit Trail report.
Generated On:	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2012 11:35 PM PST.

Table 415: Audit Trail Report Header Information (continued)

Field	Description
Generated By:	The username of the user that generated the report.
Report Name	The name of the report assigned when the report was created.
Scope	The report is generated from this view and node in the network. Perspective can be Logical view, Location view, or Device view of the network. Node represents the selected object upon which the report is based.
Report Filters	The period of time specified for data collection.

The body of the report comprises the Audit Logs by Task Type chart, the Audit Logs by User chart, and a Audit Detail table.

The Audit Logs by Task Type is a bar chart that lists all of the user and system activities over the specified time period for all users.

The Audit Logs by Users is a pie chart that graphically represents all the active users in a specified time period.

The fields in the report body table are shown in [Table 416](#).

Table 416: Audit Trail Report Fields

Field	Description
User Name	The userid of the individual associated with the activity.
User IP	The IP address of the client.
Task	A short summary of the activity, such as Backup or Login.
Description	The description of the system activity being logged. Examples of common logging activities include logging in and out of the system, modifying application settings, creating SMTP or SCP servers, or database backups.
Result	The result of the system activity: whether it is successful or not. Regularly scheduled events, such as backups, show as recurring.
Job ID	The system generated identification for applicable tasks.
Timestamp	The date and time of the activity. The date is shown in the format: [Day of the Month] [Month] [Year], while the time is shown in standard 12-hour clock format.

RELATED DOCUMENTATION

[Understanding the Scope of Reports You Can Create | 1474](#)

[Creating Reports | 1478](#)

[Managing Generated Reports | 1489](#)

[Managing Reports on SCP Servers | 1492](#)

[Viewing Audit Logs From Network Director | 104](#)

[Network Director Documentation home page](#)

Wireless Client Details Report

The Wireless Client Details report is a standardized report generated in Network Director to show information about wireless network clients. There are two portions of the report: a report header and the report body. The contents of the report header are described in [Table 417](#).

Table 417: Wireless Client Details Report Header

Field	Description
NETWORK DIRECTOR REPORT:	The type of report; in this case, the Client Details report.
Generated On:	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2013 11:35 PM PST.
Generated By:	The username of the user that generated the report.
Report Name	The name of the report assigned when the report was created.
Scope	The report is generated from this view and node in the network. <ul style="list-style-type: none"> • <i>Perspective</i>—Can be a Logical View, Location View, or a Device View of the network. • <i>Node</i>—Represents the selected object on which the report is based.
Report	The search parameters specified for the report.

The report body contains a table for each of the following types of session information about the clients included in the report:

- **Current Session Information**—Provides information about the client's current session. [Table 418](#) describes the columns in the Current Session Information tables.

- Past Session Information—Provides information about the client's past sessions. [Table 419](#) describes the columns in the Past Session Information tables.
- Roaming History—Provides information about the client's roaming history. [Table 420](#) describes the columns in the Roaming History tables.

Table 418: Wireless Client Details Report Current Session Information Tables

Row	Description
User	Name of the user session.
Session Started On	Date and time when the session started.
Elapsed Time	Length of time the session has been active.
Controller IP	IP address of the wireless controller to which the session is connected.
VLAN	VLAN the session is in.
SSID	SSID to which the session is connected.
RSSI	Received Signal Strength Indicator (RSSI) of the wireless session.
SNR	Signal-to-noise ratio (SNR) of the wireless session.
Session Location	Location of the wireless access point to which the session is connected.
Client Device Type	Client's device type.
Client Device Group	Client's device group.
Client Device Profile	Client's device profile.
Receive Unicast KBytes	Amount of unicast data the session has received, in kilobytes.
Transmit Unicast Kbytes	Amount of unicast data the session has transmitted, in kilobytes.
Receive Multicast Kbytes	Amount of multicast data the session has received, in kilobytes.
Link Local	Client device's link local address.
Data Usage (Kbytes)	Total amount of data the session has transmitted and received.

Table 419: Wireless Client Details Report Past Session Information Tables

Column	Description
User	Name of the user session.
Session Started	Date and time when the session started.
Elapsed Time	Length of time the session was active.
Controller IP	IP address of the wireless controller to which the session was connected.
Client IP	Client's IP address.
SSID	SSID to which the session was connected.
RSSI	Received Signal Strength Indicator (RSSI) of the wireless session.
SNR	Signal-to-noise ratio (SNR) of the wireless session.
Rx Unicast (Kbytes)	Amount of unicast data the session received, in kilobytes.
Tx Unicast (Kbytes)	Amount of unicast data the session transmitted, in kilobytes.
Data Usage (Kbytes)	Total amount of data the session transmitted and received.

Table 420: Wireless Client Details Report Roaming History Tables

Column	Description
Start Time	Date and time when the session connected to a wireless access point.
AP	Wireless access point to which the session connected.

RELATED DOCUMENTATION

[Understanding the Scope of Reports You Can Create | 1474](#)

[Creating Reports | 1478](#)

[Understanding Wireless Interference | 914](#)

[Managing Generated Reports | 1489](#)

[Managing Reports on SCP Servers | 1492](#)

[Network Director Documentation home page](#)

Client Devices Report

The Client Devices report is a standardized report generated in Network Director to show the distribution of user sessions by device types, device groups, and device profiles. There are two portions of the report: a report header and the report body. The contents of the report header are described in [Table 421](#).

Table 421: Client Devices Report Header

Field	Description
NETWORK DIRECTOR REPORT:	The type of report; in this case, the Client Devices report.
Generated On:	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2013 11:35 PM PST.
Generated By:	The username of the user that generated the report.
Report Name	The name of the report assigned when the report was created.
Scope	The report is generated from this view and node in the network. <ul style="list-style-type: none"> • <i>Perspective</i>—Can be a Logical View, Location View, or a Device View of the network. • <i>Node</i>—Represents the selected object on which the report is based.

The report body contains sections for device types, device groups, and device profiles. Each section shows the following information:

- A pie chart of the distribution of sessions by that section's category.

The definitions of the pie chart sections are listed below the chart. The session count appears next to each chart section.

- A table of the data shown in the pie chart.

RELATED DOCUMENTATION

[Understanding the Scope of Reports You Can Create | 1474](#)

[Creating Reports | 1478](#)

[Understanding Wireless Interference | 914](#)

[Managing Generated Reports | 1489](#)

[Managing Reports on SCP Servers | 1492](#)

Device Inventory Report

The Device Inventory report is a standardized report generated in Network Director to show all devices that are visible to Network Director. There are two portions of the report: a report header and the report body. The contents of the report header are found in [Table 422](#).

Table 422: Device Inventory Report Header

NETWORK DIRECTOR REPORT:	The type of report; In this case, the Device Inventory report.
Generated On:	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2012 11:35 PM PST
Generated By:	The username of the user that generated the report.
Report Name:	The name of the report assigned by the user when the report was created.
Scope	<p>The report is generated from this view and node in the network.</p> <ul style="list-style-type: none"> • <i>Perspective</i>—Can be Logical view, Location view, or Device view of the network. • <i>Node</i>—Represents the selected object that the report is based.
Report Filters	<p>The report specified these device and connection state filters.</p> <p>Device Type—Supported device types are:</p> <ul style="list-style-type: none"> • EX Series switches • Wireless LAN controllers • Wireless access points • QFX • QFabric • All (supported device types) <p>Connection State—Device connection states for filtering are:</p> <ul style="list-style-type: none"> • Up • Down • N/A • All (connection states)

The body of the report comprises:

- Device Type Count, which is a pie chart that graphically represents the network composition. Each segment represents:
 - EX Series switches
 - Wireless LAN controllers
 - Wireless access points
- Device Type Summary, which gives the total count of each type of device covered in the node.
- Device details by logical groups.

Following the pie chart, details for each device segment are listed by device type. For descriptions of the device fields see [Table 423](#).

Table 423: Inventory Report Fields

Field	Description
HostName	The device label.
IP Address	Either the IPv4 or the IPv6 address. NOTE: Not applicable for access points.
Model	The full model number of the EX Series switch.
Software Version	The Junos software version and release number.
Serial No	The hardware serial number of the device.
Device Type	The type of hardware, such as switches, controllers, or access points. Switches can also be designated as NORMAL for standalone or VC for Virtual Chassis.
Connection State	The connection state of the switch.
Configuration State	The administrative or operational state of the device.

RELATED DOCUMENTATION

[Understanding the Scope of Reports You Can Create | 1474](#)

[Creating Reports | 1478](#)

[Managing Generated Reports | 1489](#)

[Managing Reports on SCP Servers | 1492](#)

[Viewing the Device Inventory Page | 1136](#)

[Network Director Documentation home page](#)

Fabric Analyzer Report

The Fabric Analyzer report is a standardized report generated in Network Director to show information about a QFabric system, Virtual Chassis Fabric (VCF) or Virtual Chassis. There are two portions of the report: a report header and the report body. The contents of the report header are described in [Table 424](#).

Table 424: Fabric Analyzer Report Header

Field	Description
NETWORK DIRECTOR REPORT:	The type of report; in this case, the Fabric Analyzer report.
Generated On:	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2013 11:35 PM PST.
Generated By:	The username of the user that generated the report.
Report Name	The name of the report assigned when the report was created.
Scope	The report is generated from this view and node in the network. <ul style="list-style-type: none"> • <i>Perspective</i>—Can be a Logical View, Location View, or a Device View of the network. • <i>Node</i>—Represents the selected object on which the report is based.

The contents of the report body depend on the fabric type:

For a QFabric system, the report body contains these sections:

- **Data Plane Health Check**—Shows the results of connectivity checks between each node and the fabric's Interconnect devices and all other nodes.
- **Connectivity Check between Interconnects and Node Devices**—Shows the results of connectivity checks between each node each Interconnect device.
- **Connectivity Check between Node Devices**—Shows the results of connectivity checks between each node device and each other node device.
- **Data Plane Topology**—Shows a diagram of the data plane topology.

- Control Plane Topology—Shows a diagram of the control plane topology.

For a VCF, the report body contains these sections:

- Virtual Chassis Connectivity Status—Shows summary and status information about the VCF and its members.
- Port Bandwidth Utilization—Shows information about the bandwidth used by each link between member devices.
- VCF Health Check—Shows the connection status between each leaf device and the spine devices.

For a Virtual Chassis, the report body contains these sections:

- Virtual Chassis Connectivity Status—Shows summary and status information about the Virtual Chassis and its members.
- Port Bandwidth Utilization—Shows information about the bandwidth used by each link between member devices.

RELATED DOCUMENTATION

[Understanding the Scope of Reports You Can Create | 1474](#)

[Creating Reports | 1478](#)

[Understanding Wireless Interference | 914](#)

[Managing Generated Reports | 1489](#)

[Managing Reports on SCP Servers | 1492](#)

[Network Director Documentation home page](#)

Network Device Traffic Report

IN THIS SECTION

- [Network Device Traffic Report Header | 1515](#)
- [Network Device Traffic Charts | 1515](#)

The Network Device Traffic report is a standardized report generated in Network Director to show the device traffic for a device. There are two portions of the report: a report header and the report body.

This topic describes:

Network Device Traffic Report Header

The contents of the report header are found in [Table 425](#).

Table 425: Network Device Traffic Report Header

Field	Description
NETWORK DIRECTOR REPORT:	The type of report; In this case, the Network Device Traffic report.
Generated On:	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2012 11:35 PM PST.
Generated By:	The username of the user that generated the report.
Report Name	The name of the report assigned by the user when the report was created.
Scope	<p>The report is generated from this view and node in the network.</p> <ul style="list-style-type: none"> • <i>Perspective</i>—Can be Logical view, Location view, or a device view of the network. • <i>Node</i>—Represents the selected object upon which the report is based. <p>NOTE: If you select a device that is down, the report might not contain any data.</p>
Report Filters	The period of time specified for data collection.

Network Device Traffic Charts

The body of the Network Device Traffic report is a series of four colored charts that show a comparison of data or trend information about the packets. The charts are:

- Unicast Vs Non-unicast

This pie-chart shows the percentage totals for packets over the specified period of time at the node:

- Inbound unicast packets
- Inbound non-unicast (such as broadcast and multicast) packets
- Outbound unicast packets
- Outbound non-unicast packets

The percentage of non-unicast packets is normally less than that of unicast packets. If the percentage of non-unicast packets is as high or higher than that of the unicast percentage, it means that too many non-unicast packets are being sent in the network.

- **Unicast Vs Non-Unicast Trend**

This line chart shows the trend in unicast and non-unicast packets over the specified period of the report. The x axis shows the polling period; the y axis shows the number of packets. Use this chart to see if the plots are symmetric or asymmetric. It can also be useful for identifying unusual patterns.

- **Traffic Trend**

This line chart shows the overall trend of all packets over the specified period of time. The x axis shows the polling period; the y axis shows the number of packets. Use this chart to find abnormalities in the traffic trend.

- **Error Trend**

This line chart shows the errors over the specified period of time. An error is caused by a missing packet. Missing packets can be a result of: packet loss in the network, uncorrectable packet out of sequence, packet length error, jitter buffer overflow, or jitter buffer underflow. Use this chart to see the overall trend in errors. The x axis shows the polling period; the y axis shows the number of errors.

RELATED DOCUMENTATION

[Understanding the Scope of Reports You Can Create | 1474](#)

[Creating Reports | 1478](#)

[Managing Generated Reports | 1489](#)

[Managing Reports on SCP Servers | 1492](#)

[Network Director Documentation home page](#)

Network Neighborhood Report

IN THIS SECTION

- [Network Neighborhood Report Header | 1517](#)

- [Network Neighborhood Report Tables | 1517](#)

The Network Neighborhood report is a standardized report generated in Network Director to show the Received Signal Strength Indication (RSSI) data for a location or for an access point cluster. The report is available in all views (Logical, Location, and Device), however it supports only selecting a Floor or Outdoor area in Location View, or a node that contains an access point or an access point cluster in Logical or Device view. There are two portions of the report: a report header and the report body.

This topic describes:

Network Neighborhood Report Header

The contents of the report header are described in [Table 426](#).

Table 426: Network Neighborhood Report Header

Field	Description
NETWORK DIRECTOR REPORT:	The type of report; In this case, the Network Neighborhood report.
Generated On:	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2013 11:35 PM PST.
Generated By:	The username of the user that generated the report.
Report Name	The name of the report assigned by the user when the report was created.
Scope	<p>The report is generated from this view and node in the network.</p> <ul style="list-style-type: none"> • <i>Perspective</i>—Can be Logical view, Location view, or a device view of the network. • <i>Node</i>—Represents the selected object upon which the report is based. <p>NOTE: Selecting a node that does not contain wireless equipment can cause this report to be blank.</p>

Network Neighborhood Report Tables

The body of the Network Neighborhood report is a series of four tables that provide information on radios within a selected view. The first two tables provide the RSSI values of the listeners for the radios. RSSI is the power level being received by the antenna. The last two tables list the radios that are being heard by the selected radio, which are grouped by channels.

The Radio 2.4 GHz RSSI Mapping table shows the Radio1NG values of listeners for the radios in the selected scope. The column headings of the table show the radio identifier, which is the access point name,

a colon, and the radio number. The table rows represent the lists the 2.4-GHz radios available in the scope. The signal strength is shown in decibels, with the stronger signals being less negative, or closer to zero.

The Radio 5 GHz RSSI Mapping table shows the Radio2NA values of listeners for the radios in the selected scope. The column headings of the table show the radio identifier, which is the access point name, a colon, and the radio number. The table rows represent the lists the 5-GHz radios available in the scope. The signal strength is shown in decibels, with the stronger signals being less negative, or closer to zero.

The Radio 2.4 GHz Channel Mapping and the Radio 5 GHz Channel Mapping tables shows which access points are being heard on a channel. The Channel column list the channels using the frequency; the APs Heard column list the radio identifiers being heard on the channel.

RELATED DOCUMENTATION

[Monitoring the RF Neighborhood | 1340](#)

[Network Director Documentation home page](#)

Radio Traffic Report

The Radio Traffic report is a standardized report generated in Network Director to show data about wireless radio traffic across the network over time. This data can be used to identify heavily used areas of the network or identify trends on network usage. The applicable scopes for this report are: My Network, Mobility Domain, Wireless Network, wireless controller, Site, Building, Outdoor Area, Floor, wireless access point, and Radio.

There are two portions of the report: a report header and the report body. The contents of the report header are described in [Table 427](#).

Table 427: Radio Traffic Report Header

Field	Description
NETWORK DIRECTOR REPORT:	The type of report; in this case, the Client Details report.
Generated On:	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2013 11:35 PM PST.
Generated By:	The username of the user that generated the report.
Report Name	The name of the report assigned when the report was created.

Table 427: Radio Traffic Report Header (continued)

Field	Description
Scope	<p>The report is generated from this view and node in the network.</p> <ul style="list-style-type: none"> • <i>Perspective</i>—Can be a Logical View, Location View, or a Device View of the network. • <i>Node</i>—Represents the selected object on which the report is based.
Report	The filter options specified for the report.

The report body contains a section for each wireless access point within the report's scope, unless the scope is one wireless controller or wireless access point. In that case, the information is presented at the radio level. Each of these sections shows the following information about the wireless access point or radio:

- **Throughput Chart**—A line chart that shows the throughput rate over time of the wireless access point or radio.
- **Number of Clients Chart**—A line chart that shows the number of clients over time on the wireless access point or radio.

RELATED DOCUMENTATION

[Understanding the Scope of Reports You Can Create | 1474](#)

[Creating Reports | 1478](#)

[Understanding Wireless Interference | 914](#)

[Managing Generated Reports | 1489](#)

[Managing Reports on SCP Servers | 1492](#)

[Network Director Documentation home page](#)

Port Bandwidth Utilization Report

The Port Bandwidth Utilization report is a standardized report generated in Network Director to show data for the last polled interval. There are two portions of the report: a report header and the report body. The contents of the report header are found in [Table 428](#).

Table 428: Port Bandwidth Utilization Report Header

Field	Description
NETWORK DIRECTOR REPORT:	The type of report; In this case, the Port Bandwidth Utilization report.
Generated On:	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2013 11:35 PM PST.
Generated By:	The username of the user that generated the report.
Report Name	The name of the report assigned when the report was created.
Scope	The report is generated from this view and object in the network. <ul style="list-style-type: none"> • <i>View</i>—Can be a Logical view, Location view, or a Device view of the network. • <i>Object</i>—Represents the selected object on which the report is based.
Report Filters	Shows the Percentage Utilization Exceeding value specified for the report. Only ports that exceed this percentage of their allocated bandwidth appear in the report.

The report contains a table for each device that contains ports that exceeded the specified percentage of their allocated bandwidth. The fields in these tables are described in [Table 429](#).

Table 429: Port Bandwidth Utilization Report Fields

Field	Description
Host Name	Host name of the device.
IP Address	IP address of the device.
Device Type	Device type.
Port Name	Port name.
Percentage Utilization	Percentage of allocated bandwidth the port used.

RELATED DOCUMENTATION

[Understanding the Scope of Reports You Can Create | 1474](#)

[Creating Reports | 1478](#)

[Managing Generated Reports | 1489](#)

[Managing Reports on SCP Servers | 1492](#)

[Network Director Documentation home page](#)

RF Interference Detail Report

The RF Interference Detail report is a standardized report generated in Network Director to show detailed information about RF interference on a wireless network. There are two portions of the report: a report header and the report body. The contents of the report header are described in [Table 430](#).

Table 430: RF Interference Detail Report Header

Field	Description
NETWORK DIRECTOR REPORT:	The type of report; in this case, the RF Interference Detail report.
Generated On:	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2013 11:35 PM PST.
Generated By:	The username of the user that generated the report.
Report Name	The name of the report assigned when the report was created.
Scope	The report is generated from this view and node in the network. <ul style="list-style-type: none"> • <i>Perspective</i>—Can be a Logical View, Location View, or a Device View of the network. • <i>Node</i>—Represents the selected object on which the report is based.

The report body contains one or more tables of detailed information about RF interference. [Table 431](#) describes the table columns.

Table 431: RF Interference Detail Report Tables

Column	Description
Interference Source Type	The source type of the RF interference. Possible sources of interference include Bluetooth, Continuous Wave, Microwave Oven, Unknown, and Phone FHSS.
Last Seen	Date and time the interference was last detected.

Table 431: RF Interference Detail Report Tables (*continued*)

Column	Description
Transmitter ID	If the interference is caused by an object with a MAC address, the MAC address is displayed. If the object has no MAC address, MSS calculates a MAC address, using the characteristics of the object. This way, you can correlate interference events over time.
AP	Name of the access point that detected the interference.
Listener MAC	MAC address of the access point that detected the interference.
Controller	Name of the controller that reported the interference.
RSSI	Received signal strength indication (RSSI), in decibels referred to 1 milliwatt (dBm). A higher value indicates a stronger signal.
Duty Cycle	Reported fraction of time that the source is emitting RF.
CIM	Estimated severity of interference on this channel caused by the source.

RELATED DOCUMENTATION

[Understanding the Scope of Reports You Can Create | 1474](#)

[Creating Reports | 1478](#)

[Understanding Wireless Interference | 914](#)

[Managing Generated Reports | 1489](#)

[Managing Reports on SCP Servers | 1492](#)

[Network Director Documentation home page](#)

RF Interference Summary Report

The RF Interference Summary report is a standardized report generated in Network Director to show a summary of RF interference on a wireless network. There are two portions of the report: a report header and the report body. The contents of the report header are described in [Table 432](#).

Table 432: RF Interference Summary Report Header

Field	Description
NETWORK DIRECTOR REPORT:	The type of report; in this case, the Active User Sessions report.
Generated On:	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2013 11:35 PM PST.
Generated By:	The username of the user that generated the report.
Report Name	The name of the report assigned when the report was created.
Scope	The report is generated from this view and node in the network. <ul style="list-style-type: none"> • <i>Perspective</i>—Can be a Logical View, Location View, or a Device View of the network. • <i>Node</i>—Represents the selected object on which the report is based.

The report body contains these elements:

- A table titled Interference Summary Details, which contains with the columns described in [Table 433](#).
- A pie chart titled RF Interference Summary that shows the distribution of interference source types.

Table 433: RF Interference Summary Report Interference Summary Details Table

Column	Description
Interference Source Type	The source type of the RF interference.
Count	The number of instances of RF interference of that source type.

RELATED DOCUMENTATION

[Understanding the Scope of Reports You Can Create | 1474](#)

[Creating Reports | 1478](#)

[Understanding Wireless Interference | 914](#)

[Managing Generated Reports | 1489](#)

[Managing Reports on SCP Servers | 1492](#)

[Network Director Documentation home page](#)

Rogue Summary Report

The Rogue Summary report is a standardized report generated in Network Director to show a summary of rogue device distribution on the network and distribution of rogue devices by SSID. Rogue devices can be a security threat by allowing unauthorized access to the network and prompt action is required to remove the potential threat. The applicable scope for this report is My Network.

There are two portions of the report: a report header and the report body. The contents of the report header are described in [Table 434](#).

Table 434: Rogue Summary Report Header

Field	Description
NETWORK DIRECTOR REPORT:	The type of report; in this case, the Rogue Summary report.
Generated On:	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2013 11:35 PM PST.
Generated By:	The username of the user that generated the report.
Report Name	The name of the report assigned when the report was created.
Scope	The report is generated from this view and node in the network. <ul style="list-style-type: none"> • <i>Perspective</i>—Can be a Logical View, Location View, or a Device View of the network. • <i>Node</i>—Represents the selected object on which the report is based.
Report Filters	Shows any filters that were applied to the report definition. Only results that match the filters are included in the report.

The report body contains information about the rogue devices detected on the network. It contains these sections:

- **Summary**—Shows information about all the rogue devices included in the report.
- **Distribution by Switch**—Shows the distribution of rogue devices by switch.
- **Distribution by SSID**—Shows the distribution of rogue devices by SSID.
- **List by SSID**—Shows the rogue devices detected in each SSID, identified by MAC address.

RELATED DOCUMENTATION

[Understanding the Scope of Reports You Can Create | 1474](#)

[Creating Reports | 1478](#)

[Understanding Wireless Interference | 914](#)

[Managing Generated Reports | 1489](#)

[Managing Reports on SCP Servers | 1492](#)

[Network Director Documentation home page](#)

Wireless Security Alarms Report

The Wireless Security Alarms report is a standardized report generated in Network Director to show security alarm activity. It provides information about alarms for unauthorized wireless networks, denial of service (DoS) attacks, and Intrusion Detection Systems (IDS). These alarm types alert you to possible network security issues. The applicable scope for this report is My Network.

There are two portions of the report: a report header and the report body. The contents of the report header are described in [Table 435](#).

Table 435: Wireless Security Alarms Report Header

Field	Description
NETWORK DIRECTOR REPORT:	The type of report; in this case, the Wireless Security report.
Generated On:	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2013 11:35 PM PST.
Generated By:	The username of the user that generated the report.
Report Name	The name of the report assigned when the report was created.
Scope	The report is generated from this view and node in the network. <ul style="list-style-type: none"> • <i>Perspective</i>—Can be a Logical View, Location View, or a Device View of the network. • <i>Node</i>—Represents the selected object on which the report is based.

The report body contains these sections:

- **Unauthorized Networks**—Contains a table and chart of detected unauthorized networks, divided into network types.
- **DoS Alarms Summary**—Contains a chart of active DoS security alarms.

- Detected Networks—Contains the following:
 - A line chart of detected devices over time.
 - A line chart of security alarms over time.
 - A table showing the number of active security alarms of each type.

RELATED DOCUMENTATION

[Understanding the Scope of Reports You Can Create | 1474](#)

[Creating Reports | 1478](#)

[Understanding Wireless Interference | 914](#)

[Managing Generated Reports | 1489](#)

[Managing Reports on SCP Servers | 1492](#)

[Network Director Documentation home page](#)

Top Users by Data Usage Report

IN THIS SECTION

- [Top Users by Data Usage Header | 1526](#)
- [Top Users of Data Table | 1527](#)

The Top Users by Data Usage report is a standardized report generated in Network Director. Use this report to identify the users with the highest data usage at the specified node during the specified time frame.

This topic describes:

Top Users by Data Usage Header

The contents of the report header are found in [Table 436](#).

Table 436: Top 10 Users by Data Usage Report Header

Field	Description
NETWORK DIRECTOR REPORT:	The type of report; In this case, the Top Users by Data Usage report.
Generated On:	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2012 11:35 PM PST.
Generated By:	The username of the user that generated the report.
Report Name	The name of the report assigned by the user when the report was created.
Scope	The report is generated from this view and node in the network. Perspective can be Logical view, Location view, or Device view of the network. Node represents the selected object that the report is based.
Report Filters	The count specified for the report generation. The default count is 10 users.

Top Users of Data Table

The body of the Top Users by Data Usage report is a snapshot of the users with the highest data usage. The number of users analyzed and the time interval is determined by the Reporting Options set during report definition. The report sorts users in the table from the user with the highest bandwidth usage to the least highest. The key fields of the table are described in [Table 437](#).

Table 437: Top 10 Users by Bandwidth Report Fields

Field	Description
User Name	The User ID with the highest bandwidth usage during the specified period.
Client MAC Address	The MAC address of the client.
Number of Sessions	The total sessions during this time period.
Total Data Used	The bandwidth used in megabytes.

RELATED DOCUMENTATION

[Top Sessions by MAC Address Monitor](#) | 1438

Traffic and Congestion Summary Report

The Traffic and Congestion Summary report is a standardized report generated in Network Director to show information about latency and port utilization on network devices. It provides detailed and trended information about latency and port utilization at the network, device and port level. This report supports any scope that includes devices that support high-frequency statistics.

There are two portions of the report: a report header and the report body. The contents of the report header are described in [Table 438](#).

Table 438: Traffic and Congestion Summary Report Header

Field	Description
NETWORK DIRECTOR REPORT:	The type of report; in this case, the Traffic and Congestion Summary report.
Generated On:	The date and time the report ran. The date format is: [Day of the Month] [Month] [Year]. The time format follows the Coordinated Universal Time (UTC). For example: 07 December 2013 11:35 PM PST.
Generated By:	The username of the user that generated the report.
Report Name	The name of the report assigned when the report was created.
Scope	The report is generated from this view and node in the network. <ul style="list-style-type: none"> • <i>Perspective</i>—Can be a Logical View, Location View, or a Device View of the network. • <i>Node</i>—Represents the selected object on which the report is based.
Report Filters	Shows any filters that were applied to the report definition.

The report body shows port utilization and latency information. If the report scope is My Network, the report contains sections that summarize port utilization and latency for all devices, then a section that shows more information about the top devices with the highest port utilization and latency. If the report scope is one device, the report shows more detailed information about that device, including port-level utilization and latency.

RELATED DOCUMENTATION

[Understanding the Scope of Reports You Can Create | 1474](#)

[Creating Reports | 1478](#)

[Understanding Wireless Interference | 914](#)

[Managing Generated Reports | 1489](#)

[Managing Reports on SCP Servers | 1492](#)

[Network Director Documentation home page](#)

8

PART

Working with Network Director Mobile

[About Network Director Mobile | 1531](#)

[Getting Started with Network Director Mobile | 1532](#)

[Working in the Network Director Mobile Dashboard Mode | 1535](#)

[Working in the Network Director Mobile Devices Mode | 1540](#)

About Network Director Mobile

IN THIS CHAPTER

- [Overview of Network Director Mobile | 1531](#)

Overview of Network Director Mobile

Network Director Mobile is a Network Director user interface that is optimized to run in a mobile browser. It enables you to use Network Director monitoring features on a mobile device.

Network Director Mobile provides a Dashboard View that summarizes information about your entire network. It also enables you to drill down into individual devices for detailed information about those devices and the devices and sessions they manage.

RELATED DOCUMENTATION

[Network Director Mobile System Requirements | 1532](#)

[Logging Into Network Director Mobile | 1533](#)

[Understanding the Network Director Mobile User Interface | 1533](#)

Getting Started with Network Director Mobile

IN THIS CHAPTER

- [Network Director Mobile System Requirements | 1532](#)
- [Logging Into Network Director Mobile | 1533](#)
- [Understanding the Network Director Mobile User Interface | 1533](#)
- [Configuring Network Director Mobile Settings | 1534](#)

Network Director Mobile System Requirements

Network Director Mobile runs in a mobile web browser on tablet devices. It has the following system requirements:

- On Apple iPad2, iPad3, and iPad mini devices:
 - Operating system versions 6.1.3 and 7.0.
 - Apple Safari browser versions included with the supported operating system versions.
- On Android tablet devices:
 - Android version 4.1.
 - Google Chrome browser version 29.0.1547.59 and higher.

RELATED DOCUMENTATION

[Overview of Network Director Mobile | 1531](#)

[Logging Into Network Director Mobile | 1533](#)

Logging Into Network Director Mobile

Network Director Mobile runs in a mobile browser. To log in to the Network Director server, navigate to this URL:

https://<server>/networkdirector/mobile, where <server> is the IP address or hostname of the Network Director server. Log in using your Network Director username and password.

RELATED DOCUMENTATION

[Network Director Mobile System Requirements | 1532](#)

Understanding the Network Director Mobile User Interface

IN THIS SECTION

- [Dashboard Mode | 1533](#)
- [Devices Mode | 1533](#)

Network Director Mobile is a Network Director user interface that is optimized to run in a mobile browser. It enables you to use Network Director monitoring and fault management features on a mobile device.

The user interface has two modes: Dashboard and Devices. Buttons to access the modes are always available at the bottom of the page. When you log in to the server, Dashboard mode is open by default.

On any page that has a Back button, select the **Back** button to return to the previous page.

These sections describe the modes:

Dashboard Mode

Dashboard mode contains monitors that show information about your entire network.

Devices Mode

Devices mode enables you to drill down into individual devices for detailed information about these devices and the devices and sessions they manage.

RELATED DOCUMENTATION

[Monitoring Network-Wide Activity Using Network Director Mobile | 1535](#)

[Locating a Device and Viewing Device Properties Using Network Director Mobile | 1540](#)

[Configuring Network Director Mobile Settings | 1534](#)

[Overview of Network Director Mobile | 1531](#)

Configuring Network Director Mobile Settings

To configure Network Director Mobile settings:

1. Select the settings button in the main banner.

A dialog box opens.

2. Select the **General** tab to configure general settings:

- Refresh Interval—Select how often the application refreshes its data from the Network Director server.

3. Select the **Sessions** tab to configure sessions settings:

- Session Timeout—Select how long the application will wait before it logs off the session automatically if there is no user activity.

RELATED DOCUMENTATION

[Overview of Network Director Mobile | 1531](#)

[Understanding the Network Director Mobile User Interface | 1533](#)

Working in the Network Director Mobile Dashboard Mode

IN THIS CHAPTER

- [Monitoring Network-Wide Activity Using Network Director Mobile | 1535](#)
- [Network Director Mobile Dashboard Reference | 1535](#)

Monitoring Network-Wide Activity Using Network Director Mobile

Use Dashboard mode to monitor network-wide activity. To open Dashboard mode, select the **Dashboard** button that is always available at the bottom of the page.

RELATED DOCUMENTATION

| [Network Director Mobile Dashboard Reference | 1535](#)

Network Director Mobile Dashboard Reference

IN THIS SECTION

- [Network Summary Monitor | 1536](#)
- [Alarms Monitor | 1536](#)
- [Top Sessions Monitor | 1537](#)
- [Ports Monitor | 1538](#)
- [Session Count Monitor | 1538](#)
- [Session Trend Monitor | 1538](#)
- [RF Interferences Monitor | 1538](#)

The Dashboard contains monitors that show information about your entire managed network:

Network Summary Monitor

The Network Summary monitor contains these pie charts:

- **Devices By Family**—Shows the distribution of devices based on device family.
- **Connection State**—Shows the distribution of devices based on the status of the device's connection to the Network Director server. The possible connection states are:
 - **UP**—Device is connected to Network Director.
 - **DOWN**—Device is not connected to Network Director.
 - **N/A**—Device connection state is not available.
- **Configuration State**—Shows the distribution of devices based on whether the Network Director configuration is in sync with the device configuration. The possible configuration states for a device depend on its connection state:
 - When connection state is **UP**, the configuration state can be **Out of Sync**, **Synchronizing**, **In Sync**, or **Sync Failed**.
 - When connection state is **DOWN**, the configuration state is **N/A**.

Alarms Monitor

The Alarms monitor provides a quick summary of the critical, major, minor, and info alarms currently active in the network.

Select the **Expand** button in the right corner of the title bar to see detailed information about the alarms on the Alarms page.

On the Alarms page, you have the following options:

- Select the **Graph** or **List** buttons to view the information in those formats.
 - For a description of the information presented in the list view, see [Table 439](#).
- In the graph view, you can select **By Severity**, **By Category**, or **By State** to see the distribution of active alarms by those properties.

Select **Back** to return to the Dashboard.

Table 439: Network Director Mobile Alarm Details Fields

Field	Value
Name	The alarm name.

Table 439: Network Director Mobile Alarm Details Fields (continued)

Field	Value
Alarm Severity	The severity of the alarm. Severity levels are: <ul style="list-style-type: none"> • Critical—A critical condition exists; immediate action is necessary. • Major—A major error has occurred; escalate or notify as necessary. • Minor—A minor error has occurred; notify or monitor the condition. • Info—An informational message; no action is necessary.
Entity ID	The identification of the entity responsible for causing this alarm. The Entity ID is the key for correlating events into an alarm. The Entity ID could be a MAC address of a radio or an IP address of the device.
Assigned to	If the alarm is assigned to an individual, it shows the name of that person; otherwise, it shows System to mark that the alarm is still unassigned.
Last Updated On	The date and time that the information for the alarm was last modified.

Top Sessions Monitor

The Top Sessions monitor contains a bar chart showing the eight user sessions that are currently using the most bandwidth.

Select the **Expand** button in the top right corner of the title bar to see detailed information on the Top Sessions page. On the Top Sessions page, you have the following options:

- Select **Top Sessions By User** to see sessions that are identified by their user.
- Select **Top Sessions By MAC** to see sessions that are identified by their MAC address.
- Select the **Graph** or **List** buttons to view the information in those formats.

For a description of the information presented in the list view, see [Table 440](#).

- Select the time period to display from the list in the title bar.

Select **Back** to return to the Dashboard.

Table 440: Network Director Mobile Top Session Details

Table Column	Description
Username	Client's user name.
Total Data Usage (KBytes)	The session's total data usage. Appears when any time period other than Current is selected.

Table 440: Network Director Mobile Top Session Details (continued)

Table Column	Description
Number of Sessions	Number of sessions.

Ports Monitor

The Ports monitor shows information about the network's device ports:

- Admin Status—Shows the number of ports that are up and down.
- Free Vs Used—Shows the number of ports that are free and the number that are used.

Select the **Expand** button in the top right corner of the title bar to see detailed information on the Ports page. On the Ports page, you have the option to view ports by admin status or by free versus used status.

Select **Back** to return to the Dashboard.

Session Count Monitor

The Session Count monitor shows the number of active user sessions on the network.

Session Trend Monitor

The Session Trend monitor contains a line graph that shows the number of active user sessions on the network over time.

RF Interferences Monitor

The RF Interferences monitor provides information about radio frequency (RF) interferences that the wireless network has detected. The monitor contains a pie chart showing the distribution of the sources of detected RF interferences.

Select the **Expand** button in the top right corner of the title bar to see detailed information on the RF Interference Sources page. On the RF Interference Sources page, you can select the **Graph** or **List** buttons to view the information in those formats.

For a description of the information presented in the list view, see [Table 441](#).

Select **Back** to return to the Dashboard.

Table 441: Network Director Mobile RF Interference Sources Details

Information	Description
Last Seen	Date and time the interference was last detected.
Transmitter Id	If the interference is caused by an object that has a MAC address, the MAC address is displayed. If the object has no MAC address, MSS calculates a MAC address by using the characteristics of the object. This way, you can correlate interference events over time.
Listener Id	MAC address of the access point that detected the interference.
Channel	Channel affected by the interference.
Duty Cycle	Reported fraction of time that the source is emitting RF.
Type	Possible sources of interference include Bluetooth, Continuous Wave, Microwave Oven, Unknown, and Phone FHSS.

RELATED DOCUMENTATION

[Monitoring Network-Wide Activity Using Network Director Mobile | 1535](#)

Working in the Network Director Mobile Devices Mode

IN THIS CHAPTER

- Locating a Device and Viewing Device Properties Using Network Director Mobile | 1540
- Monitoring Sessions on a Device Using Network Director Mobile | 1542
- Monitoring Equipment Status on a Wireless LAN Controller Using Network Director Mobile | 1543

Locating a Device and Viewing Device Properties Using Network Director Mobile

You can locate a device and view its properties by searching or by browsing.

To locate a device and view its properties:

1. Select the **Devices** button at the bottom of the page to open Devices mode.
2. To locate the device by searching:
 - a. Enter search text in the search box (it contains the text Hostname or IP until you enter text).
 - b. Select the search button.
 - c. Locate the device in the list of search results.
 - d. For information about the device properties shown, see [Table 442](#).
3. To locate the device by browsing:
 - a. Select the device type (**Switches**, **Wireless Controllers**, or **Fabric**).
 - b. If you selected Switches or Wireless Controllers, select the device family from the list, then locate the device in the list of devices that opens.

- c. If you selected Fabric:
 - i. Locate the QFabric device in the list.
 - ii. Select the arrow next to the device to browse the list of its components.
 - iii. Select a component type from the list, then locate the device in the list of devices that opens.
- d. For information about the device properties shown, see [Table 442](#).

Table 442: Device Properties Shown in Network Director Mobile Inventory

Field	Description
Hostname	Configured name of the device.
Device Family	Device family of the device. For example, MSS, EX, or WLC. Shown only on inventory pages created by searching.
Platform	Model number of the device. Shown only on inventory pages created by searching.
Model	Type of the device. Shown only on inventory pages that you browse to, not on search results pages.
Mgmt IP	IP Address of the device.
Mgmt Status	Displays whether the device is directly manageable or not.
Connection Status	Connection status of the device in Network Director: <ul style="list-style-type: none"> ● UP-Device is connected to Network Director. ● DOWN-Device is not connected to Network Director. ● N/A-Device connection status is not available.
Serial Number	Serial number on device chassis.

RELATED DOCUMENTATION

[Monitoring Sessions on a Device Using Network Director Mobile | 1542](#)

[Monitoring Equipment Status on a Wireless LAN Controller Using Network Director Mobile | 1543](#)

Monitoring Sessions on a Device Using Network Director Mobile

To monitor session activity on a device:

1. Locate the device as described in [“Locating a Device and Viewing Device Properties Using Network Director Mobile” on page 1540.](#)

2. Select the device from the list.

3. Select **Session Details**.

The Session Details page for the device opens. You can select the **Graph** or **List** buttons to view the information in those formats.

4. To see historical session data, select **Session Trend**.

You can select the time period to view by selecting a time period from the list in the page’s title bar.

For a description of the information presented in the list view, see [Table 443](#).

5. To see current session data, select **Current Sessions**.

For a description of the information presented in the list view, see [Table 444](#).

6. (On wireless LAN controllers only) To see the ten sessions that are using the most bandwidth, select **Top 10 Sessions**.

For a description of the information presented in the list view, see [Table 445](#).

Table 443: Session Trend Details

Table Column	Description
Time	Time when a poll occurred.
Min Session Count	Minimum session count for the time period.
Avg Session Count	Average session count.
Max Session Count	Maximum session count for the time period.

Table 444: Current Session Details

Table Column	Description
Username	Client’s user name

Table 444: Current Session Details (*continued*)

Table Column	Description
MAC Address	Client's MAC address.
Controller IP	IP address of the controller to which the client is connected.
AP ID	Name of the wireless access point to which the client is connected.
SSID	SSID to which a wireless client is connected.
Incremental Data Usage (KBytes)	The session's current incremental data usage.

Table 445: Top 10 Session Details

Table Column	Description
Username	Client's user name
MAC Address	Client's MAC address.
Time	Time when the session data was polled from the device.
Incremental Data Usage (KBytes)	The session's current incremental data usage.

RELATED DOCUMENTATION

[Locating a Device and Viewing Device Properties Using Network Director Mobile | 1540](#)

[Monitoring Equipment Status on a Wireless LAN Controller Using Network Director Mobile | 1543](#)

Monitoring Equipment Status on a Wireless LAN Controller Using Network Director Mobile

You can monitor the status and session activity of equipment managed by a wireless LAN controller (including wireless access points and radios).

To monitor equipment status on a wireless LAN controller:

1. Locate the device as described in [“Locating a Device and Viewing Device Properties Using Network Director Mobile” on page 1540](#).
2. Select the device in the list.
3. Select **Equipment Status**.

The Equipment Status page for the device opens. This page lists all the wireless access points that the wireless LAN controller manages.

For a description of the information shown, see [Table 446](#).

4. Select a wireless access point from the list to view details about it in the lower pane.
5. To see detailed information about the wireless access point:
 - a. Select **AP Details**.
For a description of the information shown, see [Table 447](#).
 - b. To see information about the sessions managed by the selected access point, select any row in the AP Details table, then select **AP Session Details**.
 - c. Select **Top 10 Sessions**, **Session Trend**, or **Current Sessions** to select how to view the session information for the access point.
6. To see information about the wireless access point's radios:
 - a. Select **Radio Status**.
For a description of the information shown, see [Table 448](#).
 - b. To see information about the sessions managed by the selected radio, select **Radio Session Details**.
Select **Top 10 Sessions**, **Session Trend**, or **Current Sessions** to select how to view the session information for the access point.
 - c. To see RF information about the selected radio, select **RF Details**.
Select **Interference Sources** or **RF Neighborhood** to select which RF information to view for the access point.

Table 446: Wireless LAN Controller Equipment Status in Network Director Mobile

Table Column	Description
AP Name	Wireless access point's name.
AP Serial ID	Wireless access point's serial number.
AP Model	Wireless access point's model.
AP IP Address	Wireless access point's IP address.
AP Status	Operational status of the wireless access point: <ul style="list-style-type: none"> • Down—The access point is offline. • Up—The access point is online and enabled. • Up Redundant—The access point is online, reporting to this controller as redundant and to another controller as primary.

Table 447: Wireless Access Point Details in Network Director Mobile

Name	Value
AP Up Time	The length of time since the access point last booted.
AP Boot Loader	Wireless access point's boot loader.
Primary AP Manager	Wireless access point's primary manager.
Secondary AP Manager	Wireless access point's secondary manager.
Port1 Speed	Speed of network port 1.
Port2 Speed	Speed of network port 2.
Port1 DuplexMode	Duplex mode setting of network port 1.
Port2 DuplexMode	Duplex mode setting of network port 2.
Port 1 PoeStatus	PoE status of network port 1.
Port 2 PoeStatus	PoE status of network port 2.

Table 448: Radio Status in Network Director Mobile

Table Column	Description
AP Serial ID	Wireless access point's serial number.
Type	Radio type.
Channel	Channel the radio is using.
Mac Address	Radio's MAC address.
Status	Radio's status.

RELATED DOCUMENTATION

[Locating a Device and Viewing Device Properties Using Network Director Mobile | 1540](#)

[Monitoring Sessions on a Device Using Network Director Mobile | 1542](#)

9

PART

Working with Aruba Devices and Applications in Network Director

[About Aruba Networks Integration in Network Director | 1548](#)

[Managing Aruba Devices and Applications in Network Director | 1550](#)

About Aruba Networks Integration in Network Director

IN THIS CHAPTER

- [Understanding Aruba Airwave Integration with Network Director | 1548](#)

Understanding Aruba Airwave Integration with Network Director

Juniper Networks Junos Space Network Director integrates with Aruba Airwave management application and Aruba wireless LAN devices, including Aruba access points and controllers, and to converge Juniper Networks wired technologies such as switching and Aruba wireless LAN technologies for simplified network management.

The integration of these two platforms enables you to:

- View Aruba Device Inventory List in Network Director—Network Director retrieves the Aruba wireless LAN access points and wireless LAN controllers that are connected to Juniper Networks switches only. The Aruba wireless devices and the Juniper switches must be managed by the Aruba Airwave application for Network Director to populate the Aruba device inventory. The retrieval occurs based on the resynchronization interval specified under Preferences. Alternately, you can manually retrieve the device inventory list by clicking the **Resynchronize Now** option.
- Launch Aruba Airwave from within Network Director—You can launch the context-sensitive pages of Airwave application from within Network Director to manage Aruba wireless devices. For example, you can launch the APs/Devices > List and Reports > Generated pages of Airwave application from the Monitor mode and Report mode of Network Director respectively.

NOTE: The View Aruba device inventory and Launch Aruba Airwave options will be available in the Network user interface only after you specify the Airwave settings under Wireless preferences.

This integration leverages services such as tracking the number of users and types of devices accessing a network, and so on.

NOTE: It is recommended to configure and manage the Juniper Networks switches by using Network Director to which wireless Aruba devices are connected.

RELATED DOCUMENTATION

[Linking to the Aruba Airwave Application | 1553](#)

[Viewing Aruba Wireless Device Inventory in Network Director | 1550](#)

[Launching the Aruba Airwave Application | 1554](#)

[Understanding Wireless Network Management in Network Director | 82](#)

[Network Director Documentation home page](#)

Managing Aruba Devices and Applications in Network Director

IN THIS CHAPTER

- [Viewing Aruba Wireless Device Inventory in Network Director | 1550](#)
- [Linking to the Aruba Airwave Application | 1553](#)
- [Launching the Aruba Airwave Application | 1554](#)

Viewing Aruba Wireless Device Inventory in Network Director

The Aruba Wireless Device Inventory page in Network Director lists Aruba wireless LAN access points and wireless LAN controllers connected to Juniper Networks switches. Network Director retrieves the device inventory list based on the resynchronization interval specified under Preferences. After the resynchronization is completed successfully, Network Director populates only those Aruba wireless LAN access points and wireless LAN controllers that are connected to Juniper switches. The Aruba wireless devices and the Juniper switches must be managed by the Aruba Airwave application for Network Director to populate the Aruba device inventory. The Aruba Wireless Device Inventory page provides basic information about the devices, such as IP address and current operating status.

You can view the inventory list for the Aruba devices in the following modes of Network Director:

- **Build mode**—You navigate to **Device Management > View Aruba Wireless Device Inventory** under Build mode to view the all the Aruba wireless LAN access points and wireless LAN controllers that are connected to Juniper Networks switches.
- **Monitor mode**—You navigate to **Device Management > View Aruba Wireless Devices** in Monitor modes to view the inventory of all the Aruba wireless LAN access points and wireless LAN controllers that are connected to Juniper Networks switches.
- **Fault mode**—You can navigate to **Device Management > Impacted Aruba Wireless Devices** in Fault mode to view the inventory of Aruba wireless LAN access points and wireless LAN controllers that are connected to Juniper Networks switch ports that are down.

The scope that you select in the View pane determines which Aruba devices are listed in the Device Inventory page. For example:

- If you select My Network, all the Aruba devices connected to Juniper Networks switches are listed.
- If you select a building in Location view, only those access points connected to Juniper Networks switches assigned to that building (including the floors and closets in the building) are listed.

The Aruba device Inventory page provides two pie charts that summarize the status of the devices in your selected scope:

- **Devices by Category**—Indicates the proportion of devices such as Aruba wireless LAN controller, Instant Access Points (IAP), and others in each category.
- **Connection State**—Shows the proportion of devices that are up or down.

Mouse over a pie segment to view the actual number of devices and the percentage represented by that pie segment.

Clicking **Resynchronize Now** creates a job to synchronize the Juniper Networks and Aruba servers. After few moments, the latest Aruba device inventory with the current operating status of the devices is populated in this page.

You can also perform a search of wireless devices in the device inventory list by typing any of the field names of the device inventory table and its value in the Enter Search text box.

Table 449 describes the fields in the device inventory table:

Table 449: Fields in Aruba Device Inventory Table

Field	Description
Hostname	<p>Indicates the configured hostname of the Aruba wireless device or the IP address of the device if no hostname is configured.</p> <p>Clicking the hostname opens the APs/Devices > List page in Airwave application displaying the details of the selected access point.</p>
IP Address	Indicates the IP address of the Aruba device.
Serial Number	Indicates the serial number of the Aruba device.
Model	Indicates the Aruba device model number.
Device Type	<p>Indicates the type of device connected to the Juniper Networks switch. The following Aruba devices can be connected:</p> <ul style="list-style-type: none"> • Thin APs • Controllers • Instant Access Points
Connection State	<p>Indicates the connection status of the Aruba device in Network Director:</p> <ul style="list-style-type: none"> • UP—The Aruba device is connected to Network Director. • DOWN—The Aruba device is not connected to Network Director.
Controller Name	<p>Indicates the name of the Aruba wireless LAN controller.</p> <p>Clicking the controller name opens the APs/Devices > List page in Airwave application displaying the details of the Aruba wireless LAN controller.</p>
Controller IP Address	Indicates the IP address of the Aruba wireless LAN controller.
Switch Name	Indicates the Juniper Networks switch name to which Aruba devices are connected.
Switch IP Address	Indicates the IP address of the switch.
Switch Port Number	Indicates the number of the switch port where Aruba devices are connected.

Table 449: Fields in Aruba Device Inventory Table (continued)

Device Group	<p>Indicates the group name to which the device belongs.</p> <p>Aruba associates their devices in the form of groups, where a set of devices belonging to a single group shares same configuration. For example, a set of access points in a particular location can form a group and these groups are given unique names.</p> <p>Clicking on the group name opens the Groups > Monitor page in Airwave application displaying the details of the Aruba wireless device group.</p>
Switch Managed by ND	<p>Indicates whether the switch to which the Aruba device is connected to is managed by Network Director.</p>

RELATED DOCUMENTATION

[Understanding the Build Mode Tasks Pane | 189](#)

[Understanding the Monitor Mode Tasks Pane | 1277](#)

[Understanding the Fault Mode Tasks Pane | 1452](#)

[Understanding the Report Mode Tasks Pane | 1472](#)

[Understanding Aruba Airwave Integration with Network Director | 1548](#)

[Linking to the Aruba Airwave Application | 1553](#)

[Launching the Aruba Airwave Application | 1554](#)

[Network Director Documentation home page](#)

Linking to the Aruba Airwave Application

Sites with the Aruba Airwave application licenses can launch the Airwave application from within Network Director by supplying the Airwave application URL. After specifying the URL on the Wireless tab of Preferences, click **Launch Aruba Airwave** in Build mode, Monitor mode, Fault mode, or Report mode to load the respective pages of Airwave application.

To enable launching the Airwave application from within Network Director:

1. Type the URL required for launching the Airwave application.
2. Type the credentials that are used to log in to the Airwave application.

3. Specify the time in hours after which the Juniper Networks Network Director server synchronizes with the Aruba server to get the latest wireless devices inventory from the Aruba server.

The default synchronize interval is 24 hours. You can enter upto 720 hours as synchronize interval.

NOTE: If you specify the synchronize interval as 0 hours, the synchronization between the Juniper Networks server and Aruba server does not occur and you will need to manually synchronize these servers to retrieve the latest device inventory. To manually synchronize the servers click the **Resynchronize Now** button in the Aruba Wireless Device Inventory page.

4. Click **Check Connectivity** to check the network connection between the Aruba Aiwave application server and Network Director.
5. Click **OK**.

RELATED DOCUMENTATION

[Setting Up User and System Preferences | 108](#)

[Viewing Aruba Wireless Device Inventory in Network Director | 1550](#)

[Launching the Aruba Airwave Application | 1554](#)

[Understanding Aruba Airwave Integration with Network Director | 1548](#)

[Network Director Documentation home page](#)

Launching the Aruba Airwave Application

IN THIS SECTION

- [Launching the Aruba Airwave Application from Build Mode | 1555](#)
- [Launching the Aruba Airwave Application from Monitor Mode | 1555](#)
- [Launching the Aruba Airwave Application from Fault Mode | 1556](#)
- [Launching the Aruba Airwave Application from Reports Mode | 1556](#)
- [Launching Aruba Airwave for an Individual Aruba Wireless Device or Device Group | 1556](#)

To manage Aruba wireless devices that are connected to Network Director, use the Aruba Airwave application. You can launch the Aruba Airwave application from within Network Director. The launch URL is specified under System Preferences.

The Launch Aruba Airwave option under Device Management in the Build, Monitor, Fault, and Report modes opens the respective pages for Aruba Airwave application using which you can manage Aruba devices.

This section describes:

Launching the Aruba Airwave Application from Build Mode

From Build mode, you can launch the Aruba Airwave application home page, which displays monitoring status and configuration compliance pie charts, number of users on the network during a period of time, and other details for all the Aruba wireless devices connected to Juniper Networks switches:

1. Under Build mode, select **Device Management** and click **Launch Aruba Airwave**.

The Aruba Airwave application login page opens.

2. Type the credentials in the Aruba Airwave application login page.

The Home > Overview page of the Aruba Airwave application opens.

3. You can manage Aruba wireless devices by using this page.

Launching the Aruba Airwave Application from Monitor Mode

From Monitor mode, you can launch the Aruba Airwave application APs/Devices > List, which shows the assignment status and other details of all the Aruba wireless devices:

1. Under Monitor mode, select **Device Management** and click **Launch Aruba Airwave**.

The Aruba Airwave application login page opens.

2. Type the credentials in the Aruba Airwave application login page.

The APs/Devices > List page of the Aruba Airwave application opens, listing all the wireless devices connected to Juniper Networks switches.

3. You can monitor Aruba wireless devices by using this page.

Launching the Aruba Airwave Application from Fault Mode

From Fault mode, you can launch the Aruba Airwave application System > Alerts page, which shows alerts and alert details for all the Aruba wireless devices:

1. Under Fault mode, select **Device Management** and click **Launch Aruba Airwave**.

The Aruba Airwave application login page opens.

2. Type the credentials in the Aruba Airwave application login page.

The System > Alerts page of the Aruba Airwave application opens, listing all the impacted wireless devices connected to Juniper Networks switches.

3. You can view the alert summary and monitor alerts for the Aruba wireless devices.

Launching the Aruba Airwave Application from Reports Mode

From Reports mode, you can launch the Aruba Airwave application Reports > Generated to view and generate reports for the Aruba wireless devices:

1. Under Reports, select **Device Management** and click **Launch Aruba Airwave**.

The Aruba Airwave application login page opens.

2. Type the credentials in the Aruba Airwave application login page.

The Reports > Generated page opens in Aruba Airwave application.

3. You can view reports that have been run, view the most recent daily version of the reports, and if you have the administrative privileges, rerun the reports if needed.

Launching Aruba Airwave for an Individual Aruba Wireless Device or Device Group

Apart from launching the Aruba Airwave pages for all the Aruba wireless devices, Network Director also has the ability to launch the Airwave application pages for an individual Aruba wireless device (access points and controllers) or devices in a group connected to Juniper Networks switches from within Network Director. By default, respective monitoring pages of Aruba Airwave application opens for the selected device and device group.

You can navigate to device inventory table in the following modes and click on the device name or device group name in this table to launch the Aruba Airwave application:

- Build mode—Select **Device Management** and click **View Aruba Wireless Device Inventory**
- Monitor mode—Select **Device Management** and click **View Aruba Wireless Devices**

- Fault mode—Select **Device Management** and click **Impacted Aruba Wireless Devices**

To launch the Aruba Airwave application, which shows an individual Aruba access point details:

1. Click the hostname entry under column Hostname in device inventory table.

The Aruba Airwave application login page opens.

2. Type the credentials in the Aruba Airwave application login page.

The APs/Devices > Monitor page of the Aruba Airwave application opens, showing details of the selected Aruba access point.

3. You can manage the selected Aruba access point by using this page.

To launch the Aruba Airwave application, which shows an individual Aruba wireless LAN controller details:

1. Click the controller name entry in the column Controller Name in device inventory table.

The Aruba Airwave application login page opens.

2. Type the credentials in the Aruba Airwave application login page.

The APs/Devices > Monitor page of the Aruba Airwave application opens, showing details of the selected Aruba wireless LAN controller.

3. You can manage the selected Aruba wireless LAN controller by using this page.

To launch the Aruba Airwave application, which shows an individual Aruba wireless device group:

1. Click the group name entry in the column Device Group in device inventory table.

The Aruba Airwave application login page opens.

2. Type the credentials in the Aruba Airwave application login page.

The Groups > Monitor page of the Aruba Airwave application opens, showing details of the selected Aruba wireless device group.

3. You can manage the selected device group by using this page.

RELATED DOCUMENTATION

[Understanding the Build Mode Tasks Pane | 189](#)

[Understanding the Monitor Mode Tasks Pane | 1277](#)

[Understanding the Fault Mode Tasks Pane | 1452](#)

[Understanding the Report Mode Tasks Pane | 1472](#)

[Understanding Aruba Airwave Integration with Network Director | 1548](#)

[Viewing Aruba Wireless Device Inventory in Network Director | 1550](#)

[Linking to the Aruba Airwave Application | 1553](#)

[Network Director Documentation home page](#)

Juniper Networks Data Center Switching Management Pack for VMware vRealize Operations (vROps) User Guide

60

CHAPTER

Juniper Networks Data Center Switching Management Pack for vROps

Understanding Juniper Networks Data Center Switching Management Pack for vROps | **1561**

Adding and Configuring Juniper Networks Data Center Switching Management Pack for vROps | **1563**

Monitoring Juniper Networks Devices from vROps | **1569**

Managing Juniper Networks Data Center Infrastructure from vROps | **1582**

Performing Fault Management in vROps | **1586**

Understanding Juniper Networks Data Center Switching Management Pack for vROps

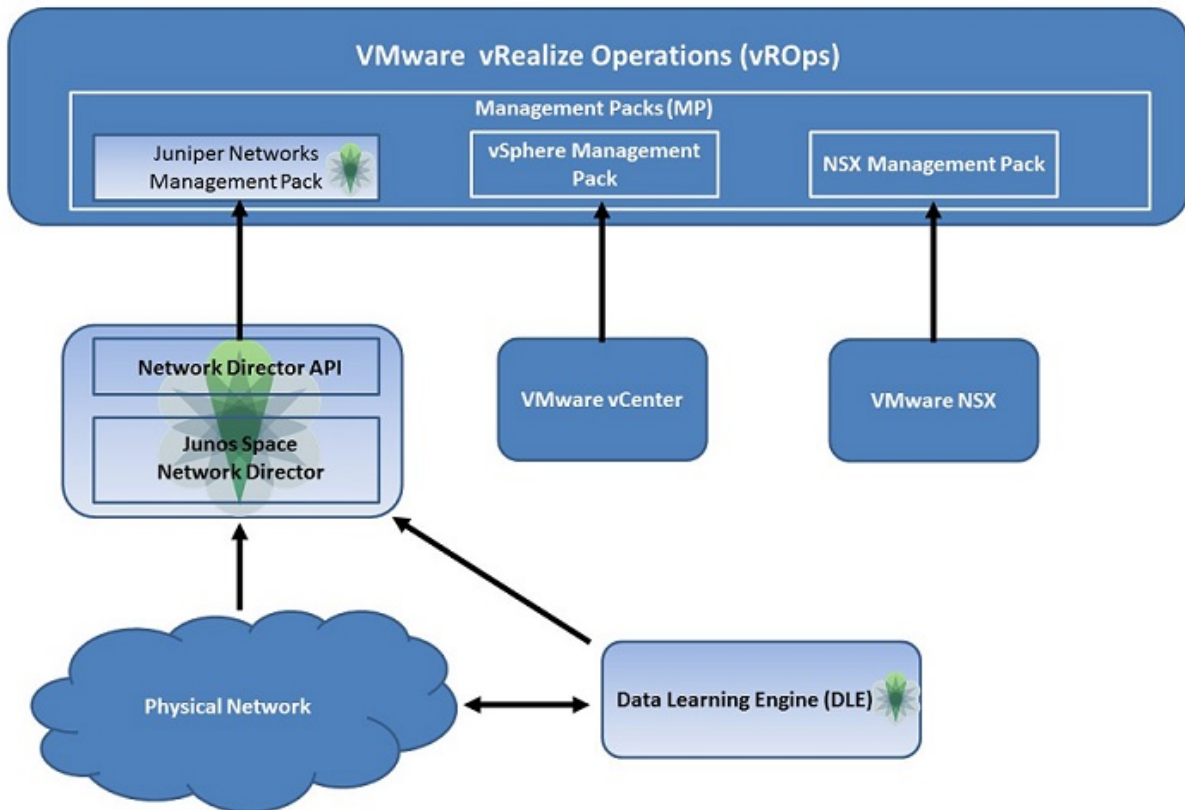
VMware vRealize Operations (vROps) is a component of VMware's vRealize suite of products. vROps provides an integrated, single-pane-of-glass(SPOG) view into the performance, capacity, and configuration management capabilities of VMware vSphere, physical, and hybrid cloud environments.

vROps management pack is part of a framework developed by VMware that enables vendors, such as Juniper Networks, to develop and integrate their plug-ins with vROps. This facilitates cloud administrators to monitor and manage the devices from the vendor along with VMware devices in a physical and cloud environment using vROps. Cloud administrators need not switch to the vendor's network management system to view device-specific information provided by the vendor.

Juniper Networks Data Center Switching Management Pack for vROps is a plug-in that you can install and integrate with vROps. After it is installed, the management pack obtains all the necessary monitoring data from Juniper Networks devices and displays the data in vROps.

[Figure 59](#) is a schematic display of how Juniper Networks Data Center Switching Management Pack for vROps interacts with the other components within Network Director to display data center details in vROps.

Figure 59: J Interaction of Juniper Networks Data Center Switching Management Pack for vROps with Network Director Components



Juniper Networks Data Center Switching Management Pack for vROps provides cloud administrators with:

- A network health dashboard.
- A correlated view of virtual and physical network components.

Depending of the size of the network and traffic volume, you can choose to install vROps as a standalone node or a multinode cluster. In a standalone node, data and administration are handled by the same node. In a multinode cluster, you have a master node that exclusively performs the administrative tasks. The additional nodes are configured to store data, collect data, provide load balancing, or provide high availability. In a multinode cluster, the master node must be online before you can add additional nodes. For instructions on installing vROps as a standalone node or multinode cluster, see the [vROps documentation](#).

NOTE: Juniper Networks Data Center Switching Management Pack for vROps monitors and troubleshoots only those fabric devices that are configured in the Datacenter View in Network Director.

Benefits of Juniper Networks Data Center Switching Management Pack for vROps

- Provides a holistic view of the data centers (that are managed by Network Director) through a single pane of glass in a vROps dashboard offering visibility into activity across the entire data center infrastructure.
- Enables administrators to identify the root cause of network fabric issues quickly and then perform deeper troubleshooting by means of a correlated health dashboard and easy navigation to the detailed views of the network provided by Junos Space Network Director.
- Provides effective and quick fault management by displaying the most volatile network fabrics based on alarms and the busiest network fabrics based on CPU and memory utilization.

RELATED DOCUMENTATION

[Monitoring Juniper Networks Devices from vROps | 1569](#)

[Managing Juniper Networks Data Center Infrastructure from vROps | 1582](#)

[Performing Fault Management in vROps | 1586](#)

Adding and Configuring Juniper Networks Data Center Switching Management Pack for vROps

IN THIS SECTION

- [Adding the Juniper Networks Data Center Switching Management Pack for vROps | 1564](#)
- [Specifying the Network Director Credentials in vROps | 1565](#)
- [Specifying the VMware vCenter Details in vROps | 1566](#)
- [Creating a Read-Only User in vROps | 1567](#)
- [Adding a VMware vCenter in Network Director | 1568](#)

Juniper Networks Data Center Switching Management Pack for vROps is a plug-in that you can add to VMware vRealize Operations (vROps) to have a single-pane of glass view into the performance, capacity,

and configuration management capabilities of the cloud data center that vROps manages. The data center can include Juniper Networks devices and VMware devices.

NOTE: You cannot uninstall or remove a management pack from any installation of vROps. However, you can upgrade a management pack to a later version.

This topic describes:

Adding the Juniper Networks Data Center Switching Management Pack for vROps

Before you begin, ensure that you have:

- vROps version 6.0, 6.0.1, 6.1, 6.2, 6.3, 6.4, or 6.5 running
- Downloaded and extracted the appropriate Juniper Networks Management Pack to a folder on your local system.

If you have vROps version 6.0, 6.0.1, or 6.1 installed, you must download and use Juniper Networks Management Pack version 1.1.

If you have vROps version 6.2, 6.3, 6.4, or 6.5 installed, you must download and use Juniper Networks Management Pack version 2.0.

You can download the management pack from the [Software Download](#) page. The solution pack has a PAK extension.

To add the Juniper Networks Management Pack:

1. Log in to vROps as a user with administrative privileges.
2. From the left navigation pane, select **Administration** and then select **Solutions**. The Solutions page opens.
3. In the Solutions tab of the Solutions page, click +.
4. Browse and select the management pack (PAK file) from the folder to which you extracted the contents of the Juniper Networks Management Pack.
5. Select the option to either install the PAK file even if it is already installed without overwriting the existing settings or reset the default settings and overwrite to a newer version.

6. Click **Upload**.
7. Click **Next** after the upload completes.
8. Read and accept the license agreement. Click **Next**.

The installation might take several minutes to complete. The progress of the installation is displayed.

9. When the installation is complete, click **Finish**.



On successful installation, vROps lists Juniper Networks Management Pack as a solution in the Solutions page.

Specifying the Network Director Credentials in vROps

After you add Juniper Networks Management Pack to vROps, you must perform a few more tasks to enable vROps to connect to Network Director and obtain all the necessary data and metrics.

Before you specify the credentials, make sure that you have added at least one vCenter using the Datacenter View in Network Director. For more information, see [Create a Data Center](#).

To perform the initial configuration:

1. Log in to vROps as a user with administrative privileges.
2. From the left navigation pane, select **Administration** and then select **Solutions**.
The Solutions page opens.
3. Select the Juniper Networks Management Pack for which you want to specify the Network Director credentials and click  on the toolbar.
The Manage Solution - Juniper Networks Management Pack page appears.
4. Select the Juniper adapter from the Adapter list.
5. Enter a display name and a description for the adapter instance in the Instance Settings section.
6. Enter the IP address of the Network Director server to which you want vROps to connect.
7. Click  to add credential details that vROps uses to log in to Network Director.

Specify the name by which you want to save the credentials. Enter the username and the password that vROps uses to connect to Network Director.

NOTE: The Network Director user that you specify must be part of the **Network Director – Admin, Super Administrator, or Network Director – Engineer** user role. You can modify user roles for Network Director users from the Junos Space user interface. For more details, see [Role-Based Access Control Overview](#).

8. Specify a username and password for a vROps user. This user must have at least read-only access to vROps.

This enables vROps to create relationships between VMware ESXi hosts and Juniper Networks data center devices, in vROps. Click **OK** to save the credentials.

For detailed steps to create a read-only user, see [“Creating a Read-Only User in vROps” on page 1567](#).

9. Click **Test Connection** to test the connection between vROps and Network Director. vROps tests and displays the connection status.
10. Click **Save Settings** to complete the initial configuration. vROps gathers data from Network Director and builds the network in vROps. This might take a while depending on the number of data centers, devices, hosts, and virtual machines that are managed by Network Director.


Specifying the VMware vCenter Details in vROps

If your data center uses VMware vCenters as the cloud infrastructure provider, then you must add those vCenters to vROps and Network Director to facilitate data collection from the virtual network. This topic describes the steps to add a vCenter in vROps. For detailed steps on adding the same vCenter server in Network Director, see [“Adding a VMware vCenter in Network Director” on page 1568](#).


To add a vCenter in vROps:

1. Log in to vROps as a user with administrative privileges.
2. From the left navigation pane, select **Administration** and then select **Solutions**.

The Solutions page opens.

3. Select the **VMware vSphere** for which you want to specify the VMware vCenter details and click  on the toolbar.


The Manage Solution - VMware vSphere page appears.

4. Select **vCenter Adapter** from the Adapter list.
5. Enter a display name and a description for the adapter instance in the Instance Settings section.
6. Enter the IP address of the vCenter server.
7. Click  to add credential details that vROps uses to log in to the vCenter.
Specify the name by which you want to save the credentials. Enter the username and the password that vROps uses to connect to the vCenter. This must be the user with administrative privileges on the vCenter server. Click **OK** to save the credentials.
8. Click **Test Connection** to test the connection between vROps and the vCenter server. vROps tests and displays the connection status.
9. Click **Save Settings** to complete the configuration.

Creating a Read-Only User in vROps

You must add a user with read-only access in vROps for the Juniper Networks Management Pack to obtain data from vROps. This user must then be added to vROps from the Solutions tab for the vCenter Adapter.

To create a read-only user in vROps:

1. Click **Administration** in the left navigation pane.
2. Click **Access Control** in the left navigation pane.
The Access Control page appears.
3. Click  to add a user account.
The Add User window opens.
4. Enter the user details and click **Next**.

5. Click the **Objects** tab and select **ReadOnly** from the Select Role drop-down list.
6. Select **Assign this role to the user**.
7. Click **Finish**.

You can add this user to vROps by following the steps given in [“Specifying the VMware vCenter Details in vROps” on page 1566](#).

Adding a VMware vCenter in Network Director

To add a VMware vCenter in Network Director:

1. Log in to the Network Director.
2. Select Build mode and select Datacenter Views, from the Views list, click **Setup Datacenter** from the Tasks pane.
The Setup Datacenter page opens.
3. Enter a name for the data center.
4. Click **Next** to specify details of the cloud infrastructure that the data center uses.
5. In the Cloud Infrastructure wizard page, click **Yes** to specify the cloud infrastructure for your data center.
6. Select **VMware vCenter** as the type of cloud infrastructure provider.
7. Enter the IP address or the hostname of the vCenter.
8. Specify the port that Network Director uses to connect to the server. The default port used to connect to a vCenter server is 443.

NOTE: You can modify the default and specify a port of your choice. If you do so, make sure to manually change the Junos Space firewall settings and apply those to the port you specify.

9. Specify the administrator username and password for the server you selected. The username and password must match the name and password configured on the server.

10. Click **Next** to specify the method that you want Network Director to use to build the data center network.

11. Select the Juniper Networks devices that you want to add to the data center from the Available Devices table and click >> to add it to the Selected Devices table.

If you want to remove a device from the Selected Devices table, select the device and click <<.

12. Click **Done** to save the data center details.

A message window opens, displaying the status of the cloud infrastructure discovery job name and job ID. Click **OK**.

You can view the status of the discovery job in the Cloud Infrastructure Discovery Jobs page.

Network Director tries to discover the servers that you specified in the Cloud Infrastructure wizard page. Network Director then adds the remaining network infrastructure based on the details you specified in the Network Infrastructure wizard page. Once these two steps are complete, Network Director lists that data center along with the devices that are part of the network infrastructure under My Datacenters in the View pane.

RELATED DOCUMENTATION

[Managing Juniper Networks Data Center Infrastructure from vROps | 1582](#)

[Performing Fault Management in vROps | 1586](#)

Monitoring Juniper Networks Devices from vROps

IN THIS SECTION

- [Using the Juniper Infrastructure Overview Dashboard | 1570](#)
- [Using the Juniper Network Fabric Monitoring Dashboard | 1575](#)
- [Using the Juniper Network Fabric Member Monitoring Dashboard | 1578](#)
- [Using the Juniper Top Network Fabrics Dashboard | 1580](#)
- [Using the Juniper Top Network Fabric Members Dashboard | 1581](#)

After you add and configure Juniper Networks Data Center Switching Management Pack for vROps, vROps discovers and builds the data center network managed by Network Director in vROps. vROps also adds five dashboards to the vROps Dashboards list. Each dashboard contains various dashboard widgets that assist administrators to monitor and manage multiple data center from vROps without having to switch to a different network management application.

vROps adds the following dashboards to the Dashboards list:

- Juniper Infrastructure Overview
- Juniper Network Fabric Monitoring
- Juniper Network Fabric Member Monitoring
- Juniper Top Network Fabrics
- Juniper Top Network Fabric Members

This topic describes:

Using the Juniper Infrastructure Overview Dashboard

IN THIS SECTION

- [View the Data Center Connectivity | 1571](#)
- [View the Top Alerts for your Data Center | 1573](#)
- [View Relationship Between Various Devices in the Data Center | 1574](#)


Enables you to view the topology and monitor the health of the data center network, by providing widgets. The Datacenter Overview widget displays all the data centers, underlying fabrics, hosts, and the connected virtual machines. If you click any data center to select it, the widget highlights the fabrics, hosts, and virtual machines that are part of that data center. You can view the top alerts, if any, for the selected data center in the Top Alerts widget. The Object Relationship widget displays the connectivity between the data center and the fabrics.

You can also filter the Datacenter Overview widget and the Object Relationship widget by two criteria—Health and Risk.

The Juniper Infrastructure Overview dashboard gives you an overview of the topology of the data center network and also enables you identify the top alerts for each data center network. The [Figure 60](#) shows the Juniper Infrastructure Overview dashboard comprising of three widgets—Datacenter Overview, Top

Alerts, and Object Relationship. These widgets display all the data centers, underlying fabrics, host devices, and virtual machines that are part of your data center network. Each of these entities are grouped under the respective categories.

Figure 60: Juniper Infrastructure Overview dashboard



You can perform the following tasks from the Juniper Infrastructure Overview dashboard:

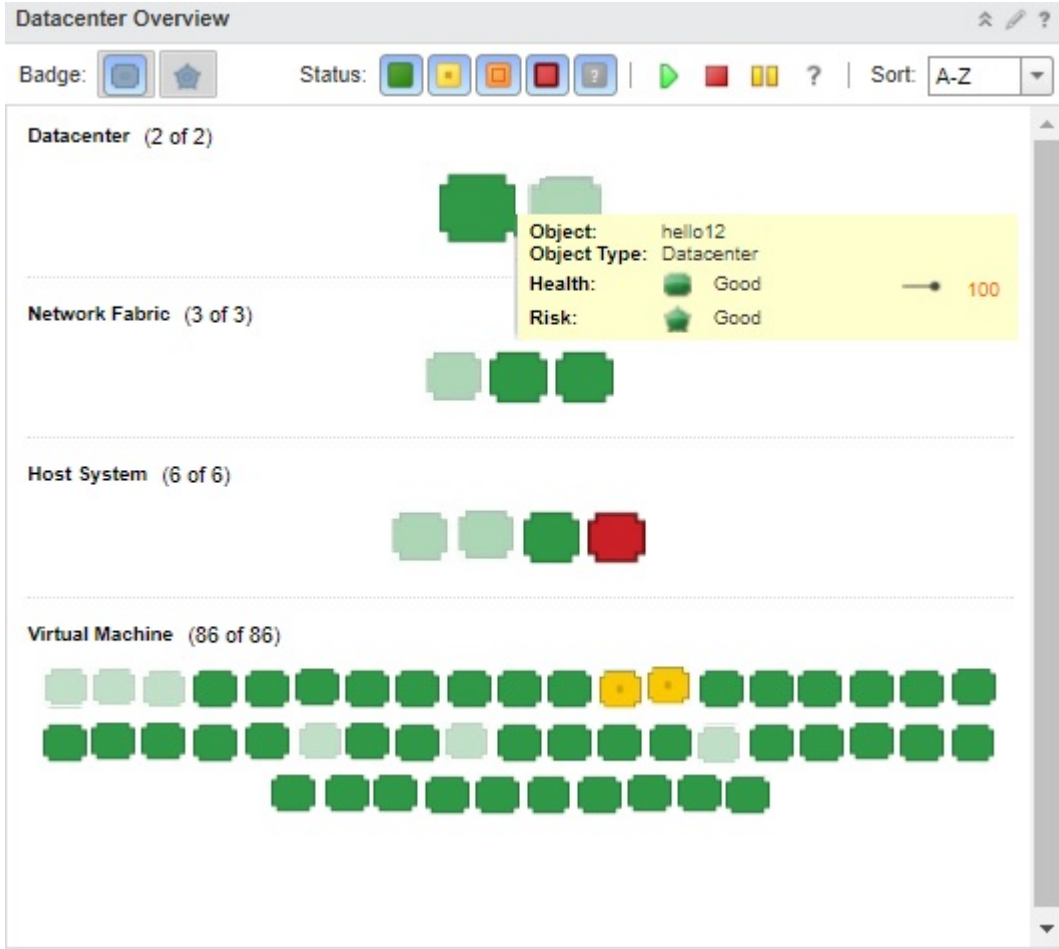
View the Data Center Connectivity

You can view the connectivity between the various devices that form your data center, by using the Datacenter Overview widget. You can do the following:

- Select a data center, network fabric, host device, or a virtual machine to display the connected devices that form that data center. For example, if you select a data center, vROps highlights the connected network fabrics, host devices, and virtual machines in the Datacenter Overview widget.

In the [Figure 61](#), you can see that the data center *hello12* is connected to two network fabrics which is connected to two host devices. The host devices host a set of virtual machines between them and these are highlighted in the Virtual Machines section.

Figure 61: Datacenter View





- View devices using two different modes—Health and Risk. Use the Health  and the Risk  buttons to view data center devices in the respective modes (see Figure 61). These modes display device icons of different colors based on the health and risk status of a device. Table 450 describes what each device icon color indicates in each of these modes.

Table 450: Descriptions for Device Icons Based on Their Colors





Device icon color	In Health Mode	In Risk Mode
	Indicates that the device works fine.	Indicates that there are no risks or threats identified as of now.

Table 450: Descriptions for Device Icons Based on Their Colors (*continued*)

Device icon color	In Health Mode	In Risk Mode
	Indicates one or more minor issues in the device that might not impact the functioning of the device—for example, memory contention.	Indicates that the device has one or more minor risks that might not impact the functioning of the device. These include high CPU utilization, high memory utilization, high port utilization, high port latency high, or high port packet drop rate.
	Indicates an issue in the device that might impact the normal functioning of the device—for example, redundant connectivity or a disk Input/Output write latency issue.	Indicates that the device has major risks that can impact the normal functioning of the device.
	Indicates a major issue in the device that needs to be fixed immediately—for example, the link between two devices is down.	Indicates major risks that need immediate attention.

View the Top Alerts for your Data Center

You can view the top alerts for your data center in the Top Alerts widget.

To view the alerts:

1. Select a data center or connected device in the Datacenter Overview widget.

vROps displays alerts for the data center, which include alerts for devices that are part of the data center, or the alerts for the selected device if you selected an individual device.

2. From the Top Alerts widget, you can:

- View the severity of the alert based on the color of the alert badge. [Table 450](#) describes what each icon color indicates in the Health mode and the Risk mode.
- Click the short description of the alert to open and view a summary of the alert and the corrective action that you must take to get rid of the alert.

vROps opens the Risk Issues page as shown in [Figure 62](#). This page displays the details about the alert and a table listing the alerts.

Figure 62: Risk Issues

Criticality	Alert Details	Triggered On	Created On	Updated On
Warning	View Details	central-ms_Soumen	8/5/17 3:04 AM	8/7/17 10:04 PM
Warning	View Details	regional-ms_Soumen	8/5/17 3:04 AM	8/5/17 3:04 AM

- Click **View Details** in the Alert Details column to view the alert details.
The Alert Details opens, which displays details about the alert.
- You can also click on the alert name in the Triggered On column to view the review the summary alert information for hosts and virtual machine.
- You can enter the alert name in the Quick filter (Name) field to display details of that only that alert in the table.
- Close the Risk Issues page to return to the Home page.

View Relationship Between Various Devices in the Data Center

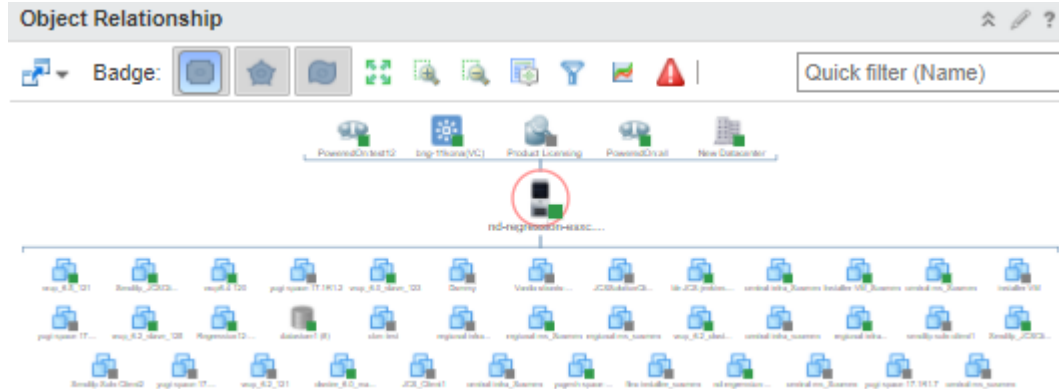
You can view the relationship between devices in the Object Relationship widget.

To view relationship between devices:





1. Select an object in the Datacenter Overview widget.

vROps displays the connection between the selected object and the connected devices. For example, if you select a host system in the Datacenter Overview widget, vROps displays the connected network fabric and the virtual machines that are part of the host system, in the Object Relationship widget as shown in [Figure 63](#).

Figure 63: Object Relationship Between Hosts and Virtual Machines



2. From the Object Relationship widget, you can:

- Toggle the Health  Risk , and Efficiency  badges to view devices based on the health, risk, or efficiency.
- Click  to view the list of alerts for the devices that are displayed.

Using the Juniper Network Fabric Monitoring Dashboard

IN THIS SECTION

- [View Data Center Fabric Details | 1576](#)
- [View CPU and Memory Utilization History and Forecast of a Fabric | 1578](#)

Enables you to monitor Juniper devices that are part of a data center, by providing dashboard widgets. You can view the list of fabrics that are part of the data center in the Network Fabrics widget. You can then select a fabric from the list to see the CPU utilization and memory utilization of the fabric in the CPU Utilization History and Forecast widget and the Memory Utilization History and Forecast widget respectively.

The Juniper Network Fabric Monitoring dashboard enables you to monitor Juniper Networks fabric devices such as Virtual Chassis, Junos Fusion data center fabric, QFabric systems, Layer 3 Fabrics, and Virtual Chassis Fabric that are part of a data center, using the following widgets:

- Network Fabrics widget

- CPU Utilization History and Forecast widget
- Memory Utilization History and Forecast widget

To open the Juniper Network Fabric Monitoring dashboard:

1. While in the vROps home page, select **Juniper Network Fabric Monitoring** from the **Dashboard List**.

The Juniper Network Fabric Monitoring dashboard opens.

You can use the Juniper Network Fabric Monitoring dashboard to:

View Data Center Fabric Details

The Network Fabrics widget displays the list of network fabrics that are part of all the data centers that are managed by Network Director. [Table 451](#) describes the fields that are displayed in the Network Fabrics widget.

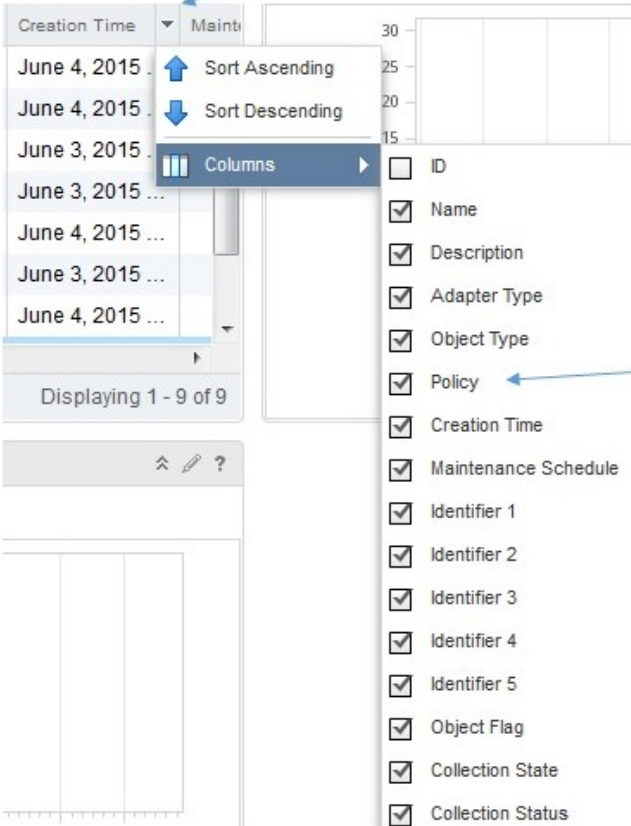
Table 451: Network Fabrics widget Field Descriptions

Field	Description
Name	Name of the device.
Description	Short description of the device.
Adapter Type	Type of adapter.
Object Type	Type of object.

By default, the table displays a few prefiltered fields. However, you can filter the table to display the fields that you want to, by following the steps shown in [Figure 64](#).

Figure 64: Filtering Fields in a Table

1. Click the DOWN arrow on any column header



2. Select Columns and from the list of available fields, clear the check box corresponding to the field that you do not want to appear in the table

Creation Time	Mainti				
June 4, 2015					
June 4, 2015					
June 3, 2015					
June 3, 2015					
June 4, 2015					
June 3, 2015					
June 4, 2015					

- ID
- Name
- Description
- Adapter Type
- Object Type
- Policy
- Creation Time
- Maintenance Schedule
- Identifier 1
- Identifier 2
- Identifier 3
- Identifier 4
- Identifier 5
- Object Flag
- Collection State
- Collection Status

View CPU and Memory Utilization History and Forecast of a Fabric

Select a fabric from the Network Fabrics table. vROps displays the CPU utilization and the memory utilization of the selected fabric in the CPU Utilization History and Forecast widget and the Memory Utilization History and Forecast widget respectively. You can use these widgets to view the historical, current, and the projected CPU and memory utilization for the selected fabric. The historic utilization data and the utilization forecast data are plotted in different colors as shown in the legend under the utilization graph.

Using the Juniper Network Fabric Member Monitoring Dashboard

IN THIS SECTION

- [View Fabric Member Details | 1579](#)
- [View CPU Utilization and Memory Utilization of a Fabric Member | 1580](#)

Enables you to monitor the fabric members of each Juniper device, that is part of a data center, by providing dashboard widgets. You can view the list of fabric members that are part of the device in the Network Fabric Members widget. You can then select a fabric member from the list to view the CPU utilization and memory utilization of the member in the CPU Utilization History and Forecast widget and the Memory Utilization History and Forecast widget respectively.

The Juniper Network Fabric Member Monitoring dashboard enables you to monitor the members of Juniper Networks fabric devices such as Virtual Chassis , Virtual Chassis Fabric that are part of a data center, using the following widgets:

- Network Fabric Members widget
- CPU Utilization History and Forecast widget
- Memory Utilization History and Forecast widget

To open the Juniper Network Fabric Member Monitoring dashboard:

1. While in the vROps home page, select **Juniper Network Fabric Member Monitoring** from the **Dashboard List**.

The Juniper Network Fabric Member Monitoring dashboard opens.

You can use the Juniper Network Fabric Member Monitoring dashboard to:

View Fabric Member Details

The Network Fabric Members widget displays the list of fabric members that are part of the network fabrics managed by Network Director. [Table 452](#) describes the fields that are displayed in the Network Fabric Members widget.

Table 452: Network Fabric Members widget Field Descriptions

Field	Description
Name	Name of the member device.
IP Address	IP address of the member device.
Member Type	Indicates whether the member device belongs to a QFabric, Virtual Chassis Fabric, or a Layer 3 Fabric.
Device Name	Name of the data center fabric to which the member belong.
Collection State	Indicates the status of data collection from the member.

View CPU Utilization and Memory Utilization of a Fabric Member

Select a fabric member from the Network Fabric Members table. vROps displays the CPU utilization and the memory utilization of the selected member in the CPU Utilization History and Forecast widget and the Memory Utilization History and Forecast widget respectively. You can use these widgets to view the historical, current, and the projected CPU and memory utilization for the selected fabric member. The historic utilization data and the utilization forecast data are plotted in different colors as shown in the legend under the utilization graph.

Using the Juniper Top Network Fabrics Dashboard


Enables you to view and identify potential performance issues in your data center. The dashboard provides four widgets—The Top Network Fabrics by CPU Utilization widget, the Top Network Fabrics by Memory Utilization, the Top Noisiest Network Fabrics based on Alerts widget, and the Top Volatile Network Fabrics based on Metrics widget. You can click a row to view more details about that fabric and possible causes for the vulnerability of the fabric, if any.

You can view the top network fabrics based on CPU utilization, memory utilization, alerts, and volatility metrics by using the Juniper Top Network Fabrics dashboard. This dashboard enables you to assess the health of your data center network by identifying the fabrics that have a high CPU or memory utilization and large number of alerts.

To open and view the Juniper Top Network Fabrics dashboard:

1. While in the vROps home page, select **Juniper Top Network Fabrics** from the **Dashboard List**.

The Juniper Top Network Fabrics dashboard opens.

2. You can view the following details by using the widgets in the Juniper Top Network Fabrics dashboard:
 - Top Network Fabrics by CPU Utilization widget—Displays the top fabrics based on CPU utilization.
 - Top Network Fabrics by Memory Utilization widget—Displays the top fabrics based on memory utilization.
 - Top Noisiest Network Fabrics based on Alerts widget—Displays the top fabrics based on the number of alerts that are generated from the fabric and the connected devices.
 - Top Volatile Network Fabrics based on Metrics widget—Displays the top fabrics based on anomalies in the CPU and memory utilization over a period of time.
3. Select a fabric and click  to view more details about the selected fabric.

Using the Juniper Top Network Fabric Members Dashboard


NOTE: This dashboard displays member details only for Virtual Chassis and Virtual Chassis Fabric devices.

Enables you to view the performance of network fabric members, in terms of CPU utilization, memory utilization, alerts, and metrics. This dashboard provides four widgets—Top Network Fabric Members by CPU Utilization widget, Top Network Fabric Members by Memory Utilization widget, Top Noisiest Network Fabric Members based on Alerts widget, and Top Volatile Network Fabric Members based on Metrics widget.

You can view the top network fabric members based on CPU utilization, memory utilization, alerts, and volatility metrics by using the Juniper Top Network Fabrics dashboard. This dashboard enables you to assess the health of your data center devices by identifying the members that have a high CPU or memory utilization and alerts.

NOTE: Juniper Networks Data Center Switching Management Pack for vROps supports only fabric devices—Virtual Chassis, Virtual Chassis Fabric, QFabric, or Layer 3 Fabric, Junos Fusion data center fabric—that are part of a datacenter in Network Director.

To open and view the Juniper Top Network Fabric Members dashboard

1. While in the vROps home page, select **Juniper Top Network Fabric Members** from the **Dashboard List**.
The Juniper Top Network Fabric Members dashboard opens.
2. You can view the following details by using the widgets in the Juniper Top Network Fabric Members dashboard:
 - Top Network Fabric Members by CPU Utilization widget—Displays the top fabric members based on CPU utilization.
 - Top Network Fabric Members by Memory Utilization widget—Displays the top fabric members based on memory utilization.
 - Top Noisiest Network Fabric Members based on Alerts widget—Displays the top fabric members based on the number of alerts that are generated from the fabric and the connected devices.
 - Top Volatile Network Fabric Members based on Metrics widget—Displays the top fabric members based on anomalies in the CPU and memory utilization over a period of time.
3. Select a fabric member and click  to view more details about the selected member.

RELATED DOCUMENTATION

| [Understanding Juniper Networks Data Center Switching Management Pack for vROps](#) | 1561

Managing Juniper Networks Data Center Infrastructure from vROps

IN THIS SECTION


- [Open the Juniper Networks Data Center Infrastructure View](#) | 1582
- [View Data Center Details](#) | 1583
- [Open Network Director from vROps](#) | 1584

You can view the Juniper Networks data center infrastructure by using the various dashboards that are installed when you add the Juniper Networks Management Pack to vROps. In addition to this, you can use the Juniper Infrastructure view to display additional details for the data center and devices.

This topic describes how to:

Open the Juniper Networks Data Center Infrastructure View

To open the Juniper Networks Data Center Infrastructure View:

1. Do one of the following:
 - Click  in the left navigation pane tool bar.
 - Select **Environment** from the list in the left navigation pane tool bar
2. From the Environment Overview pane, select **Juniper Infrastructure**.

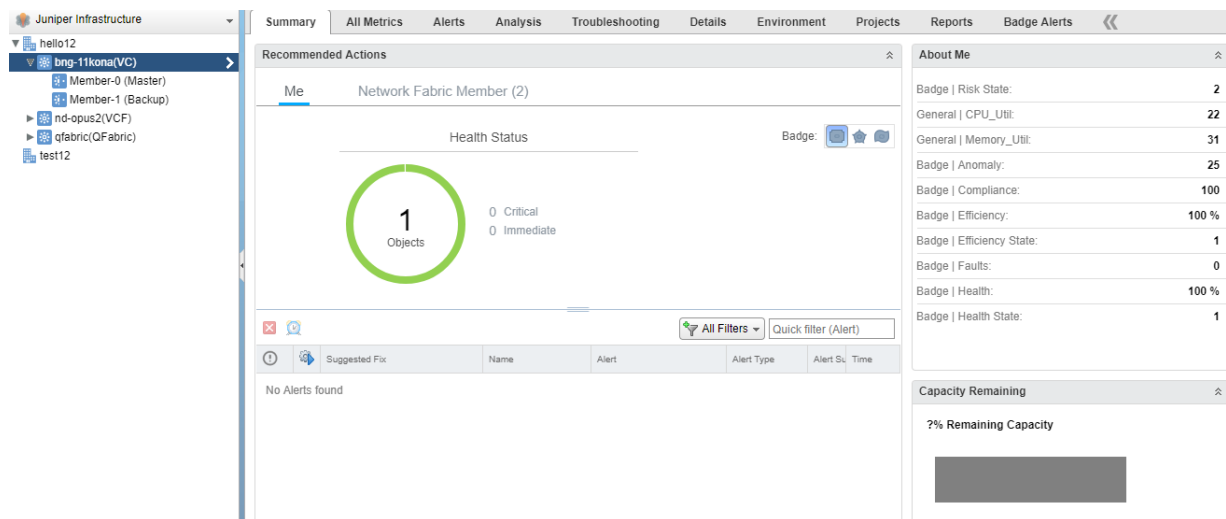
The Juniper Networks Data Center Infrastructure View page opens. The left pane displays the list of data centers that are part of Juniper infrastructure. The right pane displays details about the selected data center. Navigate the tabs in the right pane to get more details about each data center.

View Data Center Details

To view the data center details in the Juniper Networks Data Center Infrastructure View page:

1. Select a network fabric from the left navigation pane to view details about the fabric.
2. Use the following tabs in the right pane of the Juniper Networks Data Center Infrastructure View page to view more details about the selected data center:
 - Summary—Displays information summary under three categories—Health, Risk, and Efficiency. Each category displays a badge and top alerts for the fabric and descendants, for example the [Figure 65](#) shows the health status of the fabric device *bng-11kona(VC)*. The badge color indicates the status of the fabric against each category.

Figure 65: Summary - Health Status



- All Metrics—Displays the object relationship between the selected fabric and other devices. Double-click the device that you are troubleshooting, to view it in the context of parent and child objects.
- Alerts—Displays all the alerts for the selected fabric and descendants. Click the alert description for displaying details about an alert.
- Analysis—Displays anomalies in the fabric, if any.
- Troubleshooting—Enables you to troubleshoot issues in the fabric and descendants by using the data presented under four tabs:
 - Symptoms—Displays the symptoms that you can use to analyze and troubleshoot issues.
 - Timeline—Displays a customizable timeline. You can use the timeline to identify common trends over time.

- Events—Displays changes that occurred on the selected fabric or descendants because of user actions, system actions, triggered symptoms, or generated alerts.
- Details—Displays details about the network fabric and the fabric members, such as the current active alerts, CPU and memory utilization of the fabric device and fabric members, and a list of symptoms. To view any of these details, click the corresponding item in the list. vROps displays the details for the item.
- Environment—Displays the topology of the network fabric along with the data centers and the fabric members.
- Projects—Displays list of all the projects generated for the selected object, group, or application.
- Badge Alerts—Displays the Health, Risk, and Efficiency details.
- Reports—You can generate reports using an existing report template or create a new report template and generate reports using the new template.

NOTE: See the vROps online Help to understand more about the Reports functionality.

Open Network Director from vROps

You can directly launch Network Director from vROps to view the port utilization and alarm details of a member device or to view the connectivity of member devices.

To open Network Director from vROps:

1. Select a network fabric from the left navigation pane.
2. Do one of the following:

- To view the port utilization of a member of Virtual Chassis, Virtual Chassis Fabric, QFabric devices, Layer 3 Fabric, select **View Port Utilization** from the **Actions** menu.

The Network Director user interface opens in a new window displaying the Utilization for Device page. This page displays the port utilization trend for the selected devices or member. For more details, see [Utilization for Device page](#) documentation.

NOTE: This functionality is available for Layer 3 Fabric only at the member level and not at the fabric level.

- To view alarm details for the selected fabric, select **View Alarms** from the **Actions** menu.

The Network Director user interface opens in a new window. After you log in, Network Director opens the Fault mode displaying the current active alarms for the selected fabric in various alarm widgets.

NOTE: This functionality is available for Layer 3 Fabric only at the member level and not at the fabric level.

- To view connectivity between the fabric and other devices such as switches, hosts, and virtual machines, select **View Device Connectivity** from the **Actions** menu.

The Network Director user interface opens in a new window. After you log in, Network Director opens the Build mode displaying the Device Connectivity page for the selected fabric.

NOTE: The View Device Connectivity task is not available for QFabric devices.

- To view connectivity between the members of the selected fabric, select **View Fabric Internal Connectivity** from the **Actions** menu.

Network Director user interface opens in a new window. After you log in, Network Director opens the Build mode displaying the Connectivity view for the type of fabric you selected.

NOTE: The View Fabric Internal Connectivity task is available only for Virtual Chassis, Virtual Chassis Fabric devices, and QFabric devices.

RELATED DOCUMENTATION

[Understanding Juniper Networks Data Center Switching Management Pack for vROps | 1561](#)

[Adding and Configuring Juniper Networks Data Center Switching Management Pack for vROps | 1563](#)

[Managing Juniper Networks Data Center Infrastructure from vROps | 1582](#)

Performing Fault Management in vROps

IN THIS SECTION

- [Configuring Thresholds in vROps | 1587](#)
- [Modifying the Polling Interval in vROps | 1588](#)

- Juniper Networks Data Center Switching Management Pack for vROps triggers the following alarms in each category—Risk and Health:
 - Risk category
 - CPU Usage High Alert
 - Memory Usage High Alert
 - Port High Latency Alert
 - Port High Utilization alarm
 - Port High Packet Drop alarm
 - Health category
 - Port Link Down alarm
 - Device Connectivity Down alarm

NOTE: Except *CPU Usage High Alert* and *Memory Usage High Alert*, all other alarms are managed by Network Director. If alarms managed by Network Director are cleared from the Network Director user interface, they will also be cleared from vROps.

When any of the alarm thresholds are crossed, the system raises an alarm. All alarms are displayed in the Alerts page in vROps. While in the vROps home page, click **Alerts** in the left navigation pane to open the Alerts page. From this page, you can also view the timeline and object relationship for each alarm.


This topic describes:

Configuring Thresholds in vROps

When you install the Juniper Networks plugin for vROps, the thresholds for memory utilization and CPU utilization are already set. However, if you want to modify this value for your data center network, you can do so from the Metric / Supermetric Symptom Definition page in vROps.

To change the threshold values in vROps:

1. Do one of the following:

- Click  in the left pane toolbar.
- Click **Content** in the left navigation pane.


vROps opens the Content pane.

2. Click **Symptom Definition** in the Content pane.

3. If not already selected, click **Metric / Supermetric Symptom Definition** to select it. vROps opens the Metric / Supermetric Symptom Definition page in the right pane.

4. Select **Adapter Type** from the **All Filters** box and enter *Juniper* in the search box. Click **OK** to search for symptoms specific to Juniper Networks adapters.

Alternatively, scroll and locate the symptoms specific to Juniper Networks adapters—CPU Usage High Symptom, Device Connectivity Down Symptom, and Memory Usage High Symptom.



5. To change a threshold value, select a symptom and click .

6. Modify the threshold value as per your network requirements and click **Save**.

Modifying the Polling Interval in vROps

Polling interval is the time after which vROps contacts Network Director to obtain monitoring data for the data center devices. You might need to modify this value based on your network size and traffic density.

To modify the polling interval in vROps:

1. Open the **Solutions** page and note the name of the adapter instance that you gave while adding the Juniper Networks Management Pack in vROps.
2. Do one of the following:
 - Click  in the left pane toolbar.
 - Click **Administration** in the left navigation pane.
3. Click **Environment Overview** in the left navigation pane.
4. In the **Filter** box in the right pane, enter the name of the adapter instance that you noted in step 1. Select the Name from the drop-down list in the **Filter** box and press **Enter**.
vROps filters the list and displays the adapter instance with the name that you specified.
5. Select the adapter instance row and click  .
The Edit Object window opens.
6. Modify the value in the **Collection Interval** field to the required polling interval. This value must be specified in minutes.
7. Click **OK** to save the changes.

RELATED DOCUMENTATION

[Adding and Configuring Juniper Networks Data Center Switching Management Pack for vROps | 1563](#)

[Monitoring Juniper Networks Devices from vROps | 1569](#)

[Understanding Juniper Networks Data Center Switching Management Pack for vROps | 1561](#)