Dell PowerConnect 8024, 8024F, 8132, 8132F, 8164, and 8164F Switch

User's Configuration Guide

Regulatory Models: PC8024, PC8024F, PC8132, PC8132F, PC8164, PC8164F



Notes and Cautions

NOTE: A NOTE indicates important information that helps you make better use of your computer.

CAUTION: A CAUTION indicates potential damage to hardware or loss of data if instructions are not followed.

Information in this publication is subject to change without notice.
© 2012 Dell Inc. All rights reserved.

Reproduction of these materials in any manner whatsoever without the written permission of Dell Inc. is strictly forbidden.

Trademarks used in this text: Dell™, the DELL logo, EqualLogic™, PowerConnect™, and OpenManage™ are trademarks of Dell Inc. Microsoft®, Windows®, Windows Server®, MS-DOS®, and Windows Vista® are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. sFlow® is a registered trademark of InMon Corporation. Cisco® is a registered trademark of Cisco Systems. Mozilla® and Firefox® are registered trademarks of the Mozilla Foundation.

Other trademarks and trade names may be used in this publication to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

Regulatory Models: PC8024, PC8024F, PC8132, PC8132F, PC8164, PC8164F

June 2012 Rev. A03

Contents

1 Introduction 49

About This Document **49**

Audience **50**

Document Conventions **50**

Additional Documentation **51**

2 Switch Features 53

System Management Features **54**

Multiple Management Options 54

System Time Management 54

Log Messages 54

Integrated DHCP Server 55

Management of Basic Network Information 55

IPv6 Management Features 55

Dual Software Images 55

File Management 56

Switch Database Management Templates 56

Automatic Installation of Firmware and Configuration 56

sFlow 57

SNMP Alarms and Trap Logs 57

CDP Interoperability through ISDP 57

Remote Monitoring (RMON) 57

Stacking Features **58**

High Port Count 58

Single IP Management 58
Automatic Firmware Update for New Stack M embers 58
Master Failover with Transparent Transition 59
Nonstop Forwarding on the Stack 59
Hot Add/Delete and Firmware Synchronization 59

Security Features **59**
Configurable Access and Authentication Profiles 59
Password-Protected Management Access 60
Strong Password Enforcement 60
TACACS+ Client 60
RADIUS Support 60
SSH/SSL. 61
Inbound Telnet Control 61
Denial of Service 61
Port Protection 61
Captive Portal 62
Dot1x Authentication (IEEE 802.1X) 62
MAC-Based 802.1X Authentication 62
Dot1x Monitor Mode 63
MAC-Based Port Security 63
Access Control Lists (ACL) 63
Time-Based ACLs 64
IP Source Guard (IPSG). 64
DHCP Snooping 64
Dynamic ARP Inspection 64
Protected Ports (Private VLAN Edge). 65

Switching Features **65**
Flow Control Support (IEEE 802.3x) 65
Head of Line Blocking Prevention 65
Alternate Store and Forward (ASF) 65

Jumbo Frames Support 66
Auto-MDI/MDIX Support 66
VLAN-Aware MAC-based Switching 66
Back Pressure Support 66
Auto Negotiation 67
Broadcast Storm Control 67
Port Mirroring 67
Static and Dynamic MAC Address Tables 68
Link Layer Discovery Protocol (LLDP) 68
Link Layer Discovery Protocol (LLDP) for Media Endpoint Devices 68
Priority-based Flow Control (PFC) 68
Data Center Bridging Exchange (DBCx) Protocol 69
Enhanced Transmission Selection 69
Fibre Channel over Ethernet (FCoE) Initialization Protocol Snooping 69
Cisco Protocol Filtering 70
DHCP Layer 2 Relay 70

Virtual Local Area Network Supported Features 70

VLAN Support 70
Port-Based VLANs 70
IP Subnet-based VLAN 71
MAC-based VLAN 71
IEEE 802.1v Protocol-Based VLANs 71
GARP and GVRP Support 71
Voice VLAN 71
Guest VLAN 72
Double VLANs 72

Spanning Tree Protocol Features 73

Spanning Tree Protocol (STP) 73
Spanning Tree Port Settings 73
Rapid Spanning Tree 73

Multiple Spanning Tree 73
Bridge Protocol Data Unit (BPDU) Guard. 74
BPDU Filtering 74

Link Aggregation Features. **74**
Link Aggregation 74
Link Aggregate Control Protocol (LACP) 74

Routing Features **76**
Address Resolution Protocol (ARP) Table Management 76
VLAN Routing 76
IP Configuration. 76
Open Shortest Path First (OSPF) 76
BOOTP/DHCP Relay Agent 77
IP Helper and UDP Relay 77
Routing Information Protocol 77
Router Discovery 77
Routing Table 77
Virtual Router Redundancy Protocol (VRRP) 78
Tunnel and Loopback Interfaces 78

IPv6 Routing Features **78**
IPv6 Configuration 78
IPv6 Routes 79
OSPFv3 79
DHCPv6 79

Quality of Service (QoS) Features **80**
Differentiated Services (DiffServ) 80
Class Of Service (CoS) 80
Auto Voice over IP (VoIP) 80
Internet Small Computer System Interface (iSCSI) Optimization. 81

Layer 2 Multicast Features **81**
MAC Multicast Support. 81
IGMP Snooping 81
IGMP Snooping Querier 82
MLD Snooping 82
Multicast VLAN Registration 82

Layer 3 Multicast Features **83**
Distance Vector Multicast Routing Protocol 83
Internet Group Management Protocol 83
IGMP Proxy 83
Protocol Independent Multicast—Dense Mode 83
Protocol Independent Multicast—Sparse Mode 84
Protocol Independent Multicast—Source Specific Multicast 84
Protocol Independent Multicast IPv6 Support . . . 84
MLD/MLDv2 (RFC2710/RFC3810) 84

3 Hardware Overview 85

PowerConnect 8000-series and 8100-series Front Panel **85**
PowerConnect 8024 Front Panel 85
PowerConnect 8024F Front Panel 86
PowerConnect 8132 Front Panel 87
PowerConnect 8132F Front Panel 87
PowerConnect 8164 Front Panel 88
PowerConnect 8164F Front Panel 89
Hot-Pluggable Interface Modules 90
USB Port (Power Connect 8100-series switches only) 92
Port and System LEDs 92

PowerConnect 8000-series and 8100-series Back Panel **92**
Console Port 93
Out-of-Band Management Port. 93
Power Supplies 93
Ventilation System 94

LED Definitions **95**
Port LEDs 95
System LEDs 96

Switch Addresses **97**

4 Using Dell OpenManage Switch Administrator 101

About Dell OpenManage Switch Administrator **101**

Starting the Application **102**

Understanding the Interface **103**

Defining Fields **105**

Understanding the Device View **106**
Using the Device View Port Features. 106

5 Using the Command-Line Interface 107

Accessing the Switch Through the CLI **107**
Console Connection. 107
Telnet Connection. 108

Understanding Command Modes **109**

Entering CLI Commands **111**

Using the Question Mark to Get Help 111

Using Command Completion 112

Entering Abbreviated Commands 112

Negating Commands 112

Understanding Error Messages 113

Recalling Commands from the History Buffer . . . 113

6 Default Settings 115

7 Setting the IP Address and Other Basic Network Information 119

IP Address and Network Information Overview **119**

What Is the Basic Network Information? 119

Why Is Basic Network Information Needed? . . . 120

How Is Basic Network Information Configured? 121

What Is Out-of-Band Management and In-Band Management? 121

Default Network Information **123**

Configuring Basic Network Information (Web) **124**

Out-of-Band Interface 124

IP Interface Configuration (Default VLAN IP Address) 125

Route Entry Configuration (Switch Default Gateway) 126

Domain Name Server 128

Default Domain Name 129

Host Name Mapping 130

Dynamic Host Name Mapping 131

Configuring Basic Network Information (CLI) **132**

Enabling the DHCP Client on the OOB Port 132

Enabling the DHCP Client on the Default VLAN 132

Managing DHCP Leases 133

Configuring Static Network Information on the OOB Port 134

Configuring Static Network Information on the Default VLAN 134

Configuring and Viewing Additional Network Information 135

Basic Network Information Configuration Example **136**

8 Managing QSFP Ports 139

9 Managing a Switch Stack 141

Stacking Overview **141**

Creating a PowerConnect 8000/8100 Series Stack 142

How is the Stack Master Selected? 144

Adding a Switch to the Stack 145

Removing a Switch from the Stack 146

How is the Firmware Updated on the Stack? . . . 147

What is Stacking Standby? 147

What is Nonstop Forwarding? 147

Switch Stack MAC Addressing and Stack Design Considerations 150

NSF Network Design Considerations 151

Why is Stacking Needed? 151

Default Stacking Values **152**

Managing and Monitoring the Stack (Web) **153**

Unit Configuration 153

Stack Summary 155

Stack Firmware Synchronization 156

Supported Switches 157

Stack Port Summary 158

Stack Port Counters 159

Stack Port Diagnostics 159

NSF Summary 160

Checkpoint Statistics 161

Managing the Stack (CLI) **162**

Configuring Stack Member, Stack Port, and NSF Settings 162

Viewing and Clearing Stacking and NSF Information 164

Stacking and NSF Usage Scenarios **164**

Basic Failover 166

Preconfiguring a Stack Member 168

NSF in the Data Center 170

NSF and VoIP 171

NSF and DHCP Snooping 172

NSF and the Storage Access Network 173

NSF and Routed Access 175

10 Configuring Authentication, Authorization, and Accounting 177

AAA Overview **177**

Methods 178

Access Lines 179

Authentication **179**

Authorization **180**

Exec Authorization Capabilities. 181

Accounting **182**

Authentication Examples **183**

Local Authentication Example 183

TACACS+ Authentication Example 184

RADIUS Authentication Example 186

Authorization Examples **187**

Local Authorization Example—Direct Login to Privileged EXEC Mode 187

TACACS+ Authorization Example—Direct Login to Privileged EXEC Mode 187

TACACS+ Authorization Example—Administrative Profiles 188

TACACS+ Authorization Example—Custom Administrative Profile 189

TACACS+ Authorization Example—Per-command Authorization 190

RADIUS Authorization Example—Direct Login to Privileged EXEC Mode 191

RADIUS Authorization Example—Administrative Profiles 191

Using RADIUS Servers to Control Management Access **192**

How Does RADIUS Control Management Access? 192

Which RADIUS Attributes Does the Switch Support?. 194

How Are RADIUS Attributes Processed on the Switch? 196

Using TACACS+ Servers to Control Management Access **197**

Which TACACS+ Attributes Does the Switch Support? 198

Default Configurations **199**

Method Lists 199

Access Lines (AAA) 199

Access Lines (Non-AAA) 200

Administrative Profiles 200

11 Monitoring and Logging System Information 203

System Monitoring Overview **203**

What System Information Is Monitored? 203

Why Is System Information Needed? 204

Where Are Log Messages Sent? 204

What Are the Severity Levels? 205

What Are the System Startup and Operation Logs? 205

What Is the Log Message Format? 206

What Factors Should Be Considered When Configuring Logging? 207

Default Log Settings **207**

Monitoring System Information and Configuring Logging (Web) **208**

Device Information 208

System Health 209

System Resources 210

Integrated Cable Test for Copper Cables 211

Optical Transceiver Diagnostics 212

Log Global Settings 214

RAM Log 215
Log File 216
Remote Log Server 216
Email Alert Global Configuration 219
Email Alert Mail Server Configuration 220
Email Alert Subject Configuration 222
Email Alert To Address Configuration. 223
Email Alert Statistics 224

Monitoring System Information and Configuring Logging (CLI) **225**
Viewing System Information 225
Running Cable Diagnostics 225
Configuring Local Logging 226
Configuring Remote Logging 228
Configuring Mail Server Settings 229
Configuring Email Alerts for Log Messages . . . 230

Logging Configuration Examples **232**
Configuring Local and Remote Logging 232
Configuring Email Alerting 233

12 Managing General System Settings . . . 237

System Settings Overview **237**
Why Does System Information Need to Be Configured? 238
What Are SDM Templates? 239
Why is the System Time Needed? 240
How Does SNTP Work? 240

Default General System Information **240**

Configuring General System Settings (Web) **241**
System Information 241

CLI Banner 244
SDM Template Preference 245
Clock 246
SNTP Global Settings 247
SNTP Authentication 248
SNTP Server 250
Summer Time Configuration 254
Time Zone Configuration 255
Slot Summary 256
Supported Cards 257

Configuring System Settings (CLI) **258**
Configuring System Information 258
Configuring the Banner 259
Managing the SDM Template 260
Configuring SNTP Authentication and an SNTP Server 260
Setting the System Time and Date Manually 262
Viewing Slot Information 263

General System Settings Configuration Examples . . . **264**
Configuring System and Banner Information . . . 264
Configuring SNTP 267
Configuring the Time Manually 269

13 Configuring SNMP 271

SNMP Overview **271**
What Is SNMP? 271
What Are SNMP Traps? 272
Why Is SNMP Needed? 273

Default SNMP Values **273**

Configuring SNMP (Web) **275**

SNMP Global Parameters 275

SNMP View Settings 276

Access Control Group 278

SNMPv3 User Security Model (USM) 280

Communities 283

Notification Filter 285

Notification Recipients 286

Trap Flags 289

OSPFv2 Trap Flags 290

OSPFv3 Trap Flags 291

Trap Log 292

Configuring SNMP (CLI) **293**

Configuring the SNMPv3 Engine ID. 293

Configuring SNMP Views, Groups, and Users 294

Configuring Communities 297

Configuring SNMP Notifications (Traps and Informs). 299

SNMP Configuration Examples **302**

Configuring SNMPv1 and SNMPv2. 302

Configuring SNMPv3 303

14 Managing Images and Files. 307

Image and File Management Overview **307**

What Files Can Be Managed? 307

Why Is File Management Needed?. 309

What Methods Are Supported for File Management?. 311

What Factors Should Be Considered When Managing Files? 311
How Is the Running Configuration Saved? 313

Managing Images and Files (Web) 314
File System 314
Active Images 315
USB Flash Drive. 316
File Download. 317
File Upload 319
Copy Files 321

Managing Images and Files (CLI) 322
Downloading and Activating a New Image (TFTP) 322
Managing Files in Internal Flash 323
Managing Files on a USB Flash Device (PowerConnect 8100-series switches only) 324
Uploading a Configuration File (SCP). 324
Managing Configuration Scripts (SFTP) 325

File and Image Management Configuration Examples. 326
Upgrading the Firmware 326
Managing Configuration Scripts 329
Managing Files by Using the USB Flash Drive (PowerConnect 8100-series switches only) 331

15 Automatically Updating the Image and Configuration 333

Auto Configuration Overview 333
What Is USB Auto Configuration? 334
What Files Does USB Auto Configuration Use?. 334

What Is the DHCP Auto Configuration Process? 335
Monitoring and Completing the DHCP Auto Configuration Process 339
What Are the Dependencies for DHCP Auto Configuration?. 340

Default Auto Configuration Values **341**

Managing Auto Configuration (Web) **342**
Auto-Install Configuration 342

Managing Auto Configuration (CLI) **343**
Managing Auto Configuration 343

Auto Configuration Example **344**
Enabling DHCP Auto Configuration and Auto Image Download 344

16 Monitoring Switch Traffic 347

Traffic Monitoring Overview. **347**
What is sFlow Technology?. 347
What is RMON?. 350
What is Port Mirroring?. 351
Why is Traffic Monitoring Needed? 352

Default Traffic Monitoring Values **352**

Monitoring Switch Traffic (Web) **353**
sFlow Agent Summary 353
sFlow Receiver Configuration 354
sFlow Sampler Configuration 355
sFlow Poll Configuration 356
Interface Statistics 357
Etherlike Statistics 358

GVRP Statistics 359
EAP Statistics 360
Utilization Summary 361
Counter Summary 362
Switchport Statistics 363
RMON Statistics 364
RMON History Control Statistics 365
RMON History Table 367
RMON Event Control 368
RMON Event Log 370
RMON Alarms 371
Port Statistics 373
LAG Statistics 374
Port Mirroring 375

Monitoring Switch Traffic (CLI) 378
Configuring sFlow 378
Configuring RMON 380
Viewing Statistics 382
Configuring Port Mirroring 383

Traffic Monitoring Configuration Examples 384
Configuring sFlow 384
Configuring RMON 386

17 Configuring iSCSI Optimization 387

iSCSI Optimization Overview 387
What Does iSCSI Optimization Do? 388
How Does the Switch Detect iSCSI Traffic Flows? 388
How Is Quality of Service Applied to iSCSI Traffic Flows? 388
How Does iSCSI Optimization Use ACLs? 389

What Information Does the Switch Track in iSCSI Traffic Flows? 389
How Does iSCSI Optimization Interact With Dell EqualLogic Arrays? 390
What Occurs When iSCSI Optimization Is Enabled or Disabled? 390
How Does iSCSI Optimization Interact with DCBx? 391
How Does iSCSI Optimization Interact with Dell Compellent Arrays? 391

Default iSCSI Optimization Values **392**

Configuring iSCSI Optimization (Web) **393**
iSCSI Global Configuration 393
iSCSI Targets Table 394
iSCSI Sessions Table 395
iSCSI Sessions Detailed 396

Configuring iSCSI Optimization (CLI) **397**

iSCSI Optimization Configuration Examples **399**
Configuring iSCSI Optimization Between Servers and a Disk Array 399

18 Configuring a Captive Portal 401

Captive Portal Overview **401**
What Does a Captive Portal Do? 401
Is the Captive Portal Feature Dependent on Any Other Feature? 402
What Factors Should Be Considered When Designing and Configuring a Captive Portal? . . . 403
How Does Captive Portal Work? 404
What Captive Portal Pages Can Be Customized? 405

Default Captive Portal Behavior and Settings **406**

Configuring the Captive Portal (Web) **408**

Captive Portal Global Configuration 408
Captive Portal Configuration 409
Local User 414
User Group 418
Interface Association. 420
Captive Portal Global Status 421
Captive Portal Activation and Activity Status . . . 422
Interface Activation Status 423
Interface Capability Status 424
Client Summary 425
Client Detail 426
Captive Portal Interface Client Status 427
Captive Portal Client Status 428

Configuring a Captive Portal (CLI) **429**

Configuring Global Captive Portal Settings. 429
Creating and Configuring a Captive Portal 430
Configuring Captive Portal Groups and Users . . . 433
Managing Captive Portal Clients 434

Captive Portal Configuration Example **435**

Configuration Overview. 436
Detailed Configuration Procedures 437

19 Configuring Port Characteristics 439

Port Overview **439**

What Physical Port Characteristics Can Be Configured? 439
What is Link Dependency? 440
What Interface Types are Supported? 442
What is Interface Configuration Mode? 442

Default Port Values **444**

Configuring Port Characteristics (Web) **445**

Port Configuration. 445

Link Dependency Configuration 448

Link Dependency Summary. 450

Configuring Port Characteristics (CLI) **451**

Configuring Port Settings 451

Configuring Link Dependencies 452

Port Configuration Examples **454**

Configuring Port Settings 454

Configuring a Link Dependency Groups 455

20 Configuring Port and System Security 457

IEEE 802.1X **458**

What is IEEE 802.1X? 458

What are the 802.1X Port States?. 459

What is MAC-Based 802.1X Authentication?. . . 460

What is the Role of 802.1X in VLAN Assignment? 461

What is Monitor Mode?. 463

How Does the Authentication Server Assign DiffServ Filters? 465

What is the Internal Authentication Server? . . . 465

Default 802.1X Values 466

Configuring IEEE 802.1X (Web) 467

Configuring IEEE 802.1X (CLI) 476

Configuring Internal Authentication Server Users 481

IEEE 802.1X Configuration Examples 481

Port Security (Port-MAC Locking) **495**
Default 802.1X Values 495
Configuring Port Security (CLI) 498

Denial of Service **499**

21 Configuring Access Control Lists 501

ACL Overview **501**
What Are MAC ACLs? 502
What Are IP ACLs? 503
What Is the ACL Redirect Function? 503
What Is the ACL Mirror Function? 503
What Is ACL Logging 504
What Are Time-Based ACLs? 504
What Are the ACL Limitations? 505
How Are ACLs Configured? 506
Preventing False ACL Matches 506

Configuring ACLs (Web) **508**
IP ACL Configuration 508
IP ACL Rule Configuration 510
MAC ACL Configuration 512
MAC ACL Rule Configuration 514
IPv6 ACL Configuration 515
IPv6 ACL Rule Configuration 516
ACL Binding Configuration 518
Time Range Entry Configuration 519

Configuring ACLs (CLI) **521**
Configuring an IPv4 ACL 521
Configuring a MAC ACL 523
Configuring an IPv6 ACL 525
Configuring a Time Range 527

ACL Configuration Examples **529**
Configuring an IP ACL 529
Configuring a MAC ACL 531
Configuring a Time-Based ACL 533
Configuring a Management Access List 534

22 Configuring VLANs 539

VLAN Overview **539**
Switchport Modes 542
VLAN Tagging 543
GVRP 544
Double-VLAN Tagging 544
Voice VLAN 545
Private VLANs 547
Additional VLAN Features 553

Default VLAN Behavior **554**

Configuring VLANs (Web) **556**
VLAN Membership 556
VLAN Port Settings 561
VLAN LAG Settings 562
Bind MAC to VLAN 564
Bind IP Subnet to VLAN 565
GVRP Parameters 567
Protocol Group 569
Adding a Protocol Group 570
Double VLAN Global Configuration 572
Double VLAN Interface Configuration 573
Voice VLAN 575

Configuring VLANs (CLI) **576**
Creating a VLAN 576

Configuring a Port in Access Mode 577
Configuring a Port in General Mode 578
Configuring a Port in Trunk Mode 579
Configuring VLAN Settings for a LAG 581
Configuring Double VLAN Tagging 582
Configuring MAC-Based VLANs 583
Configuring IP-Based VLANs. 584
Configuring a Protocol-Based VLAN 584
Configuring GVRP. 586
Configuring Voice VLANs. 588

VLAN Configuration Examples 589

Configuring VLANs Using Dell OpenManage Administrator 592
Configure the VLANs and Ports on Switch 2 596
Configuring VLANs Using the CLI. 597
Configuring a Voice VLAN 601

23 Configuring the Spanning Tree Protocol 605

STP Overview 605

What Are Classic STP, Multiple STP, and Rapid STP? 605
How Does STP Work? 606
How Does MSTP Operate in the Network? 607
MSTP with Multiple Forwarding Paths. 611
What are the Optional STP Features? 612

Default STP Values 615

Configuring Spanning Tree (Web) 616

STP Global Settings. 616
STP Port Settings 618
STP LAG Settings 620

Rapid Spanning Tree 621
MSTP Settings 624
MSTP Interface Settings 626

Configuring Spanning Tree (CLI) 628
Configuring Global STP Bridge Settings 628
Configuring Optional STP Features 629
Configuring STP Interface Settings 630
Configuring MSTP Switch Settings 631
Configuring MSTP Interface Settings 632

STP Configuration Examples 633
Configuring STP 633
Configuring MSTP 635

24 Discovering Network Devices 637

Device Discovery Overview 637
What Is ISDP? 637
What is LLDP? 637
What is LLDP-MED? 638
Why are Device Discovery Protocols Needed? 638

Default IDSP and LLDP Values 639

Configuring ISDP and LLDP (Web) 641
ISDP Global Configuration 641
ISDP Cache Table 642
ISDP Interface Configuration 643
ISDP Statistics 645
LLDP Configuration 646
LLDP Statistics 648
LLDP Connections 649
LLDP-MED Global Configuration 651

LLDP-MED Interface Configuration 652
LLDP-MED Local Device Information 654
LLDP-MED Remote Device Information 655

Configuring ISDP and LLDP (CLI) 656
Configuring Global ISDP Settings 656
Enabling ISDP on a Port 657
Viewing and Clearing ISDP Information 657
Configuring Global LLDP Settings 658
Configuring Port-based LLDP Settings 658
Viewing and Clearing LLDP Information 659
Configuring LLDP-MED Settings 660
Viewing LLDP-MED Information 661

Device Discovery Configuration Examples 661
Configuring ISDP 661
Configuring LLDP 662

25 Configuring Port-Based Traffic Control 665

Port-Based Traffic Control Overview 665
What is Flow Control? 666
What is Storm Control? 666
What are Protected Ports? 667
What is Link Local Protocol Filtering? 667

Default Port-Based Traffic Control Values 668

Configuring Port-Based Traffic Control (Web) 669
Flow Control (Global Port Parameters) 669
Storm Control 670
Protected Port Configuration 672
LLPF Configuration 674

Configuring Port-Based Traffic Control (CLI) **676**
Configuring Flow Control and Storm Control . . . 676
Configuring Protected Ports 677
Configuring LLPF 678

Port-Based Traffic Control Configuration Example **679**

26 Configuring L2 Multicast Features 681

L2 Multicast Overview **681**
What Are the Multicast Bridging Features? . . . 681
L2 Multicast Forwarding Modes 682
What Is IP Multicast Traffic? 682
What Is IGMP Snooping? 683
What Is MLD Snooping? 683
What Is Multicast VLAN Registration? 684
When Are L3 Multicast Features Required? . . . 685
What Are GARP and GMRP? 686

Snooping Switch Restrictions **687**
Partial IGMPv3 and MLDv2 Support 687
MAC Address-Based Multicast Group 687
IGMP Snooping in a Multicast Router 687
Topologies Where the Multicast Source Is Not Directly Connected to the Querier 687

Current MLD Snooping Functional Limitations Summary **688**

Default L2 Multicast Values **689**

Configuring L2 Multicast Features (Web) **691**
Multicast Global Parameters 691
Bridge Multicast Group 692

Bridge Multicast Forwarding 695
MRouter Status 696
General IGMP Snooping 697
Global Querier Configuration 700
VLAN Querier 701
VLAN Querier Status 704
MFDB IGMP Snooping Table 705
MLD Snooping General 706
MLD Snooping Global Querier Configuration . . . 708
MLD Snooping VLAN Querier 709
MLD Snooping VLAN Querier Status 712
MFDB MLD Snooping Table 713
MVR Global Configuration 714
MVR Members 715
MVR Interface Configuration 716
MVR Statistics 719
GARP Timers 720
GMRP Parameters 722
MFDB GMRP Table 724

Configuring L2 Multicast Features (CLI) 725
Configuring Bridge Multicasting 725
Configuring IGMP Snooping 727
Configuring IGMP Snooping on VLANs 728
Configuring IGMP Snooping Querier 729
Configuring MLD Snooping 730
Configuring MLD Snooping on VLANs 731
Configuring MLD Snooping Querier 732
Configuring MVR 733
Configuring GARP Timers and GMRP 735

Case Study on a Real-World Network Topology 736
Multicast Snooping Case Study 736

27 Snooping and Inspecting Traffic 743

Traffic Snooping and Inspection Overview 743

What Is DHCP Snooping? 744

How Is the DHCP Snooping Bindings Database Populated? 745

What Is IP Source Guard? 747

What is Dynamic ARP Inspection? 748

Why Is Traffic Snooping and Inspection Necessary? 749

Default Traffic Snooping and Inspection Values . . . 749

Configuring Traffic Snooping and Inspection (Web) 751

DHCP Snooping Configuration 751

DHCP Snooping Interface Configuration 752

DHCP Snooping VLAN Configuration 754

DHCP Snooping Persistent Configuration 756

DHCP Snooping Static Bindings Configuration 757

DHCP Snooping Dynamic Bindings Summary 759

DHCP Snooping Statistics 760

IPSG Interface Configuration 761

IPSG Binding Configuration 762

IPSG Binding Summary 763

DAI Global Configuration 764

DAI Interface Configuration 765

DAI VLAN Configuration 767

DAI ACL Configuration 768

DAI ACL Rule Configuration 769

DAI Statistics 771

Configuring Traffic Snooping and Inspection (CLI) 772

Configuring DHCP Snooping 772

Configuring IP Source Guard 774
Configuring Dynamic ARP Inspection 775

Traffic Snooping and Inspection Configuration Examples 777
Configuring DHCP Snooping 777
Configuring IPSG 779

28 Configuring Link Aggregation 781

Link Aggregation Overview 781
Why Are Link Aggregation Groups Necessary? 782
What Is the Difference Between Static and Dynamic Link Aggregation? 782
What is LAG Hashing? 783
How Do LAGs Interact with Other Features? 784
LAG Configuration Guidelines 785

Default Link Aggregation Values 785

Configuring Link Aggregation (Web) 786
LAG Configuration 786
LACP Parameters 788
LAG Membership 790
LAG Hash Configuration 791
LAG Hash Summary 792

Configuring Link Aggregation (CLI) 793
Configuring LAG Characteristics 793
Configuring Link Aggregation Groups 794
Configuring LACP Parameters 796

Link Aggregation Configuration Examples 797
Configuring Dynamic LAGs 797
Configuring Static LAGs 798

29 Configuring Data Center Bridging Features 799

Data Center Bridging Technology Overview **799**

Default DCB Values 800

Priority Flow Control **801**

PFC Operation and Behavior 801

Configuring PFC Using the Web Interface 802

Configuring PFC Using the CLI 804

PFC Configuration Example 806

DCB Capability Exchange **807**

Interoperability with IEEE DCBx 808

DCBx and Port Roles 808

Configuration Source Port Selection Process 810

Disabling DCBX 811

Configuring DCBx 812

FIP Snooping **814**

Enabling and Disabling FIP Snooping. 815

Configuring the FC Map Value 815

Configuring Ports for FIP Snooping 815

Configuring FIP Snooping (CLI) 815

FIP Snooping Configuration Example 817

Enhanced Transmission Selection **819**

ETS Operation 820

Commands 823

ETS Configuration Example 824

ETS Theory of Operation 830

30 Managing the MAC Address Table 837

MAC Address Table Overview **837**
How Is the Address Table Populated? 837
What Information Is in the MAC Address Table? 838
How Is the MAC Address Table Maintained Across a Stack? 838

Default MAC Address Table Values **838**

Managing the MAC Address Table (Web) **839**
Static Address Table 839
Dynamic Address Table 841

Managing the MAC Address Table (CLI) **842**
Managing the MAC Address Table 842

31 Configuring Routing Interfaces 843

Routing Interface Overview **843**
What Are VLAN Routing Interfaces? 843
What Are Loopback Interfaces? 844
What Are Tunnel Interfaces? 845
Why Are Routing Interfaces Needed? 846

Default Routing Interface Values **848**

Configuring Routing Interfaces (Web) **849**
IP Interface Configuration 849
DHCP Lease Parameters 850
VLAN Routing Summary 850
Tunnel Configuration 851
Tunnels Summary 852
Loopbacks Configuration 853

Loopbacks Summary 854

Configuring Routing Interfaces (CLI) 855

Configuring VLAN Routing Interfaces (IPv4) . . . 855

Configuring Loopback Interfaces. 857

Configuring Tunnels. 858

32 Configuring DHCP Server Settings 859

DHCP Overview 859

How Does DHCP Work?. 859

What are DHCP Options? 860

What Additional DHCP Features Does the Switch Support? 861

Default DHCP Server Values. 861

Configuring the DHCP Server (Web). 862

DHCP Server Network Properties 862

Address Pool 864

Address Pool Options. 868

DHCP Bindings 870

DHCP Server Reset Configuration 871

DHCP Server Conflicts Information. 872

DHCP Server Statistics 873

Configuring the DHCP Server (CLI) 874

Configuring Global DHCP Server Settings 874

Configuring a Dynamic Address Pool. 875

Configuring a Static Address Pool 876

Monitoring DHCP Server Information 877

DHCP Server Configuration Examples. 878

Configuring a Dynamic Address Pool. 878

Configuring a Static Address Pool 880

33 Configuring IP Routing 883

IP Routing Overview **883**

Default IP Routing Values **885**

Configuring IP Routing Features (Web) **887**

IP Configuration. 887

IP Statistics 888

ARP Create 889

ARP Table Configuration 890

Router Discovery Configuration 891

Router Discovery Status 892

Route Table 893

Best Routes Table 894

Route Entry Configuration 895

Configured Routes 897

Route Preferences Configuration 898

Configuring IP Routing Features (CLI) **899**

Configuring Global IP Routing Settings. 899

Adding Static ARP Entries and Configuring ARP Table Settings 900

Configuring Router Discovery (IRDP). 901

Configuring Route Table Entries and Route Preferences. 902

IP Routing Configuration Example **904**

Configuring PowerConnect Switch A 905

Configuring PowerConnect Switch B 906

34 Configuring L2 and L3 Relay Features 907

L2 and L3 Relay Overview **907**

What Is L3 DHCP Relay? 907

What Is L2 DHCP Relay? 908
What Is the IP Helper Feature? 909

Default L2/L3 Relay Values **913**

Configuring L2 and L3 Relay Features (Web) **914**
DHCP Relay Global Configuration. 914
DHCP Relay Interface Configuration 915
DHCP Relay Interface Statistics 917
DHCP Relay VLAN Configuration 918
DHCP Relay Agent Configuration 919
IP Helper Global Configuration 920
IP Helper Interface Configuration 922
IP Helper Statistics 924

Configuring L2 and L3 Relay Features (CLI) **925**
Configuring L2 DHCP Relay 925
Configuring L3 Relay (IP Helper) Settings 927

Relay Agent Configuration Example **929**

35 Configuring OSPF and OSPFv3 931

OSPF Overview **932**
What Are OSPF Areas and Other OSPF Topology Features? 932
What Are OSPF Routers and LSAs? 933
How Are Routes Selected? 933
How Are OSPF and OSPFv3 Different? 933

OSPF Feature Details **934**
Max Metric 934
Static Area Range Cost 936
LSA Pacing 937
Flood Blocking 938

Default OSPF Values **940**

Configuring OSPF Features (Web) **942**

OSPF Configuration 942
OSPF Area Configuration 944
OSPF Stub Area Summary 947
OSPF Area Range Configuration 948
OSPF Interface Statistics 949
OSPF Interface Configuration 950
OSPF Neighbor Table 951
OSPF Neighbor Configuration 952
OSPF Link State Database 953
OSPF Virtual Link Configuration 953
OSPF Virtual Link Summary 955
OSPF Route Redistribution Configuration 956
OSPF Route Redistribution Summary 957
NSF OSPF Configuration 958

Configuring OSPFv3 Features (Web) **959**

OSPFv3 Configuration 959
OSPFv3 Area Configuration 960
OSPFv3 Stub Area Summary 963
OSPFv3 Area Range Configuration 964
OSPFv3 Interface Configuration 965
OSPFv3 Interface Statistics 966
OSPFv3 Neighbors 967
OSPFv3 Neighbor Table 968
OSPFv3 Link State Database 969
OSPFv3 Virtual Link Configuration 970
OSPFv3 Virtual Link Summary 972
OSPFv3 Route Redistribution Configuration 973
OSPFv3 Route Redistribution Summary 974
NSF OSPFv3 Configuration 975

Configuring OSPF Features (CLI) **976**

Configuring Global OSPF Settings 976

Configuring OSPF Interface Settings 979
Configuring Stub Areas and NSSAs 981
Configuring Virtual Links 983
Configuring OSPF Area Range Settings 985

Configuring OSPFv3 Features (CLI) 987
Configuring Global OSPFv3 Settings 987
Configuring OSPFv3 Interface Settings. 989
Configuring Stub Areas and NSSAs 991
Configuring Virtual Links 993
Configuring an OSPFv3 Area Range 994
Configuring OSPFv3 Route Redistribution Settings 995

OSPF Configuration Examples 996
Configuring an OSPF Border Router and Setting Interface Costs 996
Configuring Stub and NSSA Areas for OSPF and OSPFv3 999
Configuring a Virtual Link for OSPF and OSPFv3 1003
Interconnecting an IPv4 Backbone and Local IPv6 Network 1005
Configuring the Static Area Range Cost 1008
Configuring Flood Blocking 1013

36 Configuring RIP 1019

RIP Overview 1019
How Does RIP Determine Route Information?. 1019
What Is Split Horizon? 1020
What RIP Versions Are Supported? 1020

Default RIP Values **1021**

Configuring RIP Features (Web) **1022**

RIP Configuration 1022

RIP Interface Configuration 1023

RIP Interface Summary 1024

RIP Route Redistribution Configuration 1025

RIP Route Redistribution Summary 1026

Configuring RIP Features (CLI) **1027**

Configuring Global RIP Settings 1027

Configuring RIP Interface Settings 1028

Configuring Route Redistribution Settings 1029

RIP Configuration Example **1031**

37 Configuring VRRP 1035

VRRP Overview **1035**

How Does VRRP Work? 1035

What Is the VRRP Router Priority? 1036

What Is VRRP Preemption? 1036

What Is VRRP Accept Mode? 1037

What Are VRRP Route and Interface Tracking? 1037

Default VRRP Values **1039**

Configuring VRRP Features (Web) **1040**

VRRP Configuration 1040

VRRP Virtual Router Status 1041

VRRP Virtual Router Statistics 1042

VRRP Router Configuration 1043

VRRP Route Tracking Configuration 1044

VRRP Interface Tracking Configuration 1046

Configuring VRRP Features (CLI) **1048**
Configuring VRRP Settings 1048

VRRP Configuration Example **1050**
VRRP with Load Sharing 1050
VRRP with Route and Interface Tracking 1054

38 Configuring IPv6 Routing 1059

IPv6 Routing Overview **1059**
How Does IPv6 Compare with IPv4? 1060
How Are IPv6 Interfaces Configured? 1060

Default IPv6 Routing Values **1061**

Configuring IPv6 Routing Features (Web) **1063**
Global Configuration 1063
Interface Configuration 1064
Interface Summary 1065
IPv6 Statistics 1066
IPv6 Neighbor Table 1067
DHCPv6 Client Parameters 1068
IPv6 Route Entry Configuration 1069
IPv6 Route Table 1070
IPv6 Route Preferences 1071
Configured IPv6 Routes 1072

Configuring IPv6 Routing Features (CLI) **1073**
Configuring Global IP Routing Settings 1073
Configuring IPv6 Interface Settings 1074
Configuring IPv6 Neighbor Discovery 1075
Configuring IPv6 Route Table Entries and Route Preferences 1077
IPv6 Show Commands 1079

IPv6 Static Reject and Discard Routes **1080**

39 Configuring DHCPv6 Server and Relay Settings 1083

DHCPv6 Overview **1083**

What Is a DHCPv6 Pool? 1084

What Is a Stateless Server? 1084

What Is the DHCPv6 Relay Agent Information Option? 1084

What Is a Prefix Delegation? 1084

Default DHCPv6 Server and Relay Values **1085**

Configuring the DHCPv6 Server and Relay (Web) . . . **1086**

DHCPv6 Global Configuration 1086

DHCPv6 Pool Configuration. 1087

Prefix Delegation Configuration 1089

DHCPv6 Pool Summary 1090

DHCPv6 Interface Configuration 1091

DHCPv6 Server Bindings Summary 1093

DHCPv6 Statistics. 1094

Configuring the DHCPv6 Server and Relay (CLI) . . . **1095**

Configuring Global DHCP Server and Relay Agent Settings 1095

Configuring a DHCPv6 Pool for Stateless Server Support 1095

Configuring a DHCPv6 Pool for Specific Hosts 1096

Configuring DHCPv6 Interface Information . . . 1097

Monitoring DHCPv6 Information 1098

DHCPv6 Configuration Examples **1099**

Configuring a DHCPv6 Stateless Server 1099

Configuring the DHCPv6 Server for Prefix Delegation. 1100
Configuring an Interface as a DHCPv6 Relay Agent 1101

40 Configuring Differentiated Services . . 1103

DiffServ Overview **1103**
How Does DiffServ Functionality Vary Based on the Role of the Switch? 1104
What Are the Elements of DiffServ Configuration? 1104

Default DiffServ Values **1105**

Configuring DiffServ (Web) **1106**
DiffServ Configuration 1106
Class Configuration 1107
Class Criteria 1108
Policy Configuration 1110
Policy Class Definition 1112
Service Configuration. 1115
Service Detailed Statistics 1117
Flow-Based Mirroring 1118

Configuring DiffServ (CLI) **1119**
DiffServ Configuration (Global) 1119
DiffServ Class Configuration for IPv4. 1119
DiffServ Class Configuration for IPv6. 1121
DiffServ Policy Creation. 1122
DiffServ Policy Attributes Configuration 1123
DiffServ Service Configuration 1125

DiffServ Configuration Examples **1126**
Providing Subnets Equal Access to External Network 1126

DiffServ for VoIP 1130

41 Configuring Class-of-Service 1133

CoS Overview **1133**

What Are Trusted and Untrusted Port Modes? 1134

How Is Traffic Shaping Used on Egress Traffic? 1134

How Are Traffic Queues Defined? 1135

Which Queue Management Methods Are Supported? 1135

CoS Queue Usage 1136

Default CoS Values **1136**

Configuring CoS (Web) **1137**

Mapping Table Configuration. 1137

Interface Configuration. 1140

Interface Queue Configuration 1141

Interface Queue Drop Precedence Configuration 1142

Configuring CoS (CLI) **1144**

Mapping Table Configuration. 1144

CoS Interface Configuration Commands 1145

Interface Queue Configuration 1145

Configuring Interface Queue Drop Precedence. 1146

CoS Configuration Example **1147**

42 Configuring Auto VoIP 1151

Auto VoIP Overview **1151**
How Does Auto-VoIP Use ACLs? 1152

Default Auto VoIP Values **1152**

Configuring Auto VoIP (Web) **1153**
Auto VoIP Global Configuration. 1153
Auto VoIP Interface Configuration 1153

Configuring Auto VoIP (CLI) **1156**

43 Managing IPv4 and IPv6 Multicast . . . 1157

L3 Multicast Overview **1157**
What Is IP Multicast Traffic? 1158
What Multicast Protocols Does the Switch Support? 1159
What Are the Multicast Protocol Roles? 1159
When Is L3 Multicast Required on the Switch? 1160
What Is the Multicast Routing Table? 1160
What Is Multicast Tunneling?. 1161
What Is IGMP? 1161
What Is MLD?. 1162
What Is PIM? 1163
What Is DVMRP? 1165

Default L3 Multicast Values **1167**

Configuring General IPv4 Multicast Features (Web) **1169**

Multicast Global Configuration 1169

Multicast Interface Configuration 1170

Multicast Route Table 1171

Multicast Admin Boundary Configuration 1172

Multicast Admin Boundary Summary 1173

Multicast Static MRoute Configuration 1174

Multicast Static MRoute Summary. 1175

Configuring IPv6 Multicast Features (Web) **1176**

IPv6 Multicast Route Table 1176

Configuring IGMP and IGMP Proxy (Web) **1177**

IGMP Global Configuration 1177

IGMP Interface Configuration 1178

IGMP Interface Summary 1179

IGMP Cache Information 1180

IGMP Interface Source List Information 1181

IGMP Proxy Interface Configuration 1182

IGMP Proxy Configuration Summary. 1183

IGMP Proxy Interface Membership Info 1184

Detailed IGMP Proxy Interface Membership Information 1185

Configuring MLD and MLD Proxy (Web) **1186**

MLD Global Configuration 1186

MLD Routing Interface Configuration 1187

MLD Routing Interface Summary. 1188

MLD Routing Interface Cache Information. . . . 1189

MLD Routing Interface Source List Information 1190

MLD Traffic 1191

MLD Proxy Configuration. 1192

MLD Proxy Configuration Summary 1193
MLD Proxy Interface Membership Information 1194
Detailed MLD Proxy Interface Membership Information 1195

Configuring PIM for IPv4 and IPv6 (Web) 1196

PIM Global Configuration 1196
PIM Global Status 1198
PIM Interface Configuration 1199
PIM Interface Summary 1200
Candidate RP Configuration 1201
Static RP Configuration 1203
SSM Range Configuration 1205
BSR Candidate Configuration 1207
BSR Candidate Summary 1208

Configuring DVMRP (Web) 1209

DVMRP Global Configuration 1209
DVMRP Interface Configuration 1210
DVMRP Configuration Summary 1211
DVMRP Next Hop Summary 1212
DVMRP Prune Summary 1213
DVMRP Route Summary 1214

Configuring L3 Multicast Features (CLI) 1215

Configuring and Viewing IPv4 Multicast Information 1215
Configuring and Viewing IPv6 Multicast Route Information 1217
Configuring and Viewing IGMP 1218
Configuring and Viewing IGMP Proxy 1220
Configuring and Viewing MLD 1221
Configuring and Viewing MLD Proxy 1222
Configuring and Viewing PIM-DM for IPv4 Multicast Routing 1223

Configuring and Viewing PIM-DM for IPv6 Multicast Routing 1224
Configuring and Viewing PIM-SM for IPv4 Multicast Routing 1225
Configuring and Viewing PIM-SM for IPv6 Multicast Routing 1227
Configuring and Viewing DVMRP Information 1230

L3 Multicast Configuration Examples 1231

Configuring Multicast VLAN Routing With IGMP and PIM-SM 1231
Configuring DVMRP 1235

A Feature Limitations and Platform Constants 1237

Index 1247

1

Introduction

The Dell PowerConnect 8024, 8024F, 8132, 8132F, 8164, and 8164F switches are stackable Layer 2 and Layer 3 switches that extend the Dell PowerConnect LAN switching product range.

> **NOTE:** Throughout this document, the PowerConnect 8024 and 8024F switches are referred to as the PowerConnect 8000-series switches, and the PowerConnect 8132, 8132F, 8164, 8164F switches are referred to as the PowerConnect 8100-series switches.

These switches include the following features:

- 1U form factor, rack-mountable chassis design.
- Support for all data-communication requirements for a multi-layer switch, including layer 2 switching, IPv4 routing, IPv6 routing, IP multicast, quality of service, security, and system management features.
- High availability with hot swappable stack members.

Each of the PowerConnect 8000-series and 8100-series switches has 24 or 48 ports of 10-Gb Ethernet in 10GBase-T or SFP+ with redundant power supplies to provide high performance and high availability. PowerConnect 8000-series and 8100-series switches can be stacked with other PowerConnect switches of the same model number using the 10G SFP+ or QSFP fiber ports.

About This Document

This guide describes how to configure, monitor, and maintain a Dell PowerConnect 8000-series and 8100-series switch by using web-based Dell OpenManage Switch Administrator utility or the command-line interface (CLI).

Audience

This guide is for network administrators in charge of managing one or more PowerConnect 8024, 8024F, 8132, 8132F, 8164, and 8164F switches. To obtain the greatest benefit from this guide, you should have a basic understanding of Ethernet networks and local area network (LAN) concepts.

Document Conventions

Table 1-1 describes the typographical conventions this document uses.

Table 1-1. Document Conventions

Convention	Description
Bold	Page names, field names, menu options, button names, and CLI commands and keywords.
`courier font`	Command-line text (CLI output) and file names
[ ]	In a command line, square brackets indicate an optional entry.
{ }	In a command line, inclusive brackets indicate a selection of compulsory parameters separated by the | character. One option must be selected. For example: **spanning-tree mode {stp|rstp|mstp}** means that for the **spanning-tree mode** command you must enter either **stp**, **rstp**, or **mstp**
Italic	In a command line, indicates a variable.
<Enter>	Any individual key on the keyboard.
CTRL + Z	A keyboard combination that involves pressing the Z key while holding the CTRL key.

Additional Documentation

The following documents for the PowerConnect 8024, 8024F, 8132, 8132F, 8164, and 8164F switches are available at support.dell.com/manuals:

- *Getting Started Guide*—provides information about the switch models in the series, including front and back panel features. It also describes the installation and initial configuration procedures.
- *CLI Reference Guide*—provides information about the command-line interface (CLI) commands used to configure and manage the switch. The document provides in-depth CLI descriptions, syntax, default values, and usage guidelines.

2

Switch Features

This section describes the switch user-configurable software features.

NOTE: Before proceeding, read the release notes for this product. The release notes are part of the firmware download.

The topics covered in this section include:

- System Management Features
- Stacking Features
- Security Features
- Switching Features
- Virtual Local Area Network Supported Features
- Spanning Tree Protocol Features
- Link Aggregation Features
- Routing Features
- IPv6 Routing Features
- Quality of Service (QoS) Features
- Layer 2 Multicast Features
- Layer 3 Multicast Features

System Management Features

Multiple Management Options

You can use any of the following methods to manage the switch:

- Use a web browser to access the Dell OpenManage Switch Administrator interface. The switch contains an embedded Web server that serves HTML pages.
- Use a telnet client, SSH client, or a direct console connection to access the CLI. The CLI syntax and semantics conform as much as possible to common industry practice.
- Use a network management system (NMS) to manage and monitor the system through SNMP. The switch supports SNMP v1/v2c/v3 over the UDP/IP transport protocol.

System Time Management

You can configure the switch to obtain the system time and date through a remote Simple Network Time Protocol (SNTP) server, or you can set the time and date locally on the switch. You can also configure the time zone and information about time shifts that might occur during summer months. If you use SNTP to obtain the time, you can require communications between the switch and the SNTP server to be encrypted.

For information about configuring system time settings, see "Managing General System Settings" on page 237.

Log Messages

The switch maintains in-memory log messages as well as persistent logs. You can also configure remote logging so that the switch sends log messages to a remote log server. You can also configure the switch to send log messages to a configured SMTP server. This allows you to receive the log message in an e-mail account of your choice. Switch auditing messages, CLI command logging, Web logging, and SNMP logging can be enabled or disabled.

For information about configuring system logging, see "Monitoring and Logging System Information" on page 203.

Integrated DHCP Server

PowerConnect 8000-series and 8100-series switches include an integrated DHCP server that can deliver host-specific configuration information to hosts on the network. The switch DHCP server allows you to configure IP address pools (scopes), and when a host's DHCP client requests an address, the switch DHCP server automatically assigns the host an address from the pool.

For information about configuring the DHCP server settings, see "Configuring DHCP Server Settings" on page 859.

Management of Basic Network Information

The DHCP client on the switch allows the switch to acquire information such as the IP address and default gateway from a network DHCP server. You can also disable the DHCP client and configure static network information. Other configurable network information includes a Domain Name Server (DNS), hostname to IP address mapping, and a default domain name.

If the switch detects an IP address conflict on the management interface, it generates a trap and sends a log message.

For information about configuring basic network information, see "Setting the IP Address and Other Basic Network Information" on page 119.

IPv6 Management Features

PowerConnect 8000-series and 8100-series switches provide IPv6 support for many standard management features including HTTP, HTTPS/SSL, Telnet, SSH, SNMP, SNTP, TFTP, and traceroute.

Dual Software Images

PowerConnect 8000-series and 8100-series switches can store up to two software images. The dual image feature allows you to upgrade the switch without deleting the older software image. You designate one image as the active image and the other image as the backup image.

For information about managing the switch image, see "Managing Images and Files" on page 307.

File Management

You can upload and download files such as configuration files and system images by using HTTP (web only), TFTP, Secure FTP (SFTP), or Secure Copy (SCP). Configuration file uploads from the switch to a server are a good way to back up the switch configuration. You can also download a configuration file from a server to the switch to restore the switch to the configuration in the downloaded file.

For information about uploading, downloading, and copying files, see "Managing Images and Files" on page 307.

Switch Database Management Templates

Switch Database Management (SDM) templates enable you to reallocate system resources to support a different mix of features based on your network requirements. PowerConnect 8000-series and 8100-series switches support the following three templates:

- Dual IPv4 and IPv6 (default)
- IPv4 Routing
- IPv4 Data Center

For information about setting the SDM template, see "Managing General System Settings" on page 237.

Automatic Installation of Firmware and Configuration

The Auto Install feature allows the switch to upgrade or downgrade to a newer software image and update the configuration file automatically during device initialization with limited administrative configuration on the device. The switch can obtain the necessary information from a DHCP server on the network.

For information about Auto Install, see "Automatically Updating the Image and Configuration" on page 333.

sFlow

sFlow is the standard for monitoring high-speed switched and routed networks. sFlow technology is built into network equipment and gives complete visibility into network activity, enabling effective management and control of network resources. The PowerConnect 8000-series and 8100-series switches support sFlow version 5.

For information about configuring managing sFlow settings, see "Monitoring Switch Traffic" on page 347.

SNMP Alarms and Trap Logs

The system logs events with severity codes and timestamps. The events are sent as SNMP traps to a trap recipient list.

For information about configuring SNMP traps and alarms, see "Configuring SNMP" on page 271.

CDP Interoperability through ISDP

Industry Standard Discovery Protocol (ISDP) allows the PowerConnect switch to interoperate with Cisco devices running the Cisco Discovery Protocol (CDP). ISDP is a proprietary Layer 2 network protocol which interoperates with Cisco network equipment and is used to share information between neighboring devices (routers, bridges, access servers, and switches).

For information about configuring ISDP settings, see "Discovering Network Devices" on page 637.

Remote Monitoring (RMON)

RMON is a standard Management Information Base (MIB) that defines current and historical MAC-layer statistics and control objects, allowing real-time information to be captured across the entire network.

For information about configuring managing RMON settings, see "Monitoring Switch Traffic" on page 347.

Stacking Features

For information about creating and maintaining a stack of switches, see "Managing a Switch Stack" on page 141.

High Port Count

You can stack PowerConnect 8000-series and 8100-series switches up to six switches high, supporting up to 132 front-panel ports when two ports on each unit are configured as stacking ports. The stack can contain any combination of Power Connect 8024 and PowerConnect 8024F switches as long as all switches are running the same firmware version.

Single IP Management

When multiple switches are connected together through the stack ports, they operate as a single unit with a larger port count. The stack operates and is managed as a single entity. One switch acts as the master, and the entire stack is managed through the management interface (Web, CLI, or SNMP) of the stack master.

Automatic Firmware Update for New Stack Members

By default, if a switch is added to a stack and the switch is running a different backup version of firmware than the active version on the stack master, the backup firmware on the new member is automatically updated to match the stack master, the backup version of firmware on the new member is activated, and the new member is rebooted.

Master Failover with Transparent Transition

The stacking feature supports a *standby* or backup unit that assumes the stack master role if the stack master fails. As soon as a stack master failure is detected, the standby unit initializes the control plane and enables all other stack units with the current configuration. The standby unit maintains a synchronized copy of the running configuration for the stack.

Nonstop Forwarding on the Stack

The Nonstop Forwarding (NSF) feature allows the forwarding plane of stack units to continue to forward packets while the control and management planes restart as a result of a power failure, hardware failure, or software fault on the stack master and allows the standby switch to quickly takeover as the master.

Hot Add/Delete and Firmware Synchronization

You can add and remove units to and from the stack without cycling the power. When you add a unit, the Stack Firmware Synchronization feature automatically synchronizes the firmware version with the version running on the stack master. The synchronization operation may result in either an upgrade or a downgrade of firmware on the mismatched stack member. In addition, the running-config on the member is updated to match the master switch. The startup-config on the standby and member switches is not updated to match the master switch. The hardware configuration of every switch is updated to match the master switch (unit number, slot configuration, stack member number, etc.).

Security Features

Configurable Access and Authentication Profiles

You can configure rules to limit access to the switch management interface based on criteria such as access type and source IP address of the management host. You can also require the user to be authenticated locally or by an external server, such as a RADIUS server.

For information about configuring access and authentication profiles, see "Configuring Authentication, Authorization, and Accounting" on page 177.

Password-Protected Management Access

Access to the Web, CLI, and SNMP management interfaces is password protected, and there are no default users on the system.

For information about configuring local user accounts, see "Configuring Authentication, Authorization, and Accounting" on page 177.

Strong Password Enforcement

The Strong Password feature enforces a baseline password strength for all locally administered users. Password strength is a measure of the effectiveness of a password in resisting guessing and brute-force attacks. The strength of a password is a function of length, complexity and randomness. Using strong passwords lowers overall risk of a security breach.

For information about configuring password settings, see "Configuring Authentication, Authorization, and Accounting" on page 177.

TACACS+ Client

The switch has a TACACS+ client. TACACS+ provides centralized security for validation of users accessing the switch. TACACS+ provides a centralized user management system while still retaining consistency with RADIUS and other authentication processes.

For information about configuring TACACS+ client settings, see "Configuring Authentication, Authorization, and Accounting" on page 177.

RADIUS Support

The switch has a Remote Authentication Dial In User Service (RADIUS) client and can support up to 32 named authentication and accounting RADIUS servers. The switch also supports RADIUS Attribute 4, which is the configuration of a NAS-IP address. You can also configure the switch to accept RADIUS-assigned VLANs.

For information about configuring RADIUS client settings, see "Configuring Authentication, Authorization, and Accounting" on page 177.

SSH/SSL

The switch supports Secure Shell (SSH) for secure, remote connections to the CLI and Secure Sockets Layer (SSL) to increase security when accessing the web-based management interface.

For information about configuring SSH and SSL settings, see "Configuring Authentication, Authorization, and Accounting" on page 177.

Inbound Telnet Control

You can configure the switch to prevent new Telnet sessions from being established with the switch. Additionally, the Telnet port number is configurable.

For information about configuring inbound Telnet settings, see "Configuring Authentication, Authorization, and Accounting" on page 177.

Denial of Service

The switch supports configurable Denial of Service (DoS) attack protection for eight different types of attacks.

For information about configuring DoS settings, see "Configuring Port and System Security" on page 457.

Port Protection

A port may be put into the disabled state for any of the following reasons:

- BPDU Storm Protection: By default, if Spanning Tree Protocol (STP) bridge protocol data units (BPDUs) are received at a rate of 15pps or greater for three consecutive seconds on a port, the port will be diagnostically disabled. The threshold is not configurable.
- DHCP Snooping: If DHCP packets are received on a port at a rate that exceeds 15 pps, the port will be diagnostically disabled. The threshold is configurable up to 300 pps for up to 15s long using the **`ip dhcp snooping limit`** command. DHCP snooping is disabled by default. The default protection limit is 15 pps.

- Dynamic ARP Inspection: By default, if Dynamic ARP Inspection packets are received on a port at a rate that exceeds 15 pps for 1 second, the port will be diagnostically disabled. The threshold is configurable up to 300 pps and the burst is configurable up to 15s long using the **`ip arp inspection limit`** command.

A port that is diagnostically disabled due to exceeding one of the above limits may be returned to service using the **`no shut`** command.

Captive Portal

The Captive Portal feature blocks clients from accessing the network until user verification has been established. When a user attempts to connect to the network through the switch, the user is presented with a customized Web page that might contain username and password fields or the acceptable use policy. You can require users to be authenticated by a local or remote RADIUS database before access is granted.

For information about configuring the Captive Portal features, see "Configuring a Captive Portal" on page 401.

Dot1x Authentication (IEEE 802.1X)

Dot1x authentication enables the authentication of system users through a local internal server or an external server. Only authenticated and approved system users can transmit and receive data. Supplicants are authenticated using the Extensible Authentication Protocol (EAP). Also supported are PEAP, EAP-TTL, EAP-TTLS, and EAP-TLS.

For information about configuring IEEE 802.1X settings, see "Configuring Port and System Security" on page 457.

MAC-Based 802.1X Authentication

MAC-based authentication allows multiple supplicants connected to the same port to each authenticate individually. For example, a system attached to the port might be required to authenticate in order to gain access to the network, while a VoIP phone might not need to authenticate in order to send voice traffic through the port.

For information about configuring MAC-based 802.1X authentication, see "Configuring Port and System Security" on page 457.

Dot1x Monitor Mode

Monitor mode can be enabled in conjunction with Dot1x authentication to allow network access even when the user fails to authenticate. The switch logs the results of the authentication process for diagnostic purposes. The main purpose of this mode is to help troubleshoot the configuration of a Dot1x authentication on the switch without affecting the network access to the users of the switch.

For information about enabling the Dot1X Monitor mode, see "Configuring Port and System Security" on page 457.

MAC-Based Port Security

The port security feature limits access on a port to users with specific MAC addresses. These addresses are manually defined or learned on that port. When a frame is seen on a locked port, and the frame source MAC address is not tied to that port, the protection mechanism is invoked.

For information about configuring MAC-based port security, see "Configuring Port and System Security" on page 457.

Access Control Lists (ACL)

Access Control Lists (ACLs) ensure that only authorized users have access to specific resources while blocking off any unwarranted attempts to reach network resources. ACLs are used to provide traffic flow control, restrict contents of routing updates, decide which types of traffic are forwarded or blocked, and above all provide security for the network. The switch supports the following ACL types:

- IPv4 ACLs
- IPv6 ACLs
- MAC ACLs

For all ACL types, you can apply the ACL rule when the packet enters or exits the physical port, LAG, or VLAN interface.

For information about configuring ACLs, see "Configuring Access Control Lists" on page 501.

Time-Based ACLs

With the Time-based ACL feature, you can define when an ACL is in effect and the amount of time it is in effect.

For information about configuring time-based ACLs, see "Configuring Access Control Lists" on page 501.

IP Source Guard (IPSG)

IP source guard (IPSG) is a security feature that filters IP packets based on the source ID. The source ID may either be source IP address or a source IP address source MAC address pair.

For information about configuring IPSG, see "Snooping and Inspecting Traffic" on page 743.

DHCP Snooping

DHCP Snooping is a security feature that monitors DHCP messages between a DHCP client and DHCP server. It filters harmful DHCP messages and builds a bindings database of (MAC address, IP address, VLAN ID, port) tuples that are specified as authorized. DHCP snooping can be enabled globally and on specific VLANs. Ports within the VLAN can be configured to be trusted or untrusted. DHCP servers must be reached through trusted ports.

For information about configuring DHCP Snooping, see "Snooping and Inspecting Traffic" on page 743.

Dynamic ARP Inspection

Dynamic ARP Inspection (DAI) is a security feature that rejects invalid and malicious ARP packets. The feature prevents a class of man-in-the-middle attacks, where an unfriendly station intercepts traffic for other stations by poisoning the ARP caches of its unsuspecting neighbors. The malicious station sends ARP requests or responses mapping another station's IP address to its own MAC address.

Dynamic ARP Inspection relies on DHCP Snooping.

For information about configuring DAI, see "Snooping and Inspecting Traffic" on page 743.

Protected Ports (Private VLAN Edge)

Private VLAN Edge (PVE) ports are a Layer 2 security feature that provides port-based security between ports that are members of the same VLAN. It is an extension of the common VLAN. Traffic from protected ports is sent only to the uplink ports and cannot be sent to other ports within the VLAN.

For information about configuring IPSG, see "Configuring Port-Based Traffic Control" on page 665.

Switching Features

Flow Control Support (IEEE 802.3x)

Flow control enables lower speed switches to communicate with higher speed switches by requesting that the higher speed switch refrains from sending packets. Transmissions are temporarily halted to prevent buffer overflows.

For information about configuring flow control, see "Configuring Port-Based Traffic Control" on page 665.

Head of Line Blocking Prevention

Head of Line (HOL) blocking prevention prevents traffic delays and frame loss caused by traffic competing for the same egress port resources. HOL blocking queues packets, and the packets at the head of the queue are forwarded before packets at the end of the queue.

Alternate Store and Forward (ASF)

The Alternate Store and Forward (ASF) feature reduces latency for large packets. When ASF is enabled, the memory management unit (MMU) can forward a packet to the egress port before it has been entirely received on the Cell Buffer Pool (CBP) memory.

AFS, which is also known as cut-through mode, is configurable through the command-line interface. For information about how to configure the AFS feature, see the *CLI Reference Guide* available at support.dell.com/manuals.

Jumbo Frames Support

Jumbo frames enable transporting data in fewer frames to ensure less overhead, lower processing time, and fewer interrupts.

For information about configuring the port MTU, see "Configuring Port Characteristics" on page 439.

Auto-MDI/MDIX Support

Your switch supports auto-detection between crossed and straight-through cables. Media-Dependent Interface (MDI) is the standard wiring for end stations, and the standard wiring for hubs and switches is known as Media-Dependent Interface with Crossover (MDIX).

VLAN-Aware MAC-based Switching

Packets arriving from an unknown source address are sent to the CPU and added to the Hardware Table. Future packets addressed to or from this address are more efficiently forwarded.

Back Pressure Support

On half-duplex links, a receiver may prevent buffer overflows by jamming the link so that it is unavailable for additional traffic. On full-duplex links, a receiver may send a PAUSE frame indicating that the transmitter should cease transmission of frames for a specified period.

When flow control is enabled, the PowerConnect 8000-series and 8100-series switches will observe received PAUSE frames or jamming signals, but will not issue them when congested.

Auto Negotiation

Auto negotiation allows the switch to advertise modes of operation. The auto negotiation function provides the means to exchange information between two switches that share a point-to-point link segment, and to automatically configure both switches to take maximum advantage of their transmission capabilities.

PowerConnect 8000-series and 8100-series switches enhance auto negotiation by providing configuration of port advertisement. Port advertisement allows the system administrator to configure the port speeds that are advertised.

For information about configuring auto negotiation, see "Configuring Port Characteristics" on page 439.

Broadcast Storm Control

When Layer 2 frames are forwarded, broadcast, unknown unicast, and multicast frames are flooded to all ports on the relevant virtual local area network (VLAN). The flooding occupies bandwidth, and loads all nodes connected on all ports. Storm control limits the amount of broadcast, unknown unicast, and multicast frames accepted and forwarded by the switch.

For information about configuring Broadcast Storm Control settings, see "Configuring Port-Based Traffic Control" on page 665.

Port Mirroring

Port mirroring monitors and mirrors network traffic by forwarding copies of incoming and outgoing packets from up to four source ports to a monitoring port. The switch also supports flow-based mirroring, which allows you to copy certain types of traffic to a single destination port. This provides flexibility—instead of mirroring all ingress or egress traffic on a port the switch can mirror a subset of that traffic. You can configure the switch to mirror flows based on certain kinds of Layer 2, Layer 3, and Layer 4 information.

For information about configuring port mirroring, see "Monitoring Switch Traffic" on page 347.

Static and Dynamic MAC Address Tables

You can add static entries to the switch's MAC address table and configure the aging time for entries in the dynamic MAC address table. You can also search for entries in the dynamic table based on several different criteria.

For information about viewing and managing the MAC address table, see "Managing the MAC Address Table" on page 837.

Link Layer Discovery Protocol (LLDP)

The IEEE 802.1AB defined standard, Link Layer Discovery Protocol (LLDP), allows the switch to advertise major capabilities and physical descriptions. This information can help you identify system topology and detect bad configurations on the LAN.

For information about configuring LLDP, settings see "Discovering Network Devices" on page 637.

Link Layer Discovery Protocol (LLDP) for Media Endpoint Devices

The Link Layer Discovery Protocol for Media Endpoint Devices (LLDP-MED) provides an extension to the LLDP standard for network configuration and policy, device location, Power over Ethernet management, and inventory management.

For information about configuring LLDP-MED, settings see "Discovering Network Devices" on page 637.

Priority-based Flow Control (PFC)

The Priority-based Flow Control feature allows the switch to pause or inhibit transmission of individual priorities within a single physical link. By configuring PFC to pause a congested priority (priorities) independently, protocols that are highly loss sensitive can share the same link with traffic that has different loss tolerances. Priorities are differentiated by the priority field of the 802.1Q VLAN header. The PowerConnect 8000-series and 8100-series switches support lossless transport of frames on up to two priority classes.

NOTE: An interface that is configured for PFC is automatically disabled for 802.3x flow control.

For information about configuring the PFC feature, see "Configuring Data Center Bridging Features" on page 799.

Data Center Bridging Exchange (DBCx) Protocol

The Data Center Bridging Exchange Protocol (DCBx) is used by DCB devices to exchange configuration information with directly connected peers. The protocol is also used to detect misconfiguration of the peer DCB devices and, optionally, for configuration of peer DCB devices. For information about configuring DCBx settings, see "Configuring Data Center Bridging Features" on page 799. DCBx is a link-local protocol and operates only on individual links. When configuring FIP snooping on a port channel, ensure that all of the physical links in the port channel utilize the same auto configuration setting.

Enhanced Transmission Selection

Enhanced Transmission Selection (ETS) allows the switch to allocate bandwidth to traffic classes and share unused bandwidth with lower-priority traffic classes while coexisting with strict-priority traffic classes. ETS is supported on the PowerConnect 8000-series and 8100-series switches and can be configured manually or automatically using the auto configuration feature. For more information about ETS, see "Enhanced Transmission Selection" on page 819.

Fibre Channel over Ethernet (FCoE) Initialization Protocol Snooping

The FCoE Initialization Protocol (FIP) is used to perform the functions of FC_BB_E device discovery, initialization, and maintenance as defined in the ANSI T11 FC-BB-5 specification. The PC8024/PC8024F switch supports FIP snooping, which is a frame inspection method used by FIP Snooping Bridges to monitor FIP frames and apply policies based upon the L2 header information in those frames. For information about configuring the FIP Snooping feature, see "Configuring Data Center Bridging Features" on page 799.

Cisco Protocol Filtering

The Cisco Protocol Filtering feature (also known as Link Local Protocol Filtering) filters Cisco protocols that should not normally be relayed by a bridge. The group addresses of these Cisco protocols do not fall within the IEEE defined range of the 802.1D MAC Bridge Filtered MAC Group Addresses (01-80-C2-00-00-00 to 01-80-C2-00-00-0F).

For information about configuring LLPF, settings see "Configuring Port-Based Traffic Control" on page 665.

DHCP Layer 2 Relay

This feature permits Layer 3 Relay agent functionality in Layer 2 switched networks. The switch supports L2 DHCP relay configuration on individual ports, link aggregation groups (LAGs) and VLANs.

For information about configuring L2 DHCP Relay settings see "Configuring L2 and L3 Relay Features" on page 907.

Virtual Local Area Network Supported Features

For information about configuring VLAN features see "Configuring VLANs" on page 539.

VLAN Support

VLANs are collections of switching ports that comprise a single broadcast domain. Packets are classified as belonging to a VLAN based on either the VLAN tag or a combination of the ingress port and packet contents. Packets sharing common attributes can be groups in the same VLAN. The PowerConnect 8000-series and 8100-series switches are in full compliance with IEEE 802.1Q VLAN tagging.

Port-Based VLANs

Port-based VLANs classify incoming packets to VLANs based on their ingress port. When a port uses 802.1X port authentication, packets can be assigned to a VLAN based on the result of the 802.1X authentication a client uses when it accesses the switch. This feature is useful for assigning traffic to Guest VLANs or Voice VLANs.

IP Subnet-based VLAN

This feature allows incoming untagged packets to be assigned to a VLAN and traffic class based on the source IP address of the packet.

MAC-based VLAN

This feature allows incoming untagged packets to be assigned to a VLAN and traffic class based on the source MAC address of the packet.

IEEE 802.1v Protocol-Based VLANs

VLAN classification rules are defined on data-link layer (Layer 2) protocol identification. Protocol-based VLANs are used for isolating Layer 2 traffic for differing Layer 3 protocols.

GARP and GVRP Support

The switch supports the configuration of Generic Attribute Registration Protocol (GARP) timers GARP VLAN Registration Protocol (GVRP) relies on the services provided by GARP to provide IEEE 802.1Q-compliant VLAN pruning and dynamic VLAN creation on 802.1Q trunk ports. When GVRP is enabled, the switch registers and propagates VLAN membership on all ports that are part of the active spanning tree protocol topology.

For information about configuring GARP timers see "Configuring L2 Multicast Features" on page 681.

Voice VLAN

The Voice VLAN feature enables switch ports to carry voice traffic with defined priority. The priority level enables the separation of voice and data traffic coming onto the port. Voice VLAN is the preferred solution for enterprises wishing to deploy voice services in their network.

Guest VLAN

The Guest VLAN feature allows a switch to provide a distinguished service to unauthenticated users. This feature provides a mechanism to allow visitors and contractors to have network access to reach external network with no ability to browse information on the internal LAN.

For information about configuring the Guest VLAN see "Configuring Port and System Security" on page 457.

Double VLANs

The Double VLAN feature (IEEE 802.1QinQ) allows the use of a second tag on network traffic. The additional tag helps differentiate between customers in the Metropolitan Area Networks (MAN) while preserving individual customer's VLAN identification when they enter their own 802.1Q domain.

Spanning Tree Protocol Features

For information about configuring Spanning Tree Protocol features, see "Configuring the Spanning Tree Protocol" on page 605.

Spanning Tree Protocol (STP)

Spanning Tree Protocol (IEEE 802.1D) is a standard requirement of Layer 2 switches that allows bridges to automatically prevent and resolve L2 forwarding loops.

Spanning Tree Port Settings

The STP feature supports a variety of per-port settings including path cost, priority settings, Port Fast mode, STP Root Guard, Loop Guard, TCN Guard, and Auto Edge. These settings are also configurable per-LAG.

Rapid Spanning Tree

Rapid Spanning Tree Protocol (RSTP) detects and uses network topologies to enable faster spanning tree convergence after a topology change, without creating forwarding loops. The port settings supported by STP are also supported by RSTP.

Multiple Spanning Tree

Multiple Spanning Tree (MSTP) operation maps VLANs to spanning tree instances. Packets assigned to various VLANs are transmitted along different paths within MSTP Regions (MST Regions). Regions are one or more interconnected MSTP bridges with identical MSTP settings. The MSTP standard lets administrators assign VLAN traffic to unique paths.

The switch supports IEEE 802.1Q-2005, which is a version of corrects problems associated with the previous version, provides for faster transition-to-forwarding, and incorporates new features for a port (restricted role and restricted TCN).

Bridge Protocol Data Unit (BPDU) Guard

Spanning Tree BPDU Guard is used to disable the port in case a new device tries to enter the already existing topology of STP. Thus devices, which were originally not a part of STP, are not allowed to influence the STP topology.

BPDU Filtering

When spanning tree is disabled on a port, the BPDU Filtering feature allows BPDU packets received on that port to be dropped. Additionally, the BPDU Filtering feature prevents a port in Port Fast mode from sending and receiving BPDUs. A port in Port Fast mode is automatically placed in the forwarding state when the link is up to increase convergence time.

Link Aggregation Features

For information about configuring link aggregation (port-channel) features, see "Configuring Link Aggregation" on page 781.

Link Aggregation

Up to eight ports can combine to form a single Link Aggregated Group (LAG). This enables fault tolerance protection from physical link disruption, higher bandwidth connections and improved bandwidth granularity.

Per IEEE 802.1AX, only links with the same operational characteristics, such as speed and duplex setting, may be aggregated. PowerConnect switches aggregate links only if they have the same operational speed and duplex setting, as opposed to the configured speed and duplex setting. This allows operators to aggregate links that use auto negotiation to set values for speed and duplex. Dissimilar ports will not become active in the LAG if their operational settings do not match those of the first member of the LAG. PowerConnect switches also support setting the MTU on a LAG. When a link becomes active in a LAG, its MTU is dynamically changed to the LAG MTU. When the link leaves the LAG, its MTU reverts to the link setting.

Link Aggregate Control Protocol (LACP)

Link Aggregate Control Protocol (LACP) uses peer exchanges across links to determine, on an ongoing basis, the aggregation capability of various links, and continuously provides the maximum level of aggregation capability

achievable between a given pair of systems. LACP automatically determines, configures, binds, and monitors the binding of ports to aggregators within the system.

Routing Features

Address Resolution Protocol (ARP) Table Management

You can create static ARP entries and manage many settings for the dynamic ARP table, such as age time for entries, retries, and cache size.

For information about managing the ARP table, see "Configuring IP Routing" on page 883.

VLAN Routing

PowerConnect 8000-series and 8100-series switches support VLAN routing. You can also configure the software to allow traffic on a VLAN to be treated as if the VLAN were a router port.

For information about configuring VLAN routing interfaces, see "Configuring Routing Interfaces" on page 843.

IP Configuration

The switch IP configuration settings to allow you to configure network information for VLAN routing interfaces such as IP address and subnet mask, MTU size, and ICMP redirects. Global IP configuration settings for the switch allow you to enable or disable the generation of several types of ICMP messages and enable or disable the routing mode.

For information about managing global IP settings, see "Configuring IP Routing" on page 883.

Open Shortest Path First (OSPF)

Open Shortest Path First (OSPF) is a dynamic routing protocol commonly used within medium-to-large enterprise networks. OSPF is an interior gateway protocol (IGP) that operates within a single autonomous system.

For information about configuring OSPF, see "Configuring OSPF and OSPFv3" on page 931.

BOOTP/DHCP Relay Agent

The switch BootP/DHCP Relay Agent feature relays BootP and DHCP messages between DHCP clients and DHCP servers that are located in different IP subnets.

For information about configuring the BootP/DHCP Relay agent, see "Configuring L2 and L3 Relay Features" on page 907.

IP Helper and UDP Relay

The IP Helper and UDP Relay features provide the ability to relay various protocols to servers on a different subnet.

For information about configuring the IP helper and UDP relay features, see "Configuring L2 and L3 Relay Features" on page 907.

Routing Information Protocol

Routing Information Protocol (RIP), like OSPF, is an IGP used within an autonomous Internet system. RIP is an IGP that is designed to work with moderate-size networks.

For information about configuring RIP, see "Configuring RIP" on page 1019.

Router Discovery

For each interface, you can configure the Router Discovery Protocol (RDP) to transmit router advertisements. These advertisements inform hosts on the local network about the presence of the router.

For information about configuring router discovery, see "Configuring IP Routing" on page 883.

Routing Table

The routing table displays information about the routes that have been dynamically learned. You can configure static and default routes and route preferences. A separate table shows the routes that have been manually configured.

For information about viewing the routing table, see "Configuring IP Routing" on page 883.

Virtual Router Redundancy Protocol (VRRP)

VRRP provides hosts with redundant routers in the network topology without any need for the hosts to reconfigure or know that there are multiple routers. If the primary (master) router fails, a secondary router assumes control and continues to use the virtual router IP (VRIP) address.

VRRP Route Interface Tracking extends the capability of VRRP to allow tracking of specific route/interface IP states within the router that can alter the priority level of a virtual router for a VRRP group.

For information about configuring VRRP settings, see "Configuring VRRP" on page 1035.

Tunnel and Loopback Interfaces

PowerConnect 8000-series and 8100-series switches support the creation, deletion, and management of tunnel and loopback interfaces. Tunnel interfaces facilitate the transition of IPv4 networks to IPv6 networks. A loopback interface is always expected to be up, so you can configure a stable IP address that other network devices use to contact or identify the switch.

For information about configuring tunnel and loopback interfaces, see "Configuring Routing Interfaces" on page 843.

IPv6 Routing Features

IPv6 Configuration

The switch supports IPv6, the next generation of the Internet Protocol. You can globally enable IPv6 on the switch and configure settings such as the IPv6 hop limit and ICMPv6 rate limit error interval. You can also control whether IPv6 is enabled on a specific interface. The switch supports the configuration of many per-interface IPv6 settings including the IPv6 prefix and prefix length.

For information about configuring general IPv6 routing settings, see "Configuring IPv6 Routing" on page 1059.

IPv6 Routes

Because IPv4 and IPv6 can coexist on a network, the router on such a network needs to forward both traffic types. Given this coexistence, each switch maintains a separate routing table for IPv6 routes. The switch can forward IPv4 and IPv6 traffic over the same set of interfaces.

For information about configuring IPv6 routes, see "Configuring IPv6 Routing" on page 1059.

OSPFv3

OSPFv3 provides a routing protocol for IPv6 networking. OSPFv3 is a new routing component based on the OSPF version 2 component. In dual stack IPv6, you can configure and use both OSPF and OSPFv3 components.

For information about configuring OSPFv3, see "Configuring OSPF and OSPFv3" on page 931.

DHCPv6

DHCPv6 incorporates the notion of the “stateless” server, where DHCPv6 is not used for IP address assignment to a client, rather it only provides other networking information such as DNS, Network Time Protocol (NTP), and/or Session Initiation Protocol (SIP) information.

For information about configuring DHCPv6 settings, see "Configuring DHCPv6 Server and Relay Settings" on page 1083.

Quality of Service (QoS) Features

NOTE: Some features that can affect QoS, such as ACLs and Voice VLAN, are described in other sections within this chapter.

Differentiated Services (DiffServ)

The QoS Differentiated Services (DiffServ) feature allows traffic to be classified into streams and given certain QoS treatment in accordance with defined per-hop behaviors. PowerConnect 8000-series and 8100-series switches support both IPv4 and IPv6 packet classification.

For information about configuring DiffServ, see "Configuring Differentiated Services" on page 1103.

Class Of Service (CoS)

The Class Of Service (CoS) queueing feature lets you directly configure certain aspects of switch queuing. This provides the desired QoS behavior for different types of network traffic when the complexities of DiffServ are not required. CoS queue characteristics, such as minimum guaranteed bandwidth and transmission rate shaping, are configurable at the queue (or port) level.

For information about configuring CoS, see "Configuring Class-of-Service" on page 1133.

Auto Voice over IP (VoIP)

This feature provides ease of use for the user in setting up VoIP for IP phones on a switch. This is accomplished by enabling a VoIP profile that a user can select on a per port basis.

For information about configuring Auto VoIP, see "Configuring Auto VoIP" on page 1151.

Internet Small Computer System Interface (iSCSI) Optimization

The iSCSI Optimization feature helps network administrators track iSCSI traffic between iSCSI initiator and target systems. This is accomplished by monitoring, or snooping traffic to detect packets used by iSCSI stations in establishing iSCSI sessions and connections. Data from these exchanges may optionally be used to create classification rules to assign the traffic between the stations to a configured traffic class. This affects how the packets in the flow are queued and scheduled for egress on the destination port.

For information about configuring iSCSI settings, see "Configuring iSCSI Optimization" on page 387.

Layer 2 Multicast Features

For information about configuring L2 multicast features, see "Configuring L2 Multicast Features" on page 681.

MAC Multicast Support

Multicast service is a limited broadcast service that allows one-to-many and many-to-many connections. In Layer 2 multicast services, a single frame addressed to a specific multicast address is received, and copies of the frame to be transmitted on each relevant port are created.

IGMP Snooping

Internet Group Management Protocol (IGMP) Snooping is a feature that allows a switch to forward multicast traffic intelligently on the switch. Multicast IP traffic is traffic that is destined to a host group. Host groups are identified by class D IP addresses, which range from **224.0.0.0** to **239.255.255.255**. Based on the IGMP query and report messages, the switch forwards traffic only to the ports that request the multicast traffic. This prevents the switch from broadcasting the traffic to all ports and possibly affecting network performance.

IGMP Snooping Querier

When Protocol Independent Multicast (PIM) and IGMP are enabled in a network with IP multicast routing, the IP multicast router acts as the IGMP querier. However, if it is desirable to keep the multicast network Layer 2 switched only, the IGMP Snooping Querier can perform the query functions of a Layer 3 multicast router.

MLD Snooping

In IPv4, Layer 2 switches can use IGMP Snooping to limit the flooding of multicast traffic by dynamically configuring Layer 2 interfaces so that multicast traffic is forwarded to only those interfaces associated with IP multicast address.

In IPv6, MLD snooping performs a similar function. With MLD snooping, IPv6 multicast data is selectively forwarded to a list of ports intended to receive the data (instead of being flooded to all of the ports in a VLAN). This list is constructed by snooping IPv6 multicast control packets.

Multicast VLAN Registration

The Multicast VLAN Registration (MVR) protocol, like IGMP Snooping, allows a Layer 2 switch to listen to IGMP frames and forward the multicast traffic only to the receivers that request it. Unlike IGMP Snooping, MVR allows the switch to listen across different VLANs. MVR uses a dedicated VLAN, which is called the multicast VLAN, to forward multicast traffic over the Layer 2 network to the various VLANs that have multicast receivers as members.

Layer 3 Multicast Features

For information about configuring L3 multicast features, see "Managing IPv4 and IPv6 Multicast" on page 1157.

Distance Vector Multicast Routing Protocol

Distance Vector Multicast Routing Protocol (DVMRP) exchanges probe packets with all DVMRP-enabled routers, establishing two way neighboring relationships and building a neighbor table. It exchanges report packets and creates a unicast topology table, which is used to build the multicast routing table. This multicast route table is then used to route the multicast packets.

Internet Group Management Protocol

The Internet Group Management Protocol (IGMP) is used by IPv4 systems (hosts and routers) to report their IP multicast group memberships to any neighboring multicast routers. PowerConnect 8000-series and 8100-series switches perform the “multicast router part” of the IGMP protocol, which means it collects the membership information needed by the active multicast router.

IGMP Proxy

The IGMP Proxy feature allows the switch to act as a proxy for hosts by sending IGMP host messages on behalf of the hosts that the switch discovered through standard IGMP router interfaces.

Protocol Independent Multicast—Dense Mode

Protocol Independent Multicast (PIM) is a standard multicast routing protocol that provides scalable inter-domain multicast routing across the Internet, independent of the mechanisms provided by any particular unicast routing protocol. The Protocol Independent Multicast-Dense Mode (PIM-DM) protocol uses an existing Unicast routing table and a Join/Prune/Graft mechanism to build a tree. PIM-DM creates source-based shortest-path distribution trees, making use of reverse path forwarding (RPF).

Protocol Independent Multicast—Sparse Mode

Protocol Independent Multicast-Sparse Mode (PIM-SM) is used to efficiently route multicast traffic to multicast groups that may span wide area networks, and where bandwidth is a constraint. PIM-SM uses shared trees by default and implements source-based trees for efficiency. This data threshold rate is used to toggle between trees.

Protocol Independent Multicast—Source Specific Multicast

Protocol Independent Multicast—Source Specific Multicast (PIM-SSM) is a subset of PIM-SM and is used for one-to-many multicast routing applications, such as audio or video broadcasts. PIM-SSM does not use shared trees.

Protocol Independent Multicast IPv6 Support

PIM-DM and PIM-SM support IPv6 routes.

MLD/MLDv2 (RFC2710/RFC3810)

MLD is used by IPv6 systems (listeners and routers) to report their IP multicast addresses memberships to any neighboring multicast routers. The implementation of MLD v2 is backward compatible with MLD v1.

MLD protocol enables the IPv6 router to discover the presence of multicast listeners, the nodes that want to receive the multicast data packets, on its directly attached interfaces. The protocol specifically discovers which multicast addresses are of interest to its neighboring nodes and provides this information to the multicast routing protocol that make the decision on the flow of the multicast data packets.

3

Hardware Overview

This section provides an overview of the switch hardware. The topics covered in this section include:

- PowerConnect 8000-series and 8100-series Front Panel
- PowerConnect 8000-series and 8100-series Back Panel
- LED Definitions
- Switch Addresses

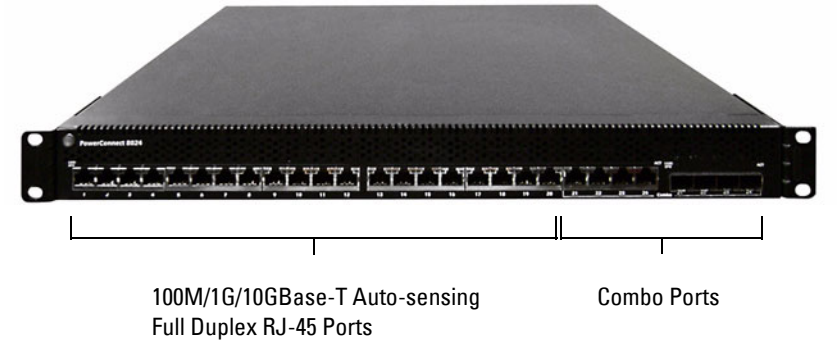
PowerConnect 8000-series and 8100-series Front Panel

The following sections describe the ports on the front panel of each switch.

PowerConnect 8024 Front Panel

The PowerConnect 8024 front panel provides 24 100M/1G/10GBase-T ports, four of which are combined with SFP/SFP+ ports.

Figure 3-1. PowerConnect 8024 Front Panel

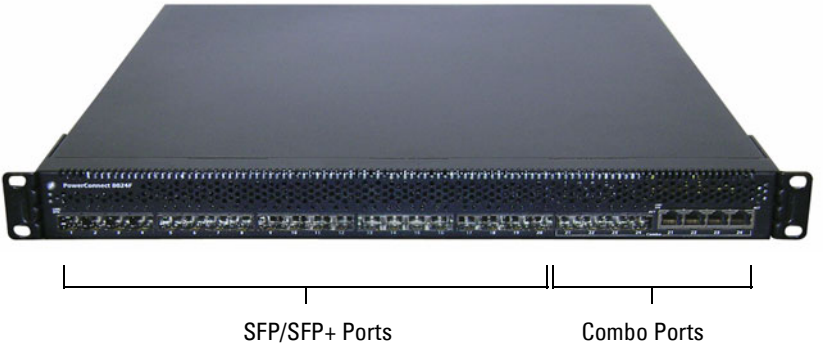


- The switch automatically detects crossed and straight-through cables on RJ-45 ports.
- RJ-45 ports support full-duplex mode 100/1000/10000 Mbps.
- PowerConnect 8024 switches can be stacked using the 10G SFP+ fiber ports. The 10G ports default to Ethernet mode and must be configured to be used as stacking ports.

PowerConnect 8024F Front Panel

The PowerConnect 8024F front panel provides 24 SFP/SFP+ ports, four of which are combined with 100M/1G/10GBase-T ports.

Figure 3-2. PowerConnect 8024F Front Panel



- The switch automatically detects crossed and straight-through cables on RJ-45 ports.
- SFP ports support both SX and LX modules.
- RJ-45 ports support full-duplex mode 100/1000/10000 Mbps.
- PowerConnect 8024F switches can be stacked with other PowerConnect 8024F switches using the 10G SFP+ fiber ports.

NOTE: A combo port may have both the RJ-45 and SFP+ ports cabled to a remote device and, if so, the SFP+ port will be the active port while the RJ-45 port will be disabled. The SFP+ ports comply with IEC60950-1, IEC60825-1, and IEC60825-2 and are contained within a fire enclosure.

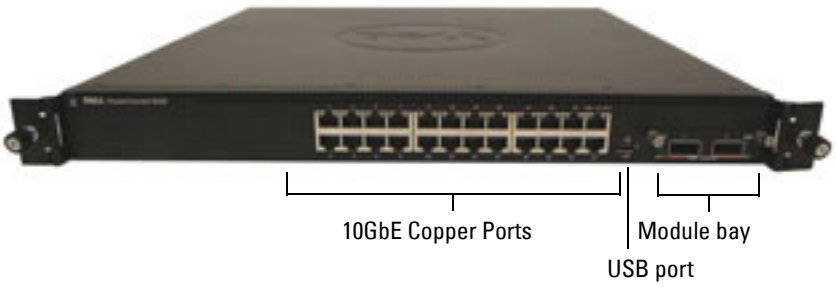
PowerConnect 8132 Front Panel

The PowerConnect 8132 front panel provides the following ports:

- 24 x 10GbE copper ports
- A USB port. See "USB Port (Power Connect 8100-series switches only)" on page 92.
- A module bay that supports the following modules:
  - 2 x 40 Gig QSFP (each QSFP may be configured as 4 x 10 Gig ports)
  - 4 x SFP+ module
  - 4 x 10GBaseT module

  See "Hot-Pluggable Interface Modules" on page 90 for more information.

Figure 3-3. PowerConnect 8132 Front Panel



PowerConnect 8132 switches can be stacked with other PowerConnect 81xx switches using 10G or 40G SFP+ or QSFP modules in the module bay.

PowerConnect 8132F Front Panel

The PowerConnect 8132F front panel provides the following ports:

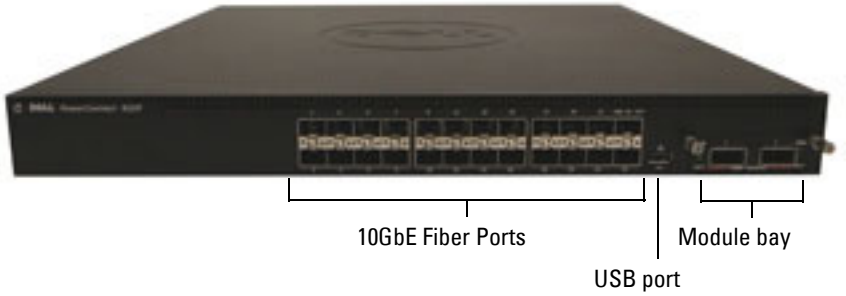
- 24 x 10GbE fiber ports
- A USB port. See "USB Port (Power Connect 8100-series switches only)" on page 92.
- A module bay that supports the following modules:
  - 2 x 40 Gig QSFP (each QSFP may be configured as 4 x 10 Gig ports)

– 4 x SFP+ module

– 4 x 10GBaseT module.

See "Hot-Pluggable Interface Modules" on page 90 for details about these modules.

Figure 3-4. PowerConnect 8132F Front Panel



PowerConnect 8132F switches can be stacked with other PowerConnect 81xx switches using 10G or 40G SFP+ or QSFP modules in the module bay.

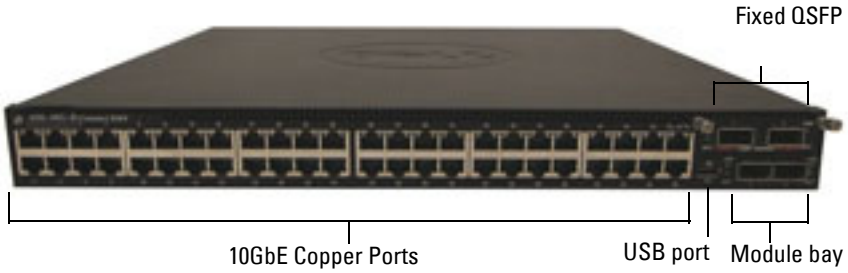
PowerConnect 8164 Front Panel

The PowerConnect 8164 front panel provides the following ports:

- 48 x 10GbE copper ports
- A USB port. See "USB Port (Power Connect 8100-series switches only)" on page 92.
- Two fixed QSFP ports, each supporting 4 x 10G or 1 x 40G connections
- One module bay that supports the following modules:
  - 2 x 40 Gig QSFP (each QSFP may be configured as 4 x 10 Gig ports)
  - 4 x SFP+ module
  - 4 x 10GBaseT module

  See "Hot-Pluggable Interface Modules" on page 90 for more information.

Figure 3-5. PowerConnect 8164 Front Panel



PowerConnect 8164 switches can be stacked with other PowerConnect 81xx switches using the 10G or 40G SFP+ or QSFP modules in the module bay or fixed QSFP ports.

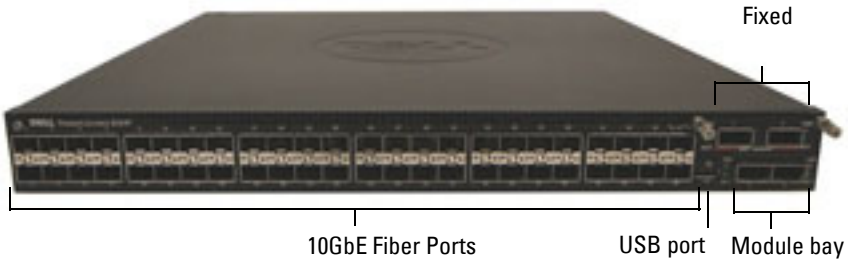
PowerConnect 8164F Front Panel

The PowerConnect 8164F front panel provides the following ports:

- 48 x 10GbE fiber ports
- A USB port. See "USB Port (Power Connect 8100-series switches only)" on page 92.
- Two fixed QSFP ports, each supporting 4 x 10G or 1 x 40G connections
- One module bay that supports the following modules:
  - 2 x 40 Gig QSFP (each QSFP may be configured as 4 x 10 Gig ports)
  - 4 x SFP+ module
  - 4 x 10GBaseT module

  See "Hot-Pluggable Interface Modules" on page 90 for more information.

Figure 3-6. PowerConnect 8164F Front Panel



PowerConnect 8164F switches can be stacked with other PowerConnect 81xx switches using the 10G or 40G SFP+ or QSFP modules in the module bay or fixed QSFP ports.

Hot-Pluggable Interface Modules

The PowerConnect 8132, 8132F, 8164, and 8164F switches support the following hot-pluggable interface modules:

- PC8100-QSFP - 2 x 40G QSFP port module - defaults to 2x40G
- PC8100-SFP+ - 4 x SFP+ port module - defaults to 4x10G mode
- PC8100-10GBT - 4 x 10GBase-T ports module - defaults to 4x10G mode
- Blank module - defaults to 10G mode

NOTE: The PowerConnect 8024 and 8024F switches do not support hot-swappable plug-in modules.

A reboot is necessary when a hot-pluggable module is replaced with a module of different type. Specifically, changing from a 40G module to a 10G module or from a 10G module to a 40G module requires a reboot.

You must execute a `no slot` or `clear config` command prior to inserting the new module. Note that changing the role of a port from stacking to Ethernet or vice-versa also requires a switch reboot.

If a `no slot` command is not issued prior to inserting a module, a message such as the following will appear:

```
Card Mismatch: Unit:1 Slot:1 Inserted-Card: Dell 2
Port QSFP Expansion Card Config-Card: Dell 4 Port
10GBase-T Expansion Card
```

The following sections provides details on each module.

Quad-Port SFP (QSFP) Uplink Module

The QSFP module supports features four ports that support 10G SFP+ transceivers. The QSFP module supports the following features:

- Four 10G ports with quad-breakout/QBO cable or one 40G port supporting CR4, SR4, and LR4 transceivers
- Front-panel port status LEDs

The QSFP interfaces can be used for stacking. Stacking is supported at distances of up to 100M.

Quad-Port SFP+ Uplink Module

The PC8100-SFP+ module features four SFP+ ports, each providing the following features:

- SFP+ SR, LR, and LRM optical interfaces
- SFP+ copper twinax interface
- Front-panel port status LEDs

The SFP+ connections can be used for stacking. Stacking is supported at distances of up to 100M.

10GBase-T Copper Uplink Module

The 10GBase-T copper module features four copper ports that can support 10GbE/1GbE/100MbE switching and provides following features:

- Complies with IEEE802.3z, IEEE 802.3, IEEE802.3u, IEEE802.3ab, IEEE802.3az, IEEE802.3an
- Four 10GBase-T/1GBase-T/100MBase-T copper ports.
- Front panel port status LEDs

USB Port (Power Connect 8100-series switches only)

The Type-A, female USB port supports a USB 2.0-compliant flash memory drive. The PowerConnect switch can read or write to a flash drive formatted as FAT-32. You can use a USB flash drive to copy switch configuration files and images between the USB flash drive and the switch. You can also use the USB flash drive to move and copy configuration files and images from one switch to other switches in the network.

The USB port does not support any other type of USB device.

Port and System LEDs

The front panel contains light emitting diodes (LEDs) to indicate port status.

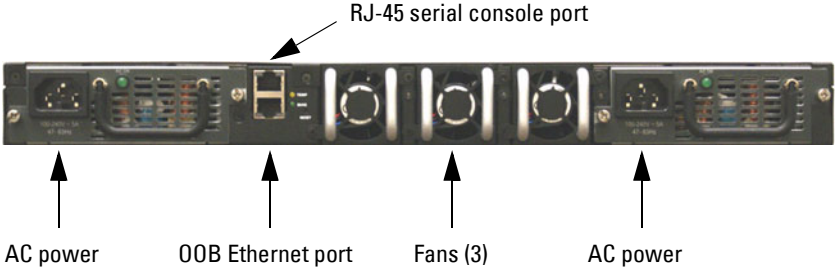
For information about the status that the LEDs indicate, see "LED Definitions" on page 95.

PowerConnect 8000-series and 8100-series Back Panel

- Console Port
- Out-of-Band Management Port
- Power Supplies
- Ventilation System

The following image show the back panel of the PowerConnect 8000-series and 8100-series switches.

Figure 3-7. PowerConnect 8000-series and 8100-series Rear Panel



Console Port

The console port is for management through a serial interface. This port provides a direct connection to the switch and allows you to access the CLI from a console terminal connected to the port through the provided serial cable (RJ-45 to female DB-9 connectors).

The console port supports asynchronous data of eight data bits, one stop bit, no parity bit, and no flow control. The default baud rate is 9600 bps.

Out-of-Band Management Port

The Out-of-Band (OOB) management port is a 10/100/1000BASE-T Ethernet port dedicated to remote switch management. Traffic on this port is segregated from operational network traffic on the switch ports and cannot be switched or routed to the operational network.

Power Supplies

Each PowerConnect 8000-series and 8100-series switch has two power supplies for redundant or loadsharing operation. Each power supply can support 300W.

CAUTION: Remove the power cable from the modules prior to removing the module itself. Power must not be connected prior to insertion in the chassis.

Ventilation System

The PowerConnect 8000-series and 8100-series switches have three removable FANs (see "PowerConnect 8000-series and 8100-series Rear Panel" on page 93), four Thermal sensors, and a FAN Speed Controller which can be used to control FAN speeds. You can verify operation by observing the LEDs.

LED Definitions

This section describes the LEDs on the front and back panels of the switch.

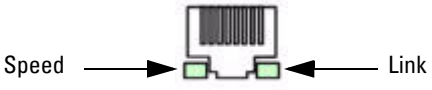
Port LEDs

Each port on a PowerConnect 8000-series and 8100-series switches includes two LEDs. One LED is on the left side of the port, and the second LED is on the right side of the port. This section describes the LEDs on the switch ports.

100/1000/10000Base-T Port LEDs

Each 100/1000/10000Base-T port has two LEDs. The following figure illustrates the 100/1000/10000Base-T port LEDs.

Figure 3-8. 100/1000/10000Base-T Port LEDs



The following table contains 100/1000/10000Base-T port LED definitions.

Table 3-1. 100/1000/10000Base-T Port Definitions

LED	Color	Definition
Speed	Green	The port is operating at 10 Gbps.
	Amber	The port is operating at 100/1000 Mbps.
Link	Blinking	Activity.

System LEDs

The system LEDs, located on the back panel, provide information about the power supplies, thermal conditions, and diagnostics.

Table 3-2 shows the System LED definitions for the 8000-series switches.

Table 3-2. System LED Definitions—PowerConnect 8000-Series Switches

LED	Color	Definition
DIAG	Flashing Green	A diagnostics test is in progress.
	Green	The diagnostics test was successfully completed.
	Red	The diagnostics test failed.
PWR	Green	Power Supply is operating correctly.
	Red	Power Supply has failed.
FAN	Green	Fans are operating correctly.
	Red	One or more fans have failed.
Temp	Amber	System temperature has exceeded threshold limit.

Table 3-3 shows the System LED definitions for the 8100-series switches.

Table 3-3. System LED Definitions—PowerConnect 8100-Series Switches

LED	Color	Definition
System	Blinking blue	The switch is booting
	Solid red	A critical system error has occurred.
	Blinking red	A noncritical system error occurred (fan or power supply failure).
Temp	Off	The switch is operating at normal temperature.
	Solid amber	The thermal sensor's system temperature threshold of 75°C has been exceeded.
Diag	Off	The switch is operating normally
	Blinking green	A diagnostic test is running.
Fan	Solid green	The fan is powered and is operating at the expected RPM.
	Solid red	A fan failure has occurred.
Stack	Solid green	The switch is in stacking master mode.
	Solid amber	The switch is in stacking slave mode.
	Off	The switch is in stand-alone mode.
Locator	Blinking green	The locator function is enabled.
	Solid green	The locator function is disabled.

Switch Addresses

The switch allocates MAC addresses from the Vital Product Data information stored locally in flash. MAC addresses are used as follows:

Table 3-4. MAC Address Use

Base	switch address
Base + 1	Out-of-band port
Base + 2	Layer 2
Base + 3	Layer 3

Shown below are three commands that display the MAC addresses used by the switch:

```
console#show system

System Description: Dell Ethernet Switch
System Up Time: 0 days, 00h:05m:11s
System Contact:
System Name:
System Location:
Burned In MAC Address: 001E.C9F0.004D
System Object ID: 1.3.6.1.4.1.674.10895.3042
System Model ID: PCT8132
Machine Type: PowerConnect 8132
Temperature Sensors:

Unit     Description       Temperature    Status
                           (Celsius)
----     -----------       -----------    ------
1        MAC               32             Good
1        CPU               31             Good
1        PHY (left side)   26             Good
1        PHY (right side)  29             Good

Fans:

Unit     Description    Status
----     -----------    ------
 1       Fan 1          OK
 1       Fan 2          OK
 1       Fan 3          OK
 1       Fan 4          OK
 1       Fan 5          OK
 1       Fan 6          No Power

Power Supplies:

Unit  Description    Status      Average     Current          Since
                                  Power       Power         Date/Time
                                  (Watts)     (Watts)
----  -----------  -----------  ----------  --------  -------------------
1     System       OK           42.0        43.4
1     Main         OK           N/A         N/A       04/06/2001 16:36:16
1     Secondary    No Power     N/A         N/A       01/01/1970 00:00:00

USB Port Power Status:
----------------------
Device Not Present
```

```
console#show ip interface out-of-band

IP Address..................................... 10.27.21.29
Subnet Mask.................................... 255.255.252.0
Default Gateway................................ 10.27.20.1
Configured IPv4 Protocol....................... DHCP
Burned In MAC Address.......................... 001E.C9F0.004E
```

```
console#show ip interface vlan 1

Routing Interface Status....................... Down
Primary IP Address............................. 1.1.1.2/255.255.255.0
Method......................................... Manual
Routing Mode................................... Enable
Administrative Mode............................ Enable
Forward Net Directed Broadcasts................ Disable
Proxy ARP...................................... Enable
Local Proxy ARP................................ Disable
Active State................................... Inactive
MAC Address.................................... 001E.C9F0.0050
Encapsulation Type............................. Ethernet
IP MTU......................................... 1500
Bandwidth...................................... 10000 kbps
Destination Unreachables....................... Enabled
ICMP Redirects................................. Enabled
```

4

Using Dell OpenManage Switch Administrator

This section describes how to use the Dell OpenManage Switch Administrator application. The topics covered in this section include:

- About Dell OpenManage Switch Administrator
- Starting the Application
- Understanding the Interface
- Using the Switch Administrator Buttons and Links
- Defining Fields

About Dell OpenManage Switch Administrator

Dell OpenManage Switch Administrator is a web-based tool to help you manage and monitor a PowerConnect 8000-series and 8100-series switches. Table 4-1 lists the web browsers that are compatible with Dell OpenManage Switch Administrator. The browsers have been tested on a PC running the Microsoft Windows operating system.

Table 4-1. Compatible Browsers

Browser	Version
Internet Explorer	v7–v8
Mozilla Firefox	v3.0 and higher
Safari	v4.0 and higher

NOTE: Additional operating systems and browsers might be compatible but have not been explicitly tested with Dell OpenManage Switch Administrator.

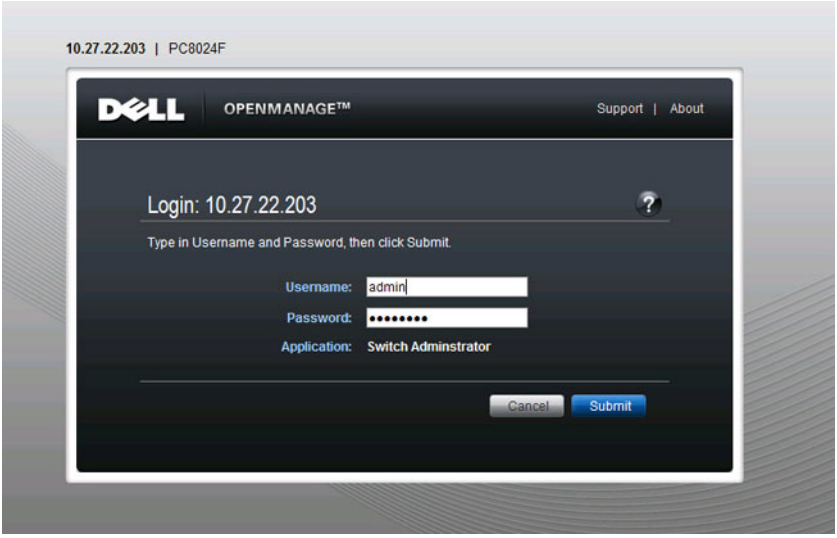
Starting the Application

To access the Dell OpenManage Switch Administrator and log on to the switch:

1. Open a web browser.
2. Enter the IP address of the switch in the address bar and press <Enter>.

   For information about assigning an IP address to a switch, see "Setting the IP Address and Other Basic Network Information" on page 119.
3. When the **Login** window displays, enter a user name and password. Passwords are both case sensitive and alpha-numeric.

Figure 4-1. Login Screen



NOTE: The switch is not configured with a default user name or password. You must connect to the CLI by using the console port to configure the initial user name and password. For information about connecting to the console, see "Console Connection" on page 107. For information about creating a user and password, see "Configuring Authentication, Authorization, and Accounting" on page 177.

4 Click **Submit**.

5 The **Dell OpenManage Switch Administrator** home page displays.

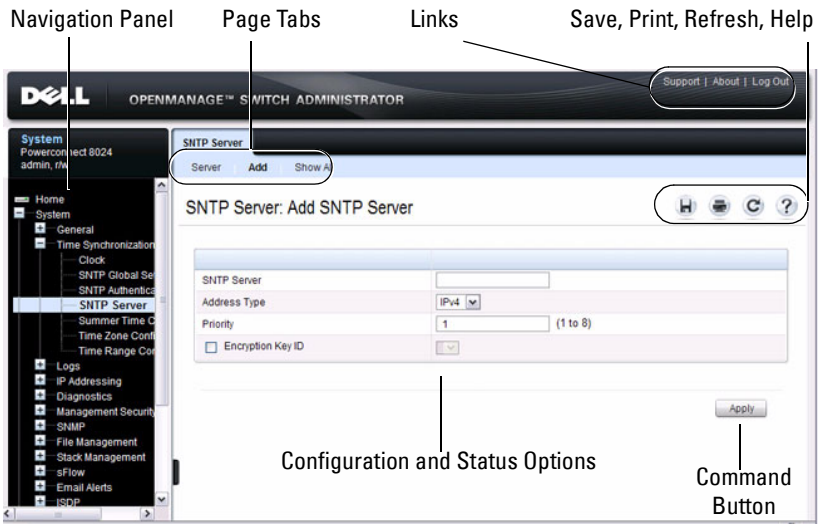
The home page is the **Device Information** page, which contains a graphical representation of the front panel of the switch. For more information about the home page, see "Device Information" on page 208.

Understanding the Interface

The Dell OpenManage Switch Administrator interface contains the following components:

- Navigation panel — Located on the left side of the page, the navigation pane provides an expandable view of features and their components.
- Configuration and status options — The main panel contains the fields you use to configure and monitor the switch.
- Page tabs — Some pages contain tabs that allow you to access additional pages related to the feature.
- Command buttons — Command buttons are located at the bottom of the page. Use the command buttons to submit changes, perform queries, or clear lists.
- Save, Print, Refresh, and Help buttons — These buttons appear on the top-right side of the main panel and are on every page.
- Support, About, and Logout links — These links appear at the top of every page.

Figure 4-2. Switch Administrator Components



Using the Switch Administrator Buttons and Links

Table 4-2 describes the buttons and links available from the **Dell OpenManage Switch Administrator** interface.

Table 4-2. Button and Link Descriptions

Button or Link	Description
Support	Opens the Dell Support page at support.dell.com
About	Contains the version and build number and Dell copyright information.
Log Out	Logs out of the application and returns to the login screen.
Save	Saves the running configuration to the startup configuration. When you click **Apply**, changes are saved to the running configuration. When the system boots, it loads the startup configuration. Any changes to the running configuration that were not saved to the startup configuration are lost across a power cycle.

Table 4-2. Button and Link Descriptions (Continued)

Button or Link	Description
Print	Opens the printer dialog box that allows you to print the current page. Only the main panel prints.
Refresh	Refreshes the screen with the current information.
Help	Online help that contains information to assist in configuring and managing the switch. The online help pages are context sensitive. For example, if the **IP Addressing** page is open, the help topic for that page displays if you click **Help**.
Apply	Updates the running configuration on the switch with the changes. Configuration changes take effect immediately.
Clear	Resets statistic counters and log files to the default configuration.
Query	Queries tables.
Left arrow and Right arrow	Moves information between lists.

NOTE: A few pages contain a button that occurs only on that page. Page-specific buttons are described in the sections that pertain to those pages.

Defining Fields

User-defined fields can contain 1–159 characters, unless otherwise noted on the Dell OpenManage Switch Administrator web page.

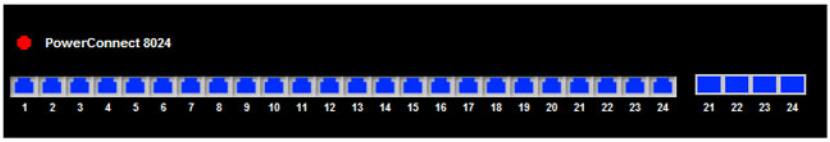
All characters may be used except for the following:

- \
- /
- :
- *
- ?
- <
- >
- |

Understanding the Device View

The Device View is a Java applet that shows various information about switch. This graphic appears on the OpenManage Switch Administrator **Home** page, which is the page that displays after a successful login. The graphic provides information about switch ports and system health.

Figure 4-3. PowerConnect 8024 Device View



Using the Device View Port Features

The switching-port coloring indicates if a port is currently active. Green indicates that the port has a link, red indicates that an error has occurred on the port, and blue indicates that the link is down. Each port image is a hyperlink to the **Port Configuration** page for the specific port.

5

Using the Command-Line Interface

This section describes how to use the Command-Line Interface (CLI) on a PowerConnect 8000-series and 8100-series switches.

The topics covered in this section include:

- Accessing the Switch Through the CLI
- Understanding Command Modes
- Entering CLI Commands

Accessing the Switch Through the CLI

The CLI provides a text-based way to manage and monitor the PowerConnect 8000-series and 8100-series switches. You can access the CLI by using a direct connection to the console port or by using a Telnet or SSH client.

To access the switch by using Telnet or Secure Shell (SSH), the switch must have an IP address, and the management station you use to access the device must be able to ping the switch IP address.

For information about assigning an IP address to a switch, see "Setting the IP Address and Other Basic Network Information" on page 119.

Console Connection

Use the following procedures to connect to the CLI by connecting to the console port. For more information about creating a serial connection, see the *Getting Started Guide* available at support.dell.com/manuals.

1. Connect the DB-9 connector of the supplied serial cable to a management station, and connect the RJ-45 connector to the switch console port.

   The console port is located on the back panel above the OOB Ethernet port.

2 Start the terminal emulator, such as Microsoft HyperTerminal, and select the appropriate serial port (for example, COM 1) to connect to the console.

3 Configure the management station serial port with the following settings:

- Data rate — 9600 baud.
- Data format — 8 data bits
- Parity — None
- Stop bits — 1
- Flow control — None

4 Power on the switch.

After the boot process completes, the `console>` prompt displays, and you can enter commands.

NOTE: By default, no authentication is required for console access. However, if an authentication method has been configured for console port access, the User: login prompt displays.

Telnet Connection

Telnet is a terminal emulation TCP/IP protocol. ASCII terminals can be virtually connected to the local device through a TCP/IP protocol network.

Telnet connections are enabled by default, and the Telnet port number is 23. The switch supports up to four simultaneous Telnet sessions. All CLI commands can be used over a Telnet session.

NOTE: SSH, which is more secure than Telnet, is disabled by default.

To connect to the switch by using Telnet, the switch must have an IP address, and the switch and management station must have network connectivity. You can use any Telnet client on the management station to connect to the switch.

You can also initiate a Telnet session from the OpenManage Switch Administrator. For more information, see "Initiating a Telnet Session from the Web Interface" on page 242.

Understanding Command Modes

The CLI groups commands into modes according to the command function. Each of the command modes supports specific software commands. The commands in one mode are not available until you switch to that particular mode, with the exception of the User EXEC mode commands. You can execute the User EXEC mode commands in the Privileged EXEC mode.

To display the commands available in the current mode, enter a question mark (?) at the command prompt. In each mode, a specific command is used to navigate from one command mode to another.

The main command modes include the following:

- User EXEC — Commands in this mode permit connecting to remote devices, changing terminal settings on a temporary basis, performing basic tests, and listing system information.
- Privileged EXEC — Commands in this mode permit you to view all switch settings and to enter the global configuration mode.
- Global Configuration — Commands in this mode manage the device configuration on a global level and apply to system features, rather than to a specific protocol or interface.
- Interface Configuration — Commands in this mode configure the settings for a specific interface or range of interfaces.
- VLAN Configuration — Commands in this mode create and remove VLANs and configure IGMP/MLD Snooping parameters for VLANs.

The CLI includes several additional command modes. For more information about the CLI command modes, including details about all modes, see the *CLI Reference Guide*.

Table 5-1 describes how to navigate between CLI Command Mode and lists the prompt that displays in each mode.

Table 5-1. Command Mode Overview

Command Mode	Access Method	Command Prompt	Exit or Access Previous Mode
User EXEC	The user is automatically in User EXEC mode unless the user is defined as a privileged user.	`console>`	logout
Privileged EXEC	From User EXEC mode, enter the **enable** command	`console#`	Use the **exit** command, or press Ctrl-Z to return to User EXEC mode.
Global Configuration	From Privileged EXEC mode, use the **configure** command.	`console(config)#`	Use the **exit** command, or press Ctrl-Z to return to Privileged EXEC mode.
Interface Configuration	From Global Configuration mode, use the **interface** command and specify the interface type and ID.	`console(config-if)#`	To exit to Global Configuration mode, use the **exit** command, or press Ctrl-Z to return to Privileged EXEC mode.
VLAN Config	From Global Configuration mode, use the **vlan database** command.	`console(config-vlan)#`	To exit to Global Configuration mode, use the **exit** command, or press Ctrl-Z to return to Privileged EXEC mode.

Entering CLI Commands

The switch CLI uses several techniques to help you enter commands.

Using the Question Mark to Get Help

Enter a question mark (?) at the command prompt to display the commands available in the current mode.

```
console(config-vlan)#?

exit            To exit from the mode.
help            Display help for various special keys.
ip              Configure IP parameters.
ipv6            Configure IPv6 parameters.
protocol        Configure the Protocols associated with
                particular Group Ids.
vlan            Create a new VLAN or delete an existing
                VLAN.
```

Enter a question mark (?) after each word you enter to display available command keywords or parameters.

```
console(config)#vlan ?

database    Type 'vlan database' to enter VLAN mode.
protocol    Configure Protocol Based VLAN parameters.
```

If the help output shows a parameter in angle brackets, you must replace the parameter with a value.

```
console#telnet ?

<ip-address|hostname>     Enter the valid host IP
                          address or Host Name.
```

If there are no additional command keywords or parameters, or if additional parameters are optional, the following message appears in the output:

```
<cr>            Press enter to execute the command.
```

You can also enter a question mark (?) after typing one or more characters of a word to list the available command or parameters that begin with the letters, as shown in the following example:

```
console#show po?

policy-map                port                          ports
```

Using Command Completion

The CLI can complete partially entered commands when you press the <Tab> or <Space> key.

```
console#show run<Tab>
console#show running-config
```

If the characters you entered are not enough for the switch to identify a single matching command, continue entering characters until the switch can uniquely identify the command. Use the question mark (?) to display the available commands matching the characters already entered.

Entering Abbreviated Commands

To execute a command, you need to enter enough characters so that the switch can uniquely identify a command. For example, to enter Global Configuration mode from Privileged EXEC mode, you can enter **con** instead of **configure**.

```
console#con

console(config)#
```

Negating Commands

For many commands, the prefix keyword **no** is entered to cancel the effect of a command or reset the configuration to the default value. All configuration commands have this capability.

Understanding Error Messages

If you enter a command and the system is unable to execute it, an error message appears. Table 5-2 describes the most common CLI error messages.

Table 5-2. CLI Error Messages

Message Text	Description
`% Invalid input detected at '^' marker.`	Indicates that you entered an incorrect or unavailable command. The carat (^) shows where the invalid text is detected. This message also appears if any of the parameters or values are not recognized.
`Command not found / Incomplete command. Use ? to list commands.`	Indicates that you did not enter the required keywords or values.
`Ambiguous command`	Indicates that you did not enter enough letters to uniquely identify the command.

If you attempt to execute a command and receive an error message, use the question mark (?) to help you determine the possible keywords or parameters that are available.

Recalling Commands from the History Buffer

Every time a command is entered in the CLI, it is recorded in an internally managed Command History buffer. By default, the history buffer is enabled and stores the last 10 commands entered. These commands can be recalled, reviewed, modified, and reissued. This buffer is not preserved after switch resets.

Table 5-3. History Buffer Navigation

Keyword	Source or Destination
Up-arrow key <Ctrl>+<P>	Recalls commands in the history buffer, beginning with the most recent command. Repeats the key sequence to recall successively older commands.
Down-arrow key <Ctrl>+<N>	Returns to more recent commands in the history buffer after recalling commands with the up-arrow key. Repeating the key sequence recalls more recent commands in succession.

6

Default Settings

This section describes the default settings for many of the software features on the PowerConnect 8000-series and 8100-series switches.

Table 6-1. Default Settings

Feature	Default
IP address	None
Subnet mask	None
Default gateway	None
DHCP client	Enabled on out-of-band (OOB) interface.
VLAN 1 Members	All switch ports
SDM template	Dual IPv4 and IPv6 routing
Users	None
Minimum password length	8 characters
IPv6 management mode	Enabled
SNTP client	Disabled
Global logging	Enabled
Switch auditing	Disabled
CLI command logging	Disabled
Web logging	Disabled
SNMP logging	Disabled
Console logging	Enabled (Severity level: debug and above)
RAM logging	Enabled (Severity level: debug and above)
Persistent (FLASH) logging	Disabled
DNS	Enabled (No servers configured)
SNMP	Enabled (SNMPv1)

Table 6-1. Default Settings (Continued)

Feature	Default
SNMP Traps	Enabled
Auto Configuration	Enabled
Auto Save	Disabled
Stacking	Enabled
Nonstop Forwarding on the Stack	Enabled
sFlow	Enabled
ISDP	Enabled (Versions 1 and 2)
RMON	Enabled
TACACS+	Not configured
RADIUS	Not configured
SSH/SSL	Disabled
Telnet	Enabled
Denial of Service Protection	Disabled
Captive Portal	Disabled
Dot1x Authentication (IEEE 802.1X)	Disabled
MAC-Based Port Security	All ports are unlocked
Access Control Lists (ACL)	None configured
IP Source Guard (IPSG)	Disabled
DHCP Snooping	Disabled
Dynamic ARP Inspection	Disabled
Protected Ports (Private VLAN Edge)	None
Flow Control Support (IEEE 802.3x)	Enabled
Head of Line Blocking Prevention	Disabled
Maximum Frame Size	1500 bytes
Auto-MDI/MDIX Support	Enabled
Auto Negotiation	Enabled
Advertised Port Speed	Maximum Capacity

Table 6-1. Default Settings (Continued)

Feature	Default
Broadcast Storm Control	Disabled
Port Mirroring	Disabled
LLDP	Enabled
LLDP-MED	Disabled
MAC Table Address Aging	300 seconds (Dynamic Addresses)
Cisco Protocol Filtering (LLPF)	No protocols are blocked
DHCP Layer 2 Relay	Disabled
Default VLAN ID	1
Default VLAN Name	Default
GVRP	Disabled
GARP Timers	Leave: 60 centiseconds Leave All: 1000 centiseconds Join: 20 centiseconds
Voice VLAN	Disabled
Guest VLAN	Disabled
RADIUS-assigned VLANs	Disabled
Double VLANs	Disabled
Spanning Tree Protocol (STP)	Enabled
STP Operation Mode	IEEE 802.1w Rapid Spanning Tree
Optional STP Features	Disabled
STP Bridge Priority	32768
Multiple Spanning Tree	Disabled
Link Aggregation	No LAGs configured
LACP System Priority	1
Routing Mode	Disabled
OSPF Admin Mode	Enabled
OSPF Router ID	0.0.0.0

Table 6-1. Default Settings (Continued)

Feature	Default
IP Helper and UDP Relay	Enabled
RIP	Enabled
VRRP	Disabled
Tunnel and Loopback Interfaces	None
IPv6 Routing	Disabled
DHCPv6	Disabled
OSPFv3	Enabled
DiffServ	Enabled
Auto VoIP	Disabled
Auto VoIP Traffic Class	6
PFC	Disabled; no classifications configured.
DCBx version	Auto detect
FIP snooping	Disabled globally and on all VLANs
iSCSI	Enabled
Bridge Multicast Filtering	Disabled
MLD Snooping	Disabled
IGMP Snooping	Disabled
IGMP Snooping Querier	Disabled
GMRP	Disabled
IPv4 Multicast	Disabled
IPv6 Multicast	Disabled

7

Setting the IP Address and Other Basic Network Information

This chapter describes how to configure basic network information for the switch, such as the IP address, subnet mask, and default gateway. The topics in this chapter include:

- IP Address and Network Information Overview
- Default Network Information
- Configuring Basic Network Information (Web)
- Configuring Basic Network Information (CLI)
- Basic Network Information Configuration Example

IP Address and Network Information Overview

What Is the Basic Network Information?

The basic network information includes settings that define the PowerConnect 8000-series and 8100-series switches in relation to the network. Table 7-1 provides an overview of the settings this chapter describes.

Table 7-1. Basic Network Information

Feature	Description
IP Address	On an IPv4 network, the a 32-bit number that uniquely identifies a host on the network. The address is expressed in dotted-decimal format, for example 192.168.10.1.
Subnet Mask	Determines which bits in the IP address identify the network, and which bits identify the host. Subnet masks are also expressed in dotted-decimal format, for example 255.255.255.0.

Table 7-1. Basic Network Information (Continued)

Feature	Description
Default Gateway	Typically a router interface that is directly connected to the switch and is in the same subnet. The switch sends IP packets to the default gateway when it does not recognize the destination IP address in a packet.
DHCP Client	Requests network information from a DHCP server on the network.
Domain Name System (DNS) Server	Translates hostnames into IP addresses. The server maintains a domain name databases and their corresponding IP addresses.
Default Domain Name	Identifies your network, such as dell.com. If you enter a hostname and do not include the domain name information, the default domain name is automatically appended to the hostname.
Host Name Mapping	Allows you to statically map an IP address to a hostname.

Additionally, this chapter describes how to view host name-to-IP address mappings that have been dynamically learned by the system.

Why Is Basic Network Information Needed?

PowerConnect 8000-series and 8100-series switches are layer 2/3 managed switches. To manage the switch remotely by using a web browser or Telnet client, the switch must have an IP address, subnet mask, and default gateway. You must also configure a username and password to be able to log into the switch from a remote host. For information about configuring users, see "Configuring Authentication, Authorization, and Accounting" on page 177. If you manage the switch only by using a console connection, configuring an IP address and user is not required.

NOTE: The configuration example in this chapter includes commands to create an administrative user with read/write access.

Configuring the DNS information, default domain name, and host name mapping help the switch identify and locate other devices on the network and on the Internet. For example, to upgrade the switch software by using a TFTP

server on the network, you must identify the TFTP server. If you configure the switch to use a DNS server to resolve hostnames into IP addresses, you can enter the hostname of the TFTP server instead of the IP address. It is often easier to remember a hostname than an IP address, and if the IP address is dynamically assigned, it might change from time-to-time.

How Is Basic Network Information Configured?

You must use a console-port connection to perform the initial switch configuration. When you boot the switch for the first time and the configuration file is empty, the Dell Easy Setup Wizard starts. The Dell Easy Setup Wizard is a CLI-based tool to help you perform the initial switch configuration. If you do not respond to the Dell Easy Setup Wizard prompt within 60 seconds, the `console>` prompt appears, and you enter User Configuration mode.

For more information about performing the initial switch configuration by using the wizard, see the *Getting Started Guide* at support.dell.com/manuals.

If you do not use the wizard to prompt you for the initial configuration information, you can enable the DHCP client on the switch to obtain network information from a DHCP server on your network, or you can statically assign the network information.

After you configure the switch with an IP address and create a user account, you can continue to use the console connection to configure basic network information, or you can log on to the switch by using a Telnet client or a web browser. You can change the IP address information and configure additional network information from the remote system.

What Is Out-of-Band Management and In-Band Management?

PowerConnect 8000-series and 8100-series switches have an external port intended solely for management of the switch. This port is the out-of-band (OOB) management port. Traffic received on the OOB port is never switched or routed to any in-band port. Likewise, traffic received on any in-band port is never forwarded or routed over the OOB port. The only applications available on the OOB port are protocols required to manage the switch, for example telnet, SSH, DHCP client, and TFTP.

Alternatively, network administrators may choose to manage their network via the production network. This is in-band management. Because in-band management traffic is mixed in with production network traffic, it is subject to all of the filtering rules usually applied on a switched/routed port such as ACLs and VLAN tagging.

You can assign an IP address to OOB management port and to any VLAN. By default, all ports are members of VLAN 1. If you assign an IP address to VLAN 1, you can connect to the switch management interface by using any of the switch ports.

Dell recommends that you use the OOB port for remote management. The following list highlights some advantages of using OOB management instead of in-band management:

- Traffic on the OOB port is segregated from traffic on the production network, so you can keep the management traffic and network traffic separate.
- If the production network is experiencing problems, you can still access the switch management interface and troubleshoot issues.
- Because the OOB port is intended to be physically isolated from the production network, configuration options are limited to just those protocols needed to manage the switch. Limiting the configuration options makes it difficult to accidentally cut off management access to the switch.

DHCP can be enabled on the OOB interface and all VLAN interfaces simultaneously, or you can configure static information. To configure static address information on the default VLAN, set the IP address and subnet mask on the VLAN interface and configure a global default gateway for the switch.

Adjusting the Management Interface MTU

When logging in to the PowerConnect switch using TCP, the switch negotiates the TCP Maximum Segment Size (MSS) using the minimum of the requested MSS or the MTU setting of the port. TCP packets are transmitted from the switch with the DF (Don't Fragment) bit set in order to receive notification of fragmentation from any transit routers. Upon receiving an ICMP *Destination Unreachable, Fragmentation needed but DF set*

notification, the switch will reduce the MSS. However, many firewalls block ICMP Destination Unreachable messages, which causes the destination to request the packet again until the connection times out.

In order to resolve this issue, you can reduce the MSS setting to a more appropriate value on the local host or alternatively, you can set the MTU on the PowerConnect management port to a smaller value.

Default Network Information

By default, no network information is configured. The DHCP client is enabled on the OOB interface by default. DNS is enabled, but no DNS servers are configured.

Configuring Basic Network Information (Web)

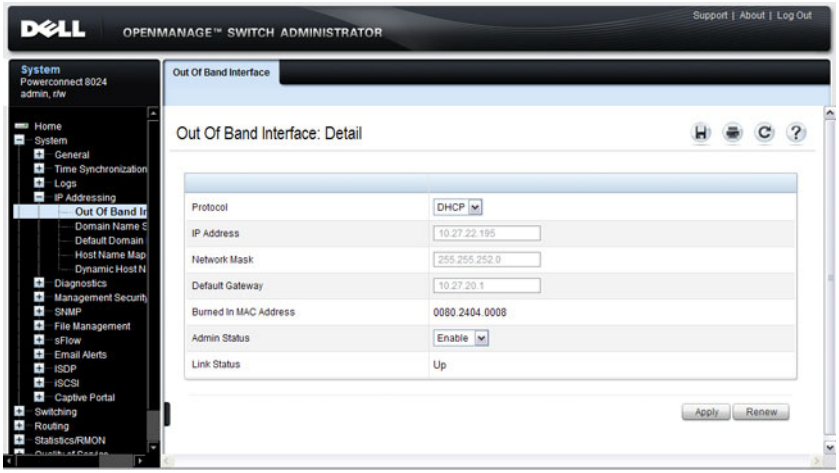
This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring basic network information on the PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

Out-of-Band Interface

Use the **Out of Band Interface** page to assign the Out of Band Interface IP address and subnet mask or to enable/disable the DHCP client for address information assignment. DHCP is enabled by default on the OOB interface.

To display the **Out of Band Interface** page, click **System** → **IP Addressing** → **Out of Band Interface** in the navigation panel.

Figure 7-1. Out of Band Interface



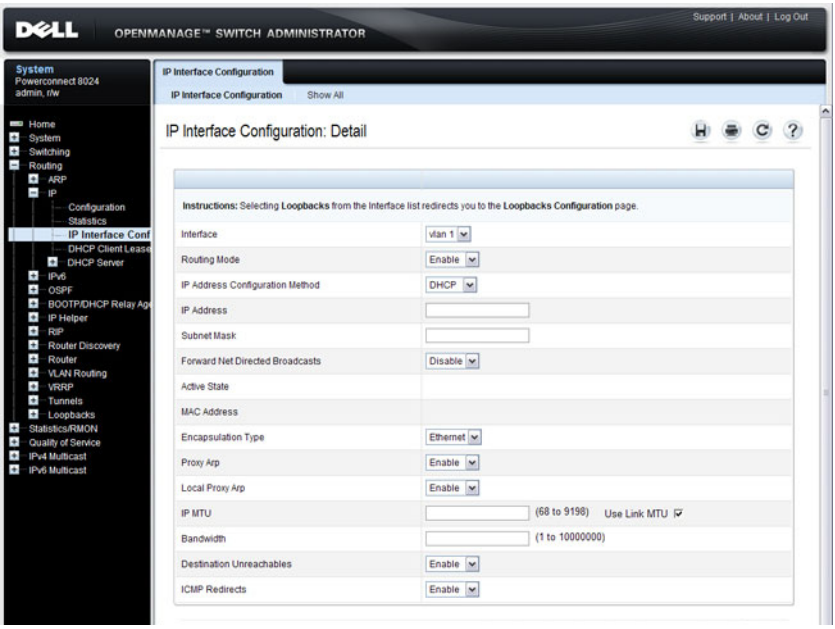
To enable the DHCP client and allow a DHCP server on your network to automatically assign the network information to the OOB interface, select DHCP from the **Protocol** menu. If you statically assign the network information, make sure the **Protocol** menu is set to None.

IP Interface Configuration (Default VLAN IP Address)

Use the **IP Interface Configuration** page to assign the Default VLAN IP address and Subnet Mask, the Default Gateway IP address, and to assign the boot protocol.

To display the **IP Interface Configuration** page, click **Routing** → **IP** → **IP Interface Configuration** in the navigation panel.

Figure 7-2. IP Interface Configuration (Default VLAN)



Assigning Network Information to the Default VLAN

To assign an IP Address and subnet mask to the default VLAN:

1. From the **Interface** menu, select VLAN 1.
2. From the **Routing Mode** field, select **Enable**.
3. From the **IP Address Configuration Method** field specify whether to assign a static IP address (Manual) or use DHCP for automatic address assignment.

4 If you select **Manual** for the configuration method, specify the **IP Address** and **Subnet Mask** in the appropriate fields.

5 Click **Apply**.

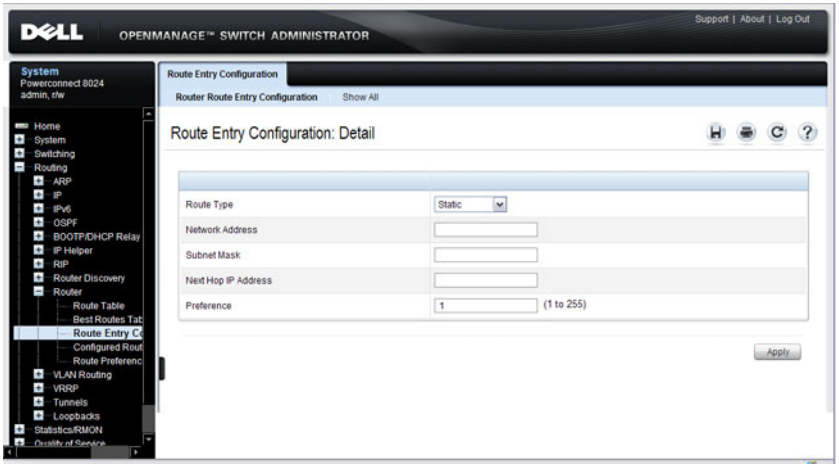
NOTE: You do not need to configure any additional fields on the page. For information about VLAN routing interfaces, see "Configuring Routing Interfaces" on page 843.

Route Entry Configuration (Switch Default Gateway)

Use the **Route Entry Configuration** page to configure the default gateway for the switch. The Default VLAN uses the switch default gateway as its default gateway.

To display the **Route Entry Configuration** page, click **Routing** → **Router** → **Route Entry Configuration** in the navigation panel.

Figure 7-3. Route Configuration (Default VLAN)

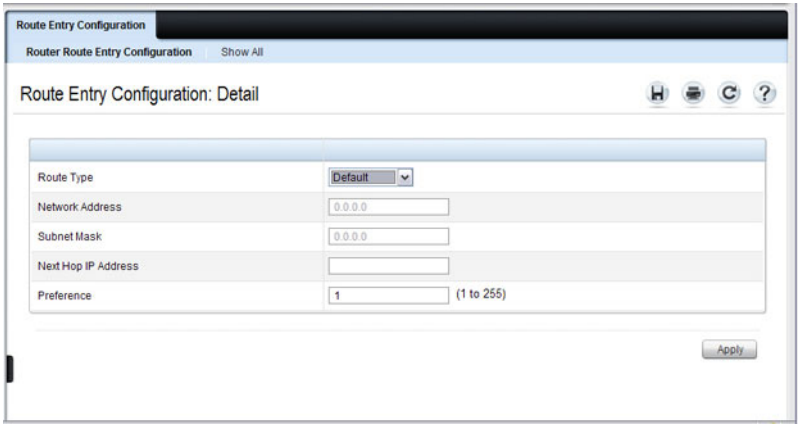


Configuring a Default Gateway for the Switch:

To configure the switch default gateway:

1. Open the **Route Entry Configuration** page.
2. From the **Route Type** field, select Default.

Figure 7-4. Default Route Configuration (Default VLAN)



3. In the **Next Hop IP Address** field, enter the IP address of the default gateway.
4. Click **Apply**.

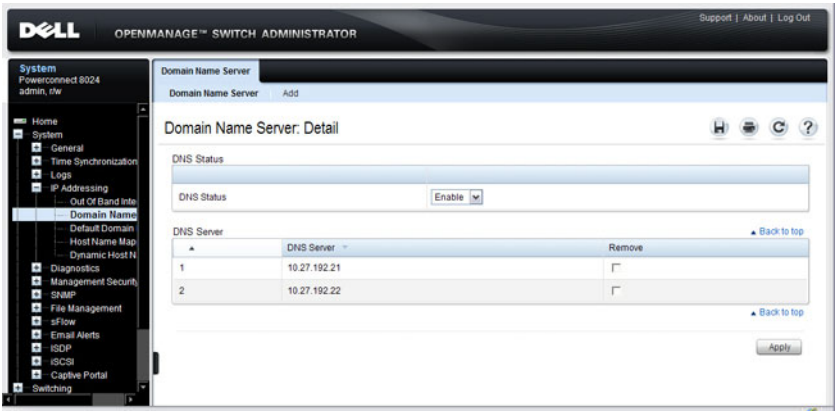
For more information about configuring routes, see "Configuring IP Routing" on page 883.

Domain Name Server

Use the **Domain Name Server** page to configure the IP address of the DNS server. The switch uses the DNS server to translate hostnames into IP addresses.

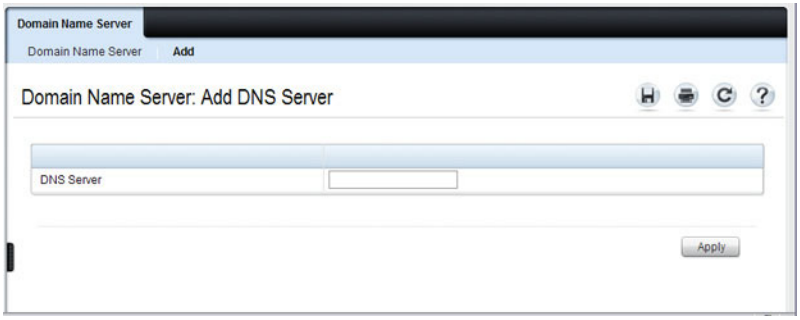
To display the **Domain Name Server** page, click **System** → **IP Addressing** → **Domain Name Server** in the navigation panel.

Figure 7-5. DNS Server



To configure DNS server information, click the **Add** link and enter the IP address of the DNS server in the available field.

Figure 7-6. Add DNS Server

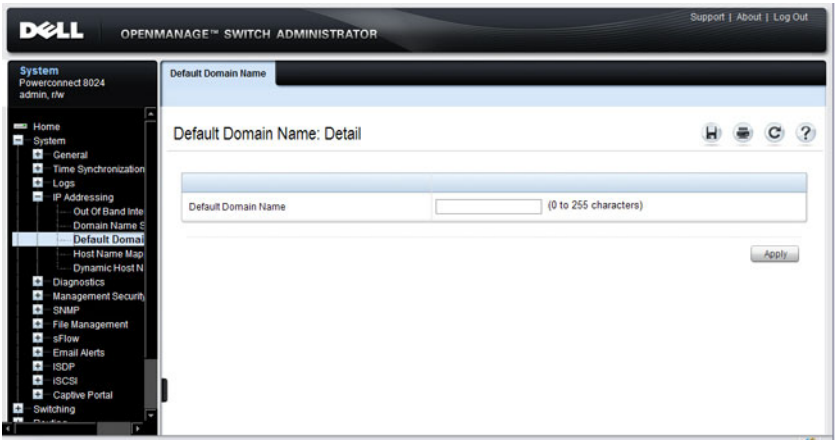


Default Domain Name

Use the **Default Domain Name** page to configure the domain name the switch adds to a local (unqualified) hostname.

To display the **Default Domain Name** page, click **System** → **IP Addressing** → **Default Domain Name** in the navigation panel.

Figure 7-7. Default Domain Name

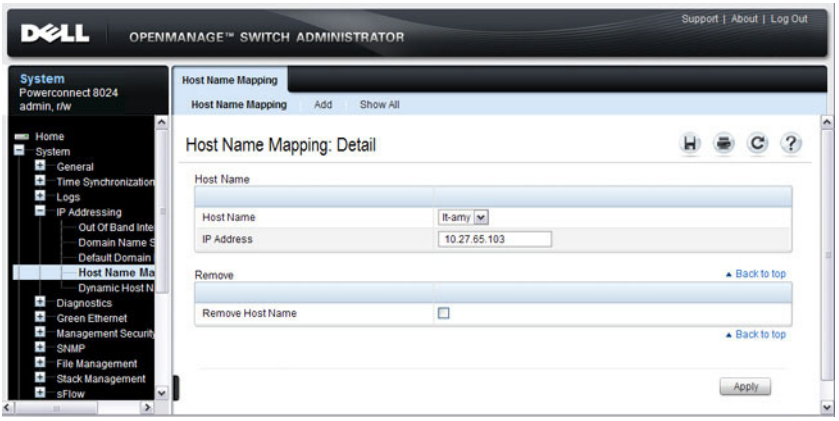


Host Name Mapping

Use the **Host Name Mapping** page to assign an IP address to a static host name. The **Host Name Mapping** page provides one IP address per host.

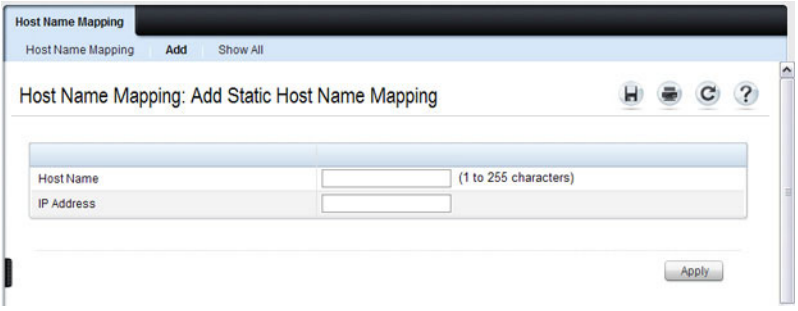
To display the **Host Name Mapping** page, click **System** → **IP Addressing** → **Host Name Mapping**.

Figure 7-8. Host Name Mapping



To map a host name to an IP address, click the **Add** link, type the name of the host and its IP address in the appropriate fields, and then click **Apply**.

Figure 7-9. Add Static Host Name Mapping



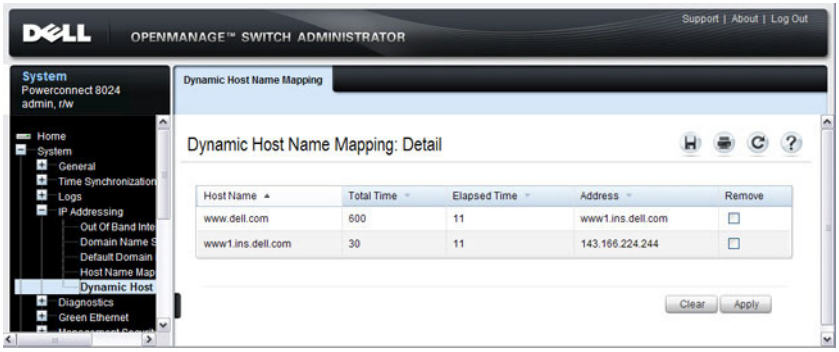
Use the **Show All** link to view all configured host name-to-IP address mappings.

Dynamic Host Name Mapping

Use the **Dynamic Host Name Mapping** page to view dynamic host entries the switch has learned. The switch learns hosts dynamically by using the configured DNS server to resolve a hostname. For example, if you ping **www.dell.com** from the CLI, the switch uses the DNS server to lookup the IP address of **dell.com** and adds the entry to the Dynamic Host Name Mapping table.

To display the **Dynamic Host Name Mapping** page, click **System → IP Addressing → Dynamic Host Name Mapping** in the navigation panel.

Figure 7-10. View Dynamic Host Name Mapping



Configuring Basic Network Information (CLI)

This section provides information about the commands you use to configure basic network information on the PowerConnect 8000-series and 8100-series switches. For more information about these commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Enabling the DHCP Client on the OOB Port

Beginning in Privileged EXEC mode, use the following commands to enable the DHCP client on the OOB port.

Command	Purpose
configure	Enter Global Configuration mode.
interface out-of-band	Enter Interface Configuration mode for the OOB port.
ip address dhcp	Enable the DHCP client.
CTRL + Z	Exit to Privileged EXEC mode.
show ip interface out-of-band	Display network information for the OOB port.

Enabling the DHCP Client on the Default VLAN

Beginning in Privileged EXEC mode, use the following commands to enable the DHCP client on the default VLAN, which is VLAN 1.

Command	Purpose
configure	Enter Global Configuration mode.
interface vlan 1	Enter Interface Configuration mode for VLAN 1.
ip address dhcp	Enable the DHCP client.
ipv6 address dhcp	Enable the DHCPv6 client.
CTRL + Z	Exit to Privileged EXEC mode.
show ip interface vlan 1	Display network information for VLAN 1.

Managing DHCP Leases

Beginning in Privileged EXEC mode, use the following commands to manage and troubleshoot DHCP leases on the switch.

Command	Purpose
release dhcp *interface*	Force the DHCPv4 client to release a leased address on the specified interface.
renew dhcp *interface*	Force the DHCP client to immediately renew an IPv4 address lease.
show dhcp lease interface [*interface*]	Display IPv4 addresses leased from a DHCP server.
show ipv6 dhcp interface [*interface*]	Display information about the IPv6 DHCP information for all interfaces or for the specified interface.
debug dhcp packet	Display debug information about DHCPv4 client activities and to trace DHCPv4 packets to and from the local DHCPv4 client.
debug ipv6 dhcp	Display debug information about DHCPv6 client activities and to trace DHCPv6 packets to and from the local DHCPv6 client.

Configuring Static Network Information on the OOB Port

Beginning in Privileged EXEC mode, use the following commands to configure a static IP address, subnet mask, and default gateway on the OOB port.

Command	Purpose
configure	Enter Global Configuration mode.
interface out-of-band	Enter Interface Configuration mode for the OOB port.
ip address *ip_address subnet_mask* [*gateway_ip*]	Configure a static IP address and subnet mask. Optionally, you can also configure a default gateway.
CTRL + Z	Exit to Privileged EXEC mode.
show ip interface out-of-band	Verify the network information for the OOB port.

Configuring Static Network Information on the Default VLAN

Beginning in Privileged EXEC mode, use the following commands to configure a static IP address, subnet mask, and default gateway on the default VLAN.

Command	Purpose
configure	Enter Global Configuration mode.
interface vlan 1	Enter Interface Configuration mode for VLAN 1.
ip address *ip_address subnet_mask*	Enter the IP address and subnet mask.
ipv6 address *prefix/prefix-length* [**eui64**]	Enter the IPv6 address and prefix.
ipv6 enable	Enable IPv6 on the interface.
exit	Exit to Global Configuration mode
ip default-gateway *ip_address*	Configure the default gateway.
exit	Exit to Privileged Exec mode.
show ip interface vlan 1	Verify the network information for VLAN 1.
show ipv6 interface vlan 1	Verify IPv6 network information for VLAN 1.

Configuring and Viewing Additional Network Information

Beginning in Privileged EXEC mode, use the following commands to configure a DNS server, the default domain name, and a static host name-to-address entry. Use the **show** commands to verify configured information and to view dynamic host name mappings.

Command	Purpose
configure	Enter Global Configuration mode.
ip domain-lookup	Enable IP DNS-based host name-to-address translation.
ip name-server *ip_address*	Enter the IP address of an available name server to use to resolve host names and IP addresses. You can specify up to six DNS servers. The first server you configure is the primary DNS server.
ip domain-name *name*	Define a default domain name to complete unqualified host names.
ip host *name ip_address*	Use to configure static host name-to-address mapping in the host cache.
ip address-conflict-detect run	Trigger the switch to run active address conflict detection by sending gratuitous ARP packets for IPv4 addresses on the switch.
CTRL + Z	Exit to Privileged EXEC mode.
show ip interface vlan 1	Verify the network information for VLAN 1.
show hosts	Verify the configured network information and view the dynamic host mappings.
show ip address-conflict	View the status information corresponding to the last detected address conflict.
clear ip address-conflict-detect	Clear the address conflict detection status in the switch.

Basic Network Information Configuration Example

In this example, an administrator at a Dell office in California decides not to use the Dell Easy Setup Wizard to perform the initial switch configuration. The administrator configures a PowerConnect 8000-series and 8100-series switches to obtain its information from a DHCP server on the network and creates the administrative user with read/write access. The administrator also configures the following information:

- Primary DNS server: 10.27.138.20
- Secondary DNS server: 10.27.138.21
- Default domain name: sunny.dell.com

The administrator also maps the administrative laptop host name to its IP address. The administrator uses the OOB port to manage the switch.

To configure the switch:

1. Connect the OOB port to the management network. DHCP is enabled by on the switch OOB interface by default. If the DHCP client on the switch has been disabled, use the following commands to enable the DHCP client on the OOB port.

   ```
   console#configure
   console(config)#interface out-of-band
   console(config-if)#ip address dhcp
   console(config-if)#exit
   ```

2. Configure the administrative user.

   ```
   console(config)#username admin password secret123
   level 15
   ```

3. Configure the DNS servers, default domain name, and static host mapping.

   ```
   console(config)#ip name-server 10.27.138.20
   10.27.138.21
   console(config)#ip domain-name sunny.dell.com
   console(config)#ip host admin-laptop 10.27.65.103
   console(config)#exit
   ```

4 View the network information that the DHCP server on the network dynamically assigned to the switch.

```
console#show ip interface out-of-band

IP Address..................... 10.27.22.153
Subnet Mask.................... 255.255.255.0
Default Gateway................ 10.27.22.1
Protocol Current............... DHCP
Burned In MAC Address.......... 001E.C9AA.AA08
```

5 View additional network information.

```
console#show hosts

Host name:
Default domain: sunny.dell.com dell.com
Name/address lookup is enabled
Name servers (Preference order): 10.27.138.20,
10.27.138.21
Configured host name-to-address mapping:
Host                      Addresses
-----------  ------------------------------------
admin-laptop              10.27.65.103

cache: TTL (Hours)

Host       Total   Elapsed Type      Addresses
---------  ------- ------- --------- ---------
No hostname is mapped to an IP address
```

6 Verify that the static hostname is correctly mapped.

```
console#ping admin-laptop
 Pinging admin-laptop with 0 bytes of data:

Reply From 10.27.65.103: icmp_seq = 0. time <10
msec.
Reply From 10.27.65.103: icmp_seq = 1. time <10
msec.
```

8

Managing QSFP Ports

QSFP ports can operate in 1 x 40G mode or in 4 x 10G mode. Appropriate cables must be used that match the selected mode. When changing from one mode to another, a switch reboot is required. The QSFP ports also support stacking over the interfaces in either 1 x 40G or 4 x 10G mode. Changing from Ethernet mode to stacking mode and vice-versa requires a reboot as well.

The ports on a QSFP plugin module are named Fo1/1/1-2 in 40-gigabit mode and Te1/1/1-8 in 10-gigabit mode. On the PC8164, the fixed QSFP ports are named Fo1/0/1-2 in 40-gigabit mode and Te1/0/49-56 in 10-gigabit mode. All of the possible populated or configured interfaces will show in the show interfaces status command regardless of the port mode, i.e. 40-gigabit or 10-gigabit. Unpopulated or unconfigured interfaces for plug in modules do not show in the `show interfaces status` command.

The default setting for a 40-gigabit Ethernet interface is nonstacking, 40-gigabit Ethernet (1 x 40G).

The commands to change 1 x 40G and 4 x 10G modes are always entered on the 40-gigabit interfaces.

The commands to change the Ethernet/stack mode are entered on the appropriate interface (tengigabitethernet or fortygigabitethernet). It is possible to configure some of the 10G ports in a 40G interface as stacking and not others.

To reconfigure a QSFP port, select the 40-gigabit port to change in Interface Config mode and enter the `hardware profile portmode` command with the selected mode. For example, to change a 1 x 40G port to 4 x 10G mode, enter the following commands on the forty-gigabit interface:

```
console(config)#interface fo1/1/1
console(config-if-Fo1/1/2)#hardware profile portmode 4x10g

This command will not take effect until the switch is rebooted.

console(config-if-Fo1/1/2)#do reload

Are you sure you want to reload the stack? (y/n)
```

To change a 4 x 10G port to 1 x 40G mode, enter the following commands on the 40-gigabit interface:

```
console(config)#interface Fo2/1/1
console(config-if-Fo2/1/1)#hardware profile portmode 1x40g

This command will not take effect until the switch is rebooted.

console(config-if-Fo1/1/2)#do reload

Are you sure you want to reload the stack? (y/n)
```

Attempting to change the port mode on the tengigabit interface will give the error "An invalid interface has been used for this function."

9

Managing a Switch Stack

This chapter describes how to configure and manage a stack of switches.

The topics covered in this chapter include:

- Stacking Overview
- Default Stacking Values
- Managing and Monitoring the Stack (Web)
- Managing the Stack (CLI)
- Stacking and NSF Usage Scenarios

Stacking Overview

PowerConnect 8000-series and 8100-series switches support high performance stacking, allowing increased capacity to be added as needed, without affecting network performance and providing a single point of management.

Up to six PowerConnect 8024/8024F units can be stacked together using the 10G SFP+ fiber ports only. In other words, the copper 10 GbaseT ports on the PC8024/PC8024F units cannot be used for stacking. When a combo port is configured in stacking mode, the corresponding copper port is disabled. The 10G ports default to Ethernet mode, so the ports must be reconfigured as stacking ports.

Also, up to six PowerConnect 8132/8132F/8164/8164F can be stacked using the 10G SFP+ fiber ports or the fixed 10G/40G QSFP ports. In other words, all the ports on the 8100 series switches can be used for stacking.

Stacking is supported at distances of up to 100M if the switch is configured to use Priority Flow Control on any port. Stacking using LR/LRM transceivers is supported up to the maximum distance supported by the transceiver/fiber combination (10 km for 10GBase-LR).

A single switch in the stack manages all the units in the stack (the stack master), and you manage the stack by using a single IP address. The IP address of the stack does not change, even if the stack master changes.

A stack is created by daisy-chaining stacking links on adjacent units. A stack of units is manageable as a single entity when the units are connected together. If a unit cannot detect a stacking partner on a port enabled for stacking, the unit automatically operates as a standalone unit. If a stacking partner is detected, the switch always operates in stacking mode. One unit in the stack is designated as the stack master. The master manages all the units in the stack. The stack master runs the user interface and switch software, and propagates changes to the member units. To manage a stack using the serial interface, you must connect to the serial port on the stack master. The serial interface on member units does not allow access to the CLI.

A second switch is designated as the standby unit, which becomes the master if the stack master is unavailable. You can manually configure which unit is selected as the standby, or the system can select the standby automatically.

When units are in a stack, the following activities occur:

- All units are checked for software version consistency.
- The switch Control Plane is active only on the master. The Control Plane is a software layer that manages system and hardware configuration and runs the network control protocols to set system configuration and state.
- The switch Data Plane is active on all units in the stack, including the master. The Data Plane is the set of hardware components that forward data packets without intervention from a control CPU.
- The running configuration is propagated to all units and the application state is synchronized between the master and standby during normal stacking operation. The startup configuration and backup configuration on the stack members are not overwritten with the master switch configuration.

In a stack of three or more switches, Dell strongly recommends connecting the stack in a ring topology so that each switch is connected to two other switches.

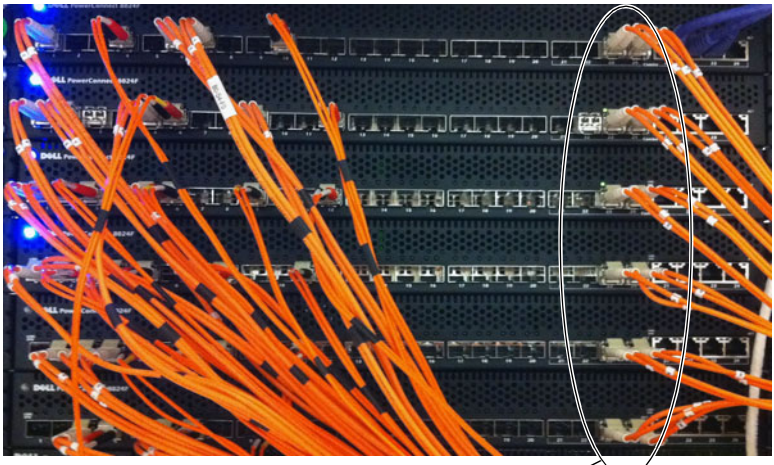
Creating a PowerConnect 8000/8100 Series Stack

You can stack up to six PC8000-series and up to six PC8100-series switches. PC8000 series switches can be stacked only with other PC8000 series switches. PC8100 series switches can only be stacked with other PC8100 series switches.

Create a stack by connecting adjacent units using the 10G ports (SFP+ ports only on the PC80xx series). It is recommended that stacking link bandwidth be at least 10 times the bandwidth of the front panel port, that is, a 10G switch (PC8100) should have 100G of stacking bandwidth to each adjacent stack member. Additional stacking connections can be made between adjacent switch units to increase the stacking bandwidth. It is strongly recommended that the stacking bandwidth be kept equal across all stacking connections; that is, avoid mixing single and double stacking connections within a stack. Up to eight links can be utilized between two stacking peers.

1. Use the CLI or web interface to configure the ports for stacking. By default, the ports are configured to operate in Ethernet mode. For more information about configuring the port mode, see "Stack Port Summary" on page 158 (Web) or "Configuring Stack Member, Stack Port, and NSF Settings" on page 162 (CLI). A reboot is required to enable the ports in stacking mode.
2. For each switch in the stack, connect one cable from a stacking port on the switch to a stacking port on the next switch. Add additional cables in parallel as needed until the desired stacking bandwidth is achieved.
3. Repeat this process until all of the devices are connected.
4. To complete the ring topology for the stack, connect one stacking port on the last switch to the remaining stacking port on the first switch. Add additional cables in parallel to achieve the desired stacking bandwidth.

Figure 9-1. Connecting a Stack of PowerConnect 8024/8024F Switches



How is the Stack Master Selected?

A stack master is elected or re-elected based on the following considerations, in order:

1. The switch is currently the stack master.
2. The switch has the higher MAC address.
3. A unit is selected as standby by the administrator, and a fail over action is manually initiated or occurs due to stack master failure.

In most cases, a switch that is added to an existing stack will become a stack member, and not the stack master. When you add a switch to the stack, one of the following scenarios takes place regarding the management status of the new switch:

- If the switch has the stack master function enabled but another stack master is already active, then the switch changes its configured stack master value to disabled.

- If the stack master function is unassigned and there is another stack master in the system then the switch changes its configured stack master value to disabled.
- If the stack master function is enabled or unassigned and there is no other stack master in the system, then the switch becomes stack master.
- If the stack master function is disabled, the unit remains a non-stack master.

If the entire stack is powered OFF and ON again, the unit that was the stack master before the reboot will remain the stack master after the stack resumes operation.

You can manually set the unit number for the switch. To avoid unit-number conflicts, one of the following scenarios takes place when you add a new member to the stack:

- If the switch has a unit number that is already in use, then the unit that you add to the stack changes its configured unit number to the lowest unassigned unit number.
- If the switch you add does not have an assigned unit number, then the switch sets its configured unit number to the lowest unassigned unit number.
- If the unit number is configured and there are no other devices using the unit number, then the switch starts using the configured unit number.
- If the switch detects that the maximum number of units already exist in the stack making it unable to assign a unit number, then the switch sets its unit number to *unassigned* and does not participate in the stack.

Adding a Switch to the Stack

When adding a new member to a stack, make sure that only the stack cables, and no network cables, are connected before powering up the new unit. Stack port configuration is stored on the member units. If stacking over Ethernet ports, configure the ports on the unit to be added to the stack as stacking ports and power the unit off prior to connecting the stacking cables. Make sure the links are not already connected to any ports of that unit. This is important because if STP is enabled and any links are UP, the STP re-convergence will take place as soon as the link is detected.

After the stack cables on the new member are connected to the stack, you can connect the power. Do not connect a new member to the stack after it is powered up. Also, do not connect two functional, powered-up stacks together. Hot insertion of units into the stack is not supported.

If a new switch is added to a stack of switches that are powered and running and already have an elected stack master, the newly added switch becomes a stack member rather than the stack master. In this situation, the firmware of the new unit may be overwritten based on the configuration of the stack master. The running configuration of the newly added unit is overwritten with the stack master configuration. Stack port configuration is always stored on the local unit and may be updated with preconfiguration information from the stack master when the unit joins the stack.

You can preconfigure information about a stack member and its ports before you add it to the stack. The preconfiguration takes place on the stack master. If there is saved configuration information on the stack master for the newly added unit, the stack master applies the configuration to the new unit; otherwise, the stack master applies the default configuration to the new unit.

Removing a Switch from the Stack

The main point to remember when you remove a unit from the stack is to disconnect all the links on the stack member to be removed. Also, be sure to take the following actions:

- Remove all the STP participating ports and wait to stabilize the STP.
- Remove all the member ports of any Port-Channels (LAGs) so there will not be any control traffic destined to those ports connected to this member.
- Statically re-route any traffic going through this unit.

When a unit in the stack fails, the stack master removes the failed unit from the stack. No changes or configuration are applied to the other stack members; however, the dynamic protocols will try to reconverge as the topology could change because of the failed unit. When there are no connected ports on the failed unit, the stack will be intact without any changes.

If you remove a unit and plan to renumber the stack, issue a **no member** *unit* command in Stack Configuration mode to delete the removed switch from the configured stack member information.

How is the Firmware Updated on the Stack?

When you add a new switch to a stack, the Stack Firmware Synchronization feature automatically synchronizes the firmware version with the version running on the stack master per the configuration on the master switch. The synchronization operation may result in either upgrade or downgrade of firmware on the mismatched stack member.

Upgrading the firmware on a stack of switches is the same as upgrading the firmware on a single switch. After you download a new image by using the File Download page or **copy** command, the downloaded image is distributed to all the connected units of the stack. For more information about downloading and installing images, see "Managing Images and Files" on page 307.

What is Stacking Standby?

A standby unit is preconfigured in the stack. If the current stack master fails, the standby unit becomes the stack master. If no switch is pre-configured as the standby unit, the software automatically selects a standby unit from the existing stack units.

When the failed master resumes normal operation, it joins the stack as a member (not a master) if the new stack master has already been elected.

The stack master copies its running configuration to the standby unit whenever it changes (subject to some restrictions to reduce overhead). This enables the standby unit to take over the stack operation with minimal interruption if the stack master becomes unavailable.

Operational state synchronization also occurs:

- when you save the running configuration to the startup configuration on the stack master.
- when the backup unit changes.

What is Nonstop Forwarding?

Networking devices are often described in terms of three semi-independent functions called the forwarding plane, the control plane, and the management plane. The forwarding plane forwards data packets and is implemented in hardware. The control plane is the set of protocols that determine how the forwarding plane should forward packets, deciding which data packets are allowed to be forwarded and where they should go.

Application software on the stack master acts as the control plane. The management plane is application software running on the stack master that provides interfaces allowing a network administrator to configure the device.

The Nonstop Forwarding (NSF) feature allows the forwarding plane of stack units to continue to forward packets while the control and management planes restart as a result of a power failure, hardware failure, or software fault on the stack master. This type of operation is called nonstop forwarding. When the stack master fails, only the switch ASICs on the stack master need to be restarted.

To prevent adjacent networking devices from rerouting traffic around the restarting device, the NSF feature uses the following three techniques:

1. A protocol can distribute a part of its control plane to stack units so that the protocol can give the appearance that it is still functional during the restart.
2. A protocol may enlist the cooperation of its neighbors through a technique known as graceful restart.
3. A protocol may simply restart after the failover if neighbors react slowly enough that they will not normally detect the outage.

The NSF feature enables the stack master unit to synchronize the running-config within 60 seconds after a configuration change has been made. However, if a lot of configuration changes happen concurrently, NSF uses a back-off mechanism to reduce the load on the switch. The **show nsf** command output includes information about when the next running-config synchronization will occur.

Initiating a Failover

The NSF feature allows you to initiate a failover, which causes the former stack master to reboot (cold start), and the new master to perform a warm restart.

Initiating a failover reloads the stack master, triggering the backup unit to take over. Before the failover, the stack master pushes application data and other important information to the backup unit. Although the handoff is controlled and causes minimal network disruption, some application state is lost, such as pending timers and other pending internal events.

Checkpointing

Switch applications (features) that build up a list of data such as neighbors or clients can significantly improve their restart behavior by remembering this data across a warm restart. This data can either be stored persistently, as DHCP server and DHCP snooping store their bindings database, or the stack master can checkpoint this data directly to the standby unit. Persistent storage allows an application on a standalone unit to retain its data across a restart, but since the amount of storage is limited, persistent storage is not always practical.

The NSF checkpoint service allows the stack master to communicate certain data to the backup unit in the stack. When the stack selects a backup unit, the checkpoint service notifies applications to start a complete checkpoint. After the initial checkpoint is done, applications checkpoint changes to their data.

NOTE: The switch cannot guarantee that a backup unit has exactly the same data that the stack master has when it fails. For example, the stack master might fail before the checkpoint service gets data to the backup if an event occurs shortly before a failover.

Table 9-1 lists the applications on the switch that checkpoint data and describes the type of data that is checkpointed.

Table 9-1. Applications that Checkpoint Data

Application	Checkpointed Data
ARP	Dynamic ARP entries
Auto VOIP	Calls in progress
Captive Portal	Authenticated clients
DHCP server	Address bindings (persistent)
DHCP snooping	DHCP bindings database
DOT1Q	Internal VLAN assignments
DOT1S	Spanning tree port roles, port states, root bridge, etc.
DOT1X	Authenticated clients
DOT3ad	Port states

Table 9-1. Applications that Checkpoint Data

Application	Checkpointed Data
IGMP/MLD Snooping	Multicast groups, list of router ports, last query data for each VLAN
IPv6 NDP	Neighbor cache entries
iSCSI	Connections
LLDP	List of interfaces with MED devices attached
OSPFv2	Neighbors and designated routers
OSPFv3	Neighbors and designated routers
Route Table Manager	IPv4 and IPv6 dynamic routes
SIM	The system's MAC addresses. System up time. IP address, network mask, default gateway on each management interface, DHCPv6 acquired IPv6 address.
Voice VLAN	VoIP phones identified by CDP or DHCP (not LLDP)

Switch Stack MAC Addressing and Stack Design Considerations

The switch stack uses the MAC addresses assigned to the stack master.

NOTE: Each switch is assigned three consecutive MAC addresses. The switch uses the MAC addresses for the service port, network port, and routing interfaces. A stack of switches uses the MAC addresses assigned to the stack master.

If the backup unit assumes control due to a stack master failure or warm restart, the backup unit continues to use the original stack master's MAC addresses. This reduces the amount of disruption to the network because ARP and other L2 entries in neighbor tables remain valid after the failover to the backup unit.

Stack units should always be connected with a ring topology (or other biconnected topology), so that the loss of a single stack link does not divide the stack into multiple stacks. If a stack is partitioned such that some units lose all connectivity to other units, then both parts of the stack start using the same MAC addresses. This can cause severe problems in the network.

If you move the stack master to a different place in the network, make sure you power down the whole stack before you redeploy the stack master so that the stack members do not continue to use the MAC address of the redeployed switch.

NSF Network Design Considerations

You can design your network to take maximum advantage of NSF. For example, by distributing a LAG's member ports across multiple units, the stack can quickly switch traffic from a port on a failed unit to a port on a surviving unit. When a unit fails, the forwarding plane of surviving units removes LAG members on the failed unit so that it only forwards traffic onto LAG members that remain up. If a LAG is left with no active members, the LAG goes down. To prevent a LAG from going down, configure LAGs with members on multiple units within the stack, when possible. If a stack unit fails, the system can continue to forward on the remaining members of the stack.

If your switch stack performs VLAN routing, another way to take advantage of NSF is to configure multiple "best paths" to the same destination on different stack members. If a unit fails, the forwarding plane removes Equal Cost Multipath (ECMP) next hops on the failed unit from all unicast forwarding table entries. If the cleanup leaves a route without any next hops, the route is deleted. The forwarding plane only selects ECMP next hops on surviving units. For this reason, try to distribute links providing ECMP paths across multiple stack units.

Why is Stacking Needed?

Stacking increases port count without requiring additional configuration. If you have multiple PowerConnect switches, stacking them helps make management of the switches easier because you configure the stack as a single unit and do not need to configure individual switches.

Default Stacking Values

Stacking is always enabled. By default, the 10G SFP+ ports are in Ethernet mode and must be configured to be used as stacking ports. Ports that are configured in stacking mode show as "detached" in the output of the s**how interfaces status** command.

Configuring an Ethernet port as a stacking port changes the default configuration of the port. The port stacking configuration does not show in the running-config. To determine the stacking configuration of a port, use the **show switch stack-ports** command.

NSF is enabled by default. You can disable NSF in order to redirect the CPU resources consumed by data checkpointing. Checkpointing only occurs when a backup unit is elected, so there is no need to disable the NSF feature on a standalone switch. When a new unit is added to the stack, the new unit takes the configuration of the stack, including the NSF setting.

Managing and Monitoring the Stack (Web)

This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring stacking on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

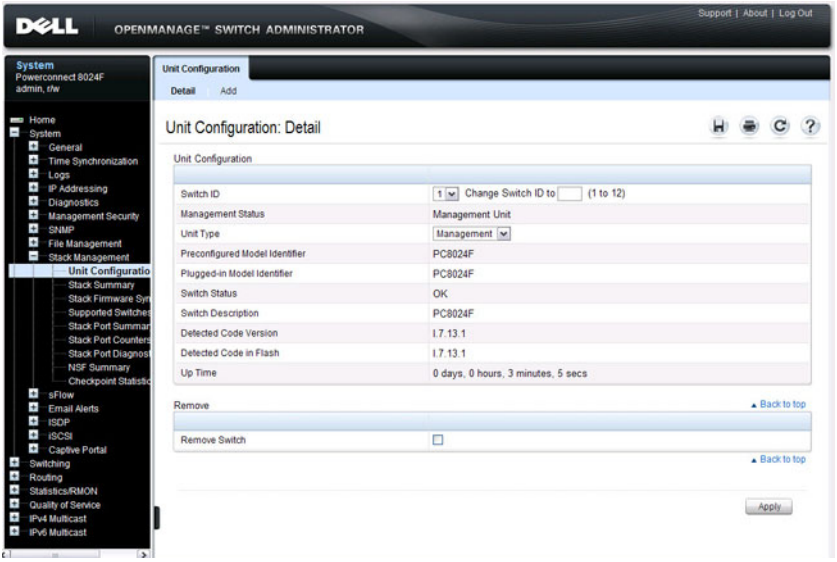
NOTE: The changes you make to the Stacking configuration pages take effect only after the device is reset.

Unit Configuration

Use the **Unit Configuration** page to change the unit number and unit type (Management, Member, or Standby).

To display the **Unit Configuration** page, click **System** → **Stack Management** → **Unit Configuration** in the navigation panel.

Figure 9-2. Stack Unit Configuration

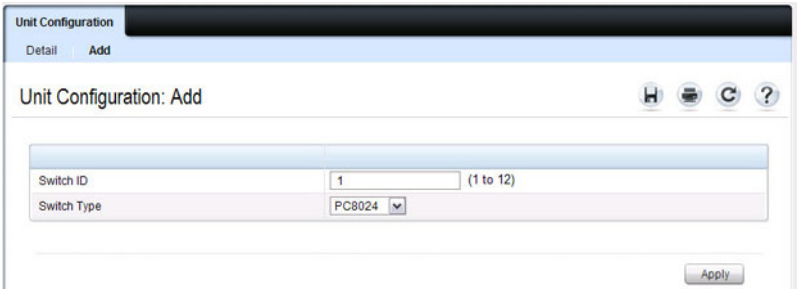


Changing the ID or Switch Type for a Stack Member

To change the switch ID or type:

1. Open the **Unit Configuration** page.
2. Click **Add** to display the **Add Unit** page.

Figure 9-3. Add Remote Log Server Settings



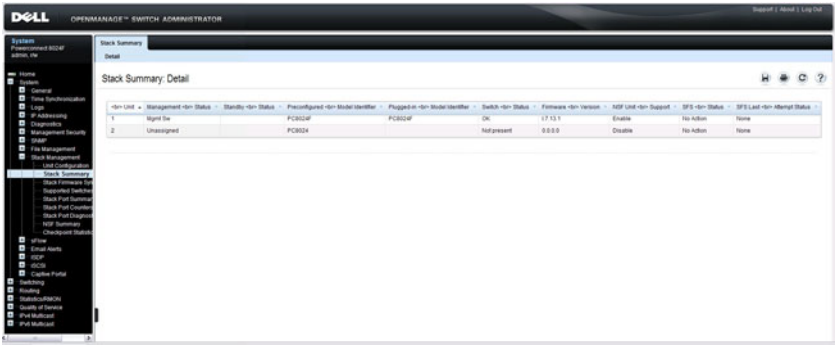
3. Specify the switch ID, and select the model number of the switch.
4. Click **Apply**.

Stack Summary

Use the **Stack Summary** page to view a summary of switches participating in the stack.

To display the **Stack Summary** page, click **System** → **Stack Management** → **Stack Summary** in the navigation panel.

Figure 9-4. Stack Summary

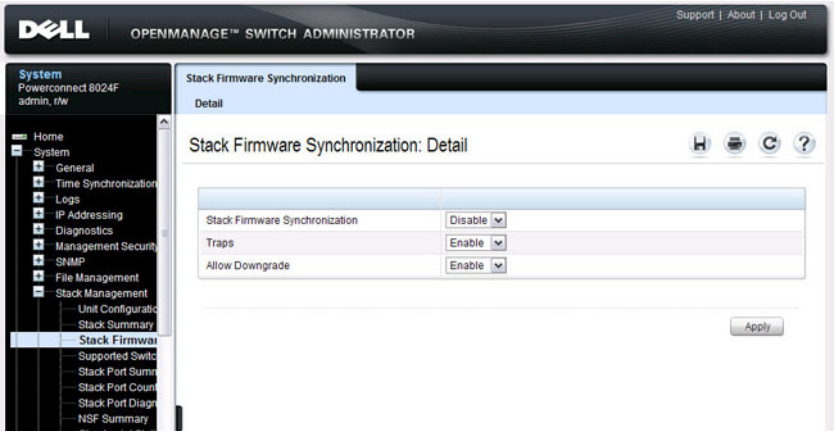


Stack Firmware Synchronization

Use the **Stack Firmware Synchronization** page to control whether the firmware image on a new stack member can be automatically upgraded or downgraded to match the firmware image of the stack master.

To display the **Stack Firmware Synchronization** page, click **System** → **Stack Management** → **Stack Firmware Synchronization** in the navigation panel.

Figure 9-5. Stack Firmware Synchronization

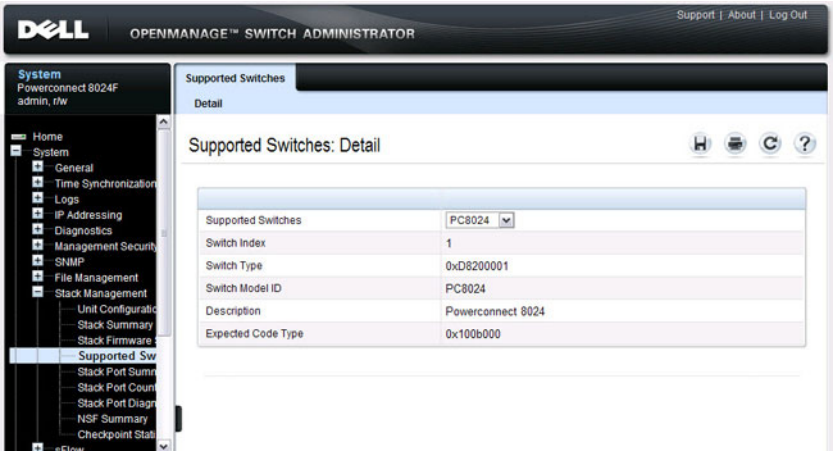


Supported Switches

Use the **Supported Switches** page to view information regarding each type of supported switch for stacking, and information regarding the supported switches.

To display the **Supported Switches** page, click **System** → **Stack Management** → **Supported Switches** in the navigation panel.

Figure 9-6. Supported Switches



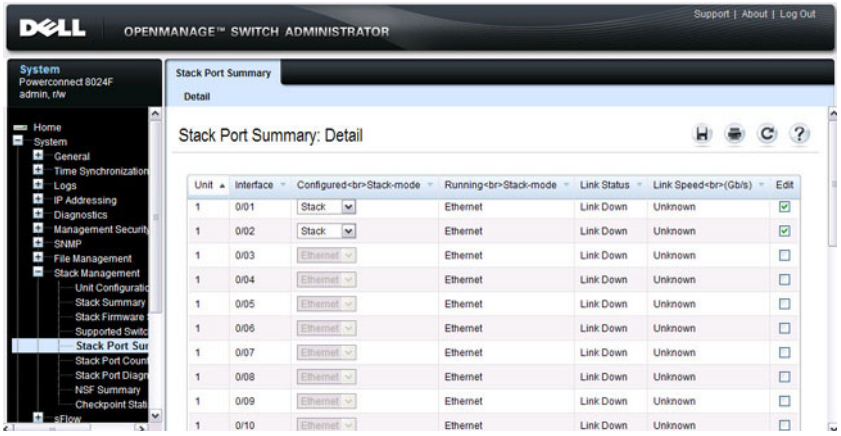
Stack Port Summary

Use the **Stack Port Summary** page to configure the stack-port mode and to view information about the stackable ports. This screen displays the unit, the stackable interface, the configured mode of the interface, the running mode as well as the link status and link speed of the stackable port.

NOTE: By default the ports are configured to operate as Ethernet ports. To configure a port as a stack port, you must change the Configured Stack Mode setting from Ethernet to Stack.

To display the **Stack Port Summary** page, click **System** → **Stack Management** → **Stack Port Summary** in the navigation panel.

Figure 9-7. Stack Port Summary

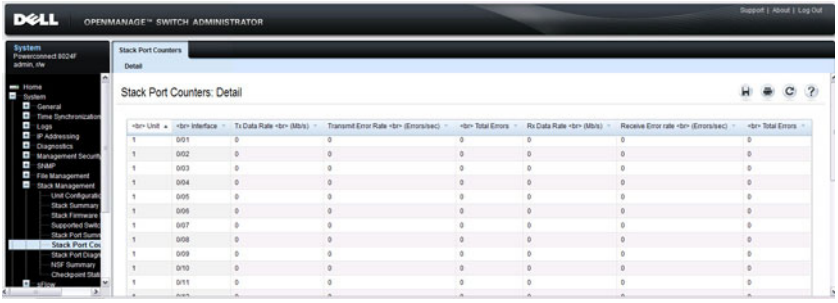


Stack Port Counters

Use the **Stack Port Counters** page to view the transmitted and received statistics, including data rate and error rate.

To display the **Stack Port Counters** page, click **System** → **Stack Management** → **Stack Point Counters** in the navigation panel.

Figure 9-8. Stack Port Counters



Stack Port Diagnostics

The **Stack Port Diagnostics** page is intended for Field Application Engineers (FAEs) and developers only

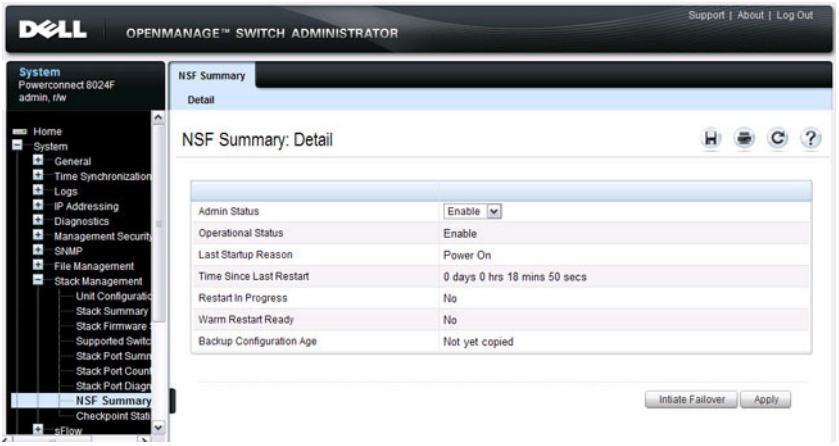
NSF Summary

Use the **NSF Summary** page to change the administrative status of the NSF feature and to view NSF information.

NOTE: The OSPF feature uses NSF to enable the hardware to continue forwarding IPv4 packets using OSPF routes while a backup unit takes over stack master responsibility. To configure NSF on a stack that uses OSPF or OSPFv3, see "NSF OSPF Configuration" on page 958 and "NSF OSPFv3 Configuration" on page 975.

To display the **NSF Summary** page, click **System** → **Stack Management** → **NSF Summary** in the navigation panel.

Figure 9-9. NSF Summary



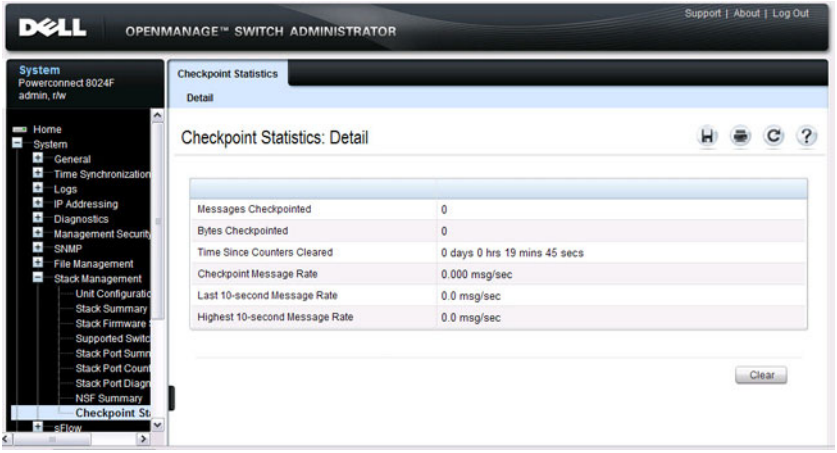
To cause the maser unit to failover to the standby unit, click **Initiate Failover**. The failover results in a warm restart of the stack master. Initiating a failover reloads the stack master, triggering the backup unit to take over.

Checkpoint Statistics

Use the **Checkpoint Statistics** page to view information about checkpoint messages generated by the stack master.

To display the **Checkpoint Statistics** page, click **System → Stack Management → Checkpoint Statistics** in the navigation panel.

Figure 9-10. Checkpoint Statistics



Managing the Stack (CLI)

This section provides information about the commands you use to manage the stack and view information about the switch stack. For more information about these commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Configuring Stack Member, Stack Port, and NSF Settings

Beginning in Privileged EXEC mode, use the following commands to configure stacking and NSF settings.

Command	Purpose
configure	Enter Global Configuration mode.
switch *current_ID* **renumber** *new_ID*	Change the switch ID number. The valid range is 1-10. **NOTE:** Changing the ID number causes all switches in the stack to be reset to perform stack master renumbering. The running configuration is cleared when the units reset.
stack	Enter Global Stack Configuration mode.
movemanagement *from_unit to_unit*	Move the management switch functionality from one switch to another.
standby *unit*	Specify the stack member that will come up as the master if a stack failover occurs.
set description *unit*	Configure a description for the specified stack member.

Command	Purpose
member *unit SID*	Add a switch to the stack and specify the model of the new stack member. • *unit* - The switch unit ID • *SID* - The index into the database of the supported switch types, indicating the type of the switch being preconfigured. **Note**: Member configuration displayed in the running config may be learned from the physical stack. Member configuration is not automatically saved in the startup-config. Save the configuration to retain the current member settings. To view the SID associated with the supported switch types, use the **show supported switchtype** command in Privileged EXEC mode.
stack-port tengigabitethernet *unit/slot/port* {**ethernet** | **stack**}	Set the mode of the port to either Ethernet or stacking.
nsf	Enable nonstop forwarding on the stack.
exit	Exit to Global Config mode.
boot auto-copy-sw	Enable the Stack Firmware Synchronization feature.
boot auto-copy-sw allow-downgrade	Allow the firmware version on the newly added stack member to be downgraded if the firmware version on manager is older.
exit	Exit to Privileged EXEC mode.
show auto-copy-sw	View the Stack Firmware Synchronization settings for the stack.
reload *unit*	If necessary, reload the specified stack member.

NOTE: The OSPF feature uses NSF to enable the hardware to continue forwarding IPv4 packets using OSPF routes while a backup unit takes over stack master responsibility. Additional NSF commands are available in OSPF and OSPFv3 command modes. For more information, see "NSF OSPF Configuration" on page 958 and "NSF OSPFv3 Configuration" on page 975

Viewing and Clearing Stacking and NSF Information

Beginning in Privileged EXEC mode, use the following commands to view stacking information and to clear NSF statistics.

Command	Purpose
show switch [*stack-member-number*]	View information about all stack members or the specified member.
show stack-standby	View the ID of the switch that will assume the role of the stack master if it goes down.
show stack-port	View information about the stacking ports.
show stack-port counters	View the statistics about the data the stacking ports have transmitted and received.
show supported switchtype	View the PowerConnect models that are supported in the stack and the switch index (SID) associated with each model.
show nsf	View summary information about the NSF feature.
show checkpoint statistics	View information about checkpoint messages generated by the stack master.
clear checkpoint statistics	Reset the checkpoint statistics counters to zero.

Stacking and NSF Usage Scenarios

Only a few settings are available to control the stacking configuration, such as the designation of the standby unit or enabling/disabling NSF. The examples in this section describe how the stacking and NSF feature act in various environments.

This section contains the following examples:

- Basic Failover
- Preconfiguring a Stack Member
- NSF in the Data Center
- NSF and VoIP
- NSF and DHCP Snooping

- NSF and the Storage Access Network
- NSF and Routed Access

Basic Failover

In this example, the stack has four members that are connected through a daisy-chain, as Figure 9-11 shows.

Figure 9-11. Basic Stack Failover

When all four units are up and running, the **show switch** CLI command gives the following output:

```
console#show switch

SW   Management     Standby   Preconfig  Plugged-  Switch  Code
     Status         Status    Model ID   in Model  Status  Version
                                         ID
---  ---------      -------   ---------  --------- ------- --------
                                                    ---
1    Stack Member             PC8024     PC8024    OK      9.19.0.2
2    Stack Member             PC8024     PC8024    OK      9.19.0.2
3    Mgmt Switch              PC8024     PC8024    OK      9.19.0.2
4    Stack Member             PC8024F    PC8024F   OK      9.19.0.2
```

At this point, if Unit 2 is powered off or rebooted due to an unexpected failure, **show switch** gives the following output:

```
console#show switch
```

When the failed unit resumes normal operation, the previous configuration that exists for that unit is reapplied by the stack master.

To permanently remove the unit from the stack, enter into Stack Config Mode and use the member command, as the following example shows.

```
console#configure
console(config)#stack
console(config-stack)#no member 2
console(config-stack)#exit
console(config)#exit
console#show switch

SW   Management    Standby  Preconfig  Plugged-  Switch      Code
     Status        Status   Model ID   in Model  Status      Version
                                       ID
---  ---------     -------  --------   --------------------  --------
1    Stack Member           PC8024     PC8024    OK          9.19.0.2
3    Mgmt Switch            PC8024     PC8024    OK          9.19.0.2
4    Stack Member           PC8024F    PC8024F   OK          9.19.0.2
```

Preconfiguring a Stack Member

To preconfigure a stack member before connecting the physical unit to the stack, use the **show support switchtype** command to obtain the SID of the unit to be added.

The example in this section demonstrates pre-configuring a PowerConnect 8024F switch on a stand-alone PowerConnect 8024 switch.

To configure the switch:

1. View the list of SIDs to determine which SID identifies the switch to preconfigure.

```
console#show supported switchtype

SID    Switch Mode ID            Code Type
------ ------------------------- ----------
1      PC8024                    0x100b000
2      PC8024F                   0x100b000
```

2. Preconfigure the 8024F switch (SID = 2) as member number 2 in the stack.

```
console#configure
console(config)#stack
console(config-stack)#member 2 2
console(config-stack)#exit
console(config)#exit
```

3 Confirm the stack configuration. Some of the fields have been omitted from the following output due to space limitations.

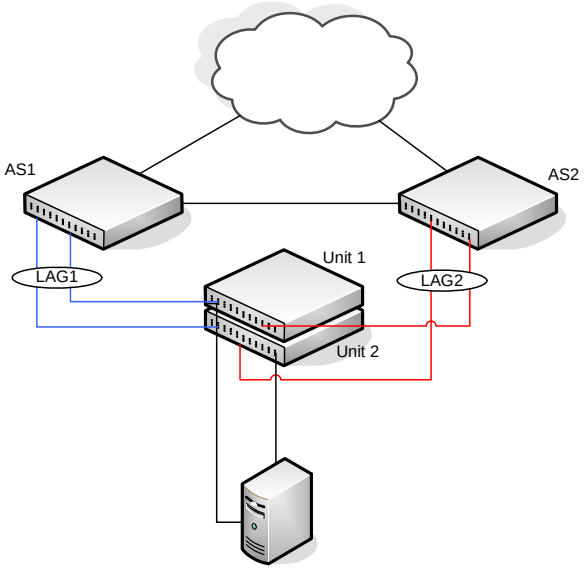
```
console#show switch

SW  Management Standby  Preconfig  Plugged-in Switch      Code
    Status     Status   Model ID   Model ID   Status      Version
--- ---------- -------- ---------- ---------- ----------  --------
1   Mgmt Sw             PC8024     PC8024     OK           M.10.2
2   Unassigned          PC8024F               Not Present  0.0.0.0
```

NSF in the Data Center

Figure 9-12 illustrates a data center scenario, where the stack of two PowerConnect switches acts as an access switch. The access switch is connected to two aggregation switches, AS1 and AS2. The stack has a link from two different units to each aggregation switch, with each pair of links grouped together in a LAG. The two LAGs and link between AS1 and AS2 are members of the same VLAN. Spanning tree is enabled on the VLAN. Assume spanning tree selects AS1 as the root bridge. Assume the LAG to AS1 is the root port on the stack and the LAG to AS2 is discarding. Unit 1 is the stack master. If unit 1 fails, the stack removes the Unit 1 link to AS1 from its LAG. The stack forwards outgoing packets through the Unit 2 link to AS1 during the failover. During the failover, the stack continues to send BPDUs and LAG PDUs on its links on Unit 2. The LAGs stay up (with one remaining link in each), and spanning tree on the aggregation switches does not see a topology change.

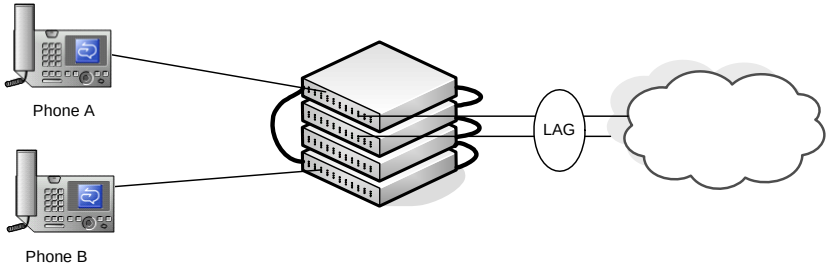
Figure 9-12. Data Center Stack Topology



NSF and VoIP

Figure 9-13 shows how NSF maintains existing voice calls during a stack master failure. Assume the top unit is the stack master. When the stack master fails, the call from phone A is immediately disconnected. The call from phone B continues. On the uplink, the forwarding plane removes the failed LAG member and continues using the remaining LAG member. If phone B has learned VLAN or priority parameters through LLDP-MED, it continues to use those parameters. The stack resumes sending LLDPDUs with MED TLVs once the control plane restarts. Phone B may miss an LLDPDU from the stack, but should not miss enough PDUs to revert its VLAN or priority, assuming the administrator has not reduced the LLDPDU interval or hold count. If phone B is receiving quality of service from policies installed in the hardware, those policies are retained across the stack master restart.

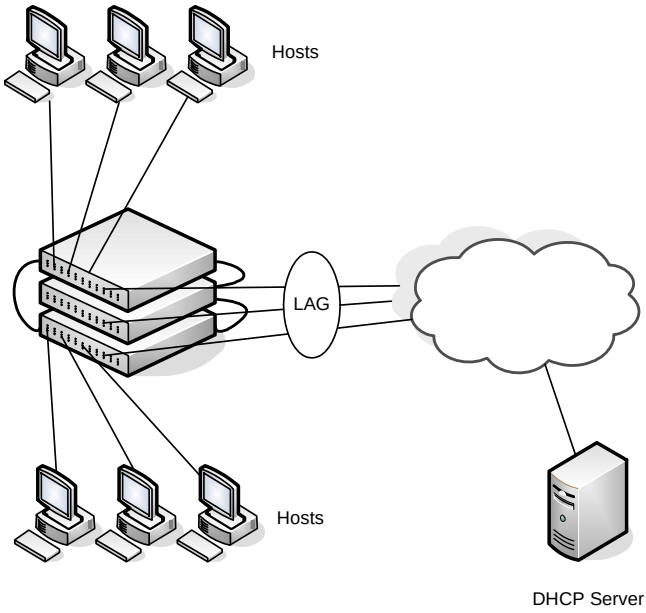
Figure 9-13. NSF and VoIP



NSF and DHCP Snooping

Figure 9-14 illustrates an L2 access switch running DHCP snooping. DHCP snooping only accepts DHCP server messages on ports configured as *trusted* ports. DHCP snooping listens to DHCP messages to build a bindings database that lists the IP address the DHCP server has assigned to each host. IP Source Guard (IPSG) uses the bindings database to filter data traffic in hardware based on source IP address and source MAC address. Dynamic ARP Inspection (DAI) uses the bindings database to verify that ARP messages contain a valid sender IP address and sender MAC address. DHCP snooping checkpoints its bindings database.

Figure 9-14. NSF and DHCP Snooping



If the stack master fails, all hosts connected to that unit lose network access until that unit reboots. The hardware on surviving units continues to enforce source filters IPSG installed prior to the failover. Valid hosts continue to communicate normally. During the failover, the hardware continues to drop data packets from unauthorized hosts so that security is not compromised.

If a host is in the middle of an exchange with the DHCP server when the failover occurs, the exchange is interrupted while the control plane restarts. When DHCP snooping is enabled, the hardware traps all DHCP packets to the CPU. The control plane drops these packets during the restart. The DHCP client and server retransmit their DHCP messages until the control plane has resumed operation and messages get through. Thus, DHCP snooping does not miss any new bindings during a failover.

As DHCP snooping applies its checkpointed DHCP bindings, IPSG confirms the existence of the bindings with the hardware by reinstalling its source IP address filters.

If Dynamic ARP Inspection is enabled on the access switch, the hardware traps ARP packets to the CPU on untrusted ports. During a restart, the control plane drops ARP packets. Thus, new traffic sessions may be briefly delayed until after the control plane restarts.

If IPSG is enabled and a DHCP binding is not checkpointed to the backup unit before the failover, that host will not be able to send data packets until it renews its IP address lease with the DHCP server.

NSF and the Storage Access Network

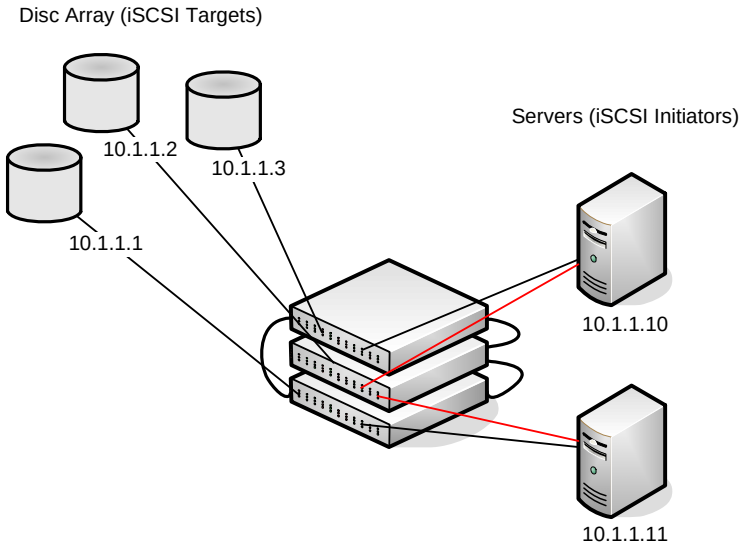
Figure 9-15 illustrates a stack of three PowerConnect switches connecting two servers (iSCSI initiators) to a disk array (iSCSI targets). There are two iSCSI connections as follows:

Session A: 10.1.1.10 to 10.1.1.3

Session B: 10.1.1.11 to 10.1.1.1

An iSCSI application running on the stack master (the top unit in the diagram) has installed priority filters to ensure that iSCSI traffic that is part of these two sessions receives priority treatment when forwarded in hardware.

Figure 9-15. NSF and a Storage Area Network



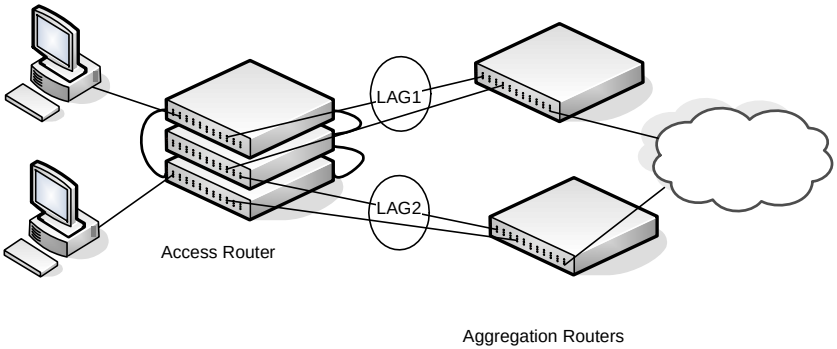
When the stack master fails, session A drops. The initiator at 10.1.1.10 detects a link down on its primary NIC and attempts to reestablish the session on its backup NIC to a different IP address on the disk array. The hardware forwards the packets to establish this new session, but assuming the session is established before the control plane is restarted on the backup unit, the new session receives no priority treatment in the hardware.

Session B remains established and fully functional throughout the restart and continues to receive priority treatment in the hardware.

NSF and Routed Access

Figure 9-16 shows a stack of three units serving as an access router for a set of hosts. Two LAGs connect the stack to two aggregation routers. Each LAG is a member of a VLAN routing interface. The stack has OSPF and PIM adjacencies with each of the aggregation routers. The top unit in the stack is the stack master.

Figure 9-16. NSF and Routed Access



If the stack master fails, its link to the aggregation router is removed from the LAG. When the control plane restarts, both routing interfaces come back up by virtue of the LAGs coming up. OSPF sends grace LSAs to inform its OSPF neighbors (the aggregation routers) that it is going through a graceful restart.

NOTE: The graceful restart feature for OSPF is disabled by default. For information about the web pages and commands to configure NSF for OSPF or OSPFv3, see "Configuring OSPF and OSPFv3" on page 931.

The grace LSAs reach the neighbors before they drop their adjacencies with the access router. PIM starts sending hello messages to its neighbors on the aggregation routers using a new generation ID to prompt the neighbors to quickly resend multicast routing information. PIM neighbors recognize the new generation ID and immediately relay the group state back to the restarting router. IGMP sends queries to relearn the hosts' interest in multicast groups. IGMP tells PIM the group membership, and PIM sends

JOIN messages upstream. The control plane updates the driver with checkpointed unicast routes. The forwarding plane reconciles L3 hardware tables.

The OSPF graceful restart finishes, and the control plane deletes any stale unicast routes not relearned at this point. The forwarding plane reconciles L3 multicast hardware tables. Throughout the process, the hosts continue to receive their multicast streams, possibly with a short interruption as the top aggregation router learns that one of its LAG members is down. The hosts see no more than a 50 ms interruption in unicast connectivity.

10

Configuring Authentication, Authorization, and Accounting

This chapter describes how to control access to the switch management interface using authentication and authorization. It also describes how to record this access using accounting. Together the three services are referred to by the acronym AAA.

The topics covered in this chapter include:

- AAA Overview
- Authentication
- Authorization
- Accounting
- Authentication Examples
- Authorization Examples
- Using RADIUS Servers to Control Management Access
- Using TACACS+ Servers to Control Management Access
- Default Configurations

AAA Overview

AAA is a framework for configuring management security in a consistent way. Three services make up AAA:

- Authentication—Validates the user identity. Authentication takes place before the user is allowed access to switch services.
- Authorization—Determines which services the user is allowed to access.
- Accounting—Collects and sends security information about users and commands.

Each service is configured using method lists. The method lists define how each service is to be performed by specifying the methods available to perform a service. The first method in a list is tried first. If the first method returns an error, the next method in the list is tried. This continues until all methods in the list have been attempted. If no method can perform the service, then the service fails. A method may return an error due to lack of network access, misconfiguration of a server, and other reasons. If there is no error, the method returns success if the user is allowed access to the service and failure if the user is not.

AAA gives the user flexibility in configuration by allowing different method lists to be assigned to different access lines. In this way, it is possible to configure different security requirements for the serial console than for telnet, for example.

Methods

A method performs the configured service. Not every method is available for every service. Some methods require a username and password and other methods only require a password. Table 10-1 summarizes the various methods:

Table 10-1. AAA Methods

Method	Username?	Password?	Can Return an Error?
enable	no	yes	yes
ias	yes	yes	no
line	no	yes	yes
local	yes	yes	no
none	no	no	no
radius	yes	yes	yes
tacacs	yes	yes	yes

Methods that never return an error cannot be followed by any other methods in a method list.

- The **enable** method uses the enable password. If there is no enable password defined, then the enable method will return an error.

- The **ias** method is a special method that is only used for 802.1X. It uses a local database (separate from the local users) that acts like an 802.1X authentication server. This method never returns an error. It will always pass or deny a user.
- The **line** method uses the password for the access line on which the user is accessing the switch. If there is no line password defined for the access line, then the line method will return an error.
- The **local** method uses the local user database. If there is no such user in the local database, then access is denied. This method never returns an error. It will always pass or deny a user.
- The **none** method does not perform any service, but instead always returns a result as if the service had succeeded. This method never returns an error.
- The **radius** and **tacacs** methods communicate with servers running the RADIUS and TACACS+ protocols, respectively. These methods can return an error if the switch is unable to contact the server.

Access Lines

There are five access lines: console, telnet, SSH, HTTP, and HTTPS. HTTP and HTTPS are not configured using AAA method lists. Instead, the authentication list for HTTP and HTTPS is configured directly (authorization and accounting are not supported). The default method lists for both the HTTP and HTTPS access lines consist of only the local method. Each of the other access lines may be assigned method lists independently for the AAA services.

Authentication

Authentication is the process of validating a user's identity. During the authentication process, only identity validation is done. There is no determination made of which switch services the user is allowed to access. This is true even when RADIUS is used for authentication; RADIUS cannot perform separate transactions for authentication and authorization. However, the RADIUS server can provide attributes during the authentication process that are used in the authorization process.

There are three types of authentication:

- **Login**— Login authentication grants access to the switch if the user credentials are validated. Access is granted only at privilege level one.
- **Enable**—Enable authentication grants access to a higher privilege level if the user credentials are validated for the higher privilege level. When RADIUS is used for enable authentication, the username for this request is always $enab15$. The username used to log into the switch is not used for RADIUS enable authentication.
- **dot1x**—Dot1x authentication is used to grant an 802.1X supplicant access to the network. For more information about 802.1X, see "Configuring Port and System Security" on page 457.

Table 10-2 shows the valid methods for each type of authentication:

Table 10-2. Valid Methods for Authentication Types

Method	Login	Enable	dot1x
enable	yes	yes	no
ias	no	no	yes
line	yes	yes	no
local	yes	no	no
none	yes	yes	yes
radius	yes	yes	yes
tacacs	yes	yes	no

Authorization

Authorization is used to determine which services the user is allowed to access. For example, the authorization process may assign a user's privilege level, which determines the set of commands the user can execute. There are three kinds of authorization: commands, exec, and network.

- **Commands**: Command authorization determines which CLI commands the user is authorized to execute.
- **Exec**: Exec authorization determines what the user is authorized to do on the switch; that is, the user's privilege level and an administrative profile.

- **Network**: Network authorization enables a RADIUS server to assign a particular 802.1X supplicant to a VLAN. For more information about 802.1X, see "Configuring Port and System Security" on page 457.

Table 10-3 shows the valid methods for each type of authorization:

Table 10-3. Authorization Methods

Method	Commands	Exec	Network
local	no	yes	no
none	yes	yes	no
radius	no	yes	yes
tacacs	yes	yes	no

Exec Authorization Capabilities

PowerConnect switches support two types of service configuration with exec authorization: privilege level and administrative profiles.

Privilege Level

By setting the privilege level during exec authorization, a user can be placed directly into Privileged EXEC mode when they log into the command line interface.

Administrative Profiles

The Administrative Profiles feature allows the network administrator to define a list of rules that control the CLI commands available to a user. These rules are collected in a “profile.” The rules in a profile can define the set of commands, or a command mode, to which a user is permitted or denied access.

Within a profile, rule numbers determine the order in which the rules are applied. When a user enters a CLI command, rules within the first profile assigned to the user are applied in descending order until there is a rule that matches the input. If no rule permitting the command is found, then the other profiles assigned to the user (if any) are searched for rules permitting the command. Rules may use regular expressions for command matching. All

profiles have an implicit “deny all” rule, such that any command that does not match any rule in the profile is considered to have been denied by that profile.

A user can be assigned to more than one profile. If there are conflicting rules in profiles, the “permit” rule always takes precedence over the “deny” rule. That is, if any profile assigned to a user permits a command, then the user is permitted access to that command. A user may be assigned up to 16 profiles.

A number of profiles are provided by default. These profiles cannot be altered by the switch administrator. See "Administrative Profiles" on page 200 for the list of default profiles.

If the successful authorization method does not provide an administrative profile for a user, then the user is permitted access based upon the user's privilege level. This means that, if a user successfully passes enable authentication or if exec authorization assigns a privilege level, the user is permitted access to all commands. This is also true if none of the administrative profiles provided are configured on the switch. If some, but not all, of the profiles provided in the authentication are configured on the switch, then the user is assigned the profiles that exist, and a message is logged that indicates which profiles could not be assigned.

Accounting

Accounting is used to record security events, such as a user logging in or executing a command. Accounting records may be sent upon completion of an event (stop-only) or at both the beginning and end of an event (start-stop). There are three types of accounting: commands, dot1x, and exec.

- **Commands**—Sends accounting records for command execution.
- **Dot1x**—Sends accounting records for network access.
- **Exec**—Sends accounting records for management access (logins).

For more information about the data sent in accounting records, see "Which RADIUS Attributes Does the Switch Support?" on page 194 and "Using TACACS+ Servers to Control Management Access" on page 197.

Table 10-4 shows the valid methods for each type of accounting:

Table 10-4. Accounting Methods

Method	Commands	Dot1x	Exec
radius	no	yes	yes
tacacs	yes	no	yes

Authentication Examples

It is important to understand that during authentication, all that happens is that the user is validated. If any attributes are returned from the server, they are not processed during authentication. In the examples below, it is assumed that the default configuration of authorization—that is, no authorization—is used.

Local Authentication Example

Use the following configuration to require local authentication when logging in over a telnet connection:

```
aaa authentication login "loc" local
line telnet
login authentication loc
exit
username guest password password
passwords strength minimum numeric-characters 2
passwords strength minimum character-classes 4
passwords strength-check
username admin password paSS1&word2 privilege 15
passwords lock-out 3
```

The following describes each line of this code:

- The `aaa authentication login "loc" local` command creates a login authentication list called "loc" that contains the method local.
- The `line telnet` command enters the configuration mode for the telnet line.
- The `login authentication loc` command assigns the loc login authentication list to be used for users accessing the switch via telnet.

- The `username guest password password` command creates a user with the name "guest" and password "password". A simple password can be configured here, since strength-checking has not yet been enabled.
- The `passwords strength minimum numeric-characters 2` command sets the minimum number of numeric characters required when password strength checking is enabled. This parameter is enabled only if the `passwords strength minimum character-classes` parameter is set to something greater than its default value of 0.
- The `passwords strength minimum character-classes 4` command sets the minimum number of character classes that must be present in the password. The possible character classes are: upper-case, lower-case, numeric and special.
- The `passwords strength-check` command enables password strength checking.
- The `username admin password paSS1&word2 privilege 15` command creates a user with the name "admin" and password "paSS1&word2". This user is enabled for privilege level 15. Note that, because password strength checking was enabled, the password was required to have at least two numeric characters, one uppercase character, one lowercase character, and one special character.
- The `passwords lock-out 3` command locks out a local user after three failed login attempts.

This configuration allows either user to log into the switch. Both users will have privilege level 1. Neither user will be able to successfully execute the `enable` command, which grants access to Privileged EXEC mode, because there is no enable password set by default (the default method list for telnet enable authentication is only the "enable" method).

NOTE: It is recommend that the password strength checking and password lockout features be enabled when using local users.

TACACS+ Authentication Example

Use the following configuration to require TACACS+ authentication when logging in over a telnet connection:

```
aaa authentication login "tacplus" tacacs
```

```
aaa authentication enable "tacp"
tacacs-server host 1.2.3.4
key "secret"
exit
line telnet
login authentication tacplus
enable authentication tacp
exit
```

The following describes each line in the above configuration:

- The `aaa authentication login "tacplus" tacacs` command creates a login authentication list called "tacplus" that contains the method tacacs. If this method returns an error, the user will fail to login.
- The `aaa authentication enable "tacp" tacacs` command creates an enable authentication list called "tacp" that contains the method tacacs. If this method fails, then the user will fail to execute the enable command.
- The `tacacs-server host 1.2.3.4` command is the first step in defining a TACACS+ server at IP address 1.2.3.4. The result of this command is to place the user in tacacs-server mode to allow further configuration of the server.
- The `key "secret"` command defines the shared secret. This must be the same as the shared secret defined on the TACACS+ server.
- The `line telnet` command enters the configuration mode for the telnet line.
- The `login authentication tacplus` command assigns the tacplus login authentication method list to be used for users accessing the switch via telnet.
- The `enable authentication tacp` command assigns the tacp enable authentication method list to be used for users executing the enable command when accessing the switch via telnet.

NOTE: A user logging in with this configuration would be placed in User EXEC mode with privilege level 1. To access Privileged EXEC mode with privilege level 15, use the enable command.

RADIUS Authentication Example

Use the following configuration to require RADIUS authentication to login over a telnet connection:

```
aaa authentication login "rad" radius
aaa authentication enable "raden" radius
radius-server host 1.2.3.4
key "secret"
exit
line telnet
login authentication rad
enable authentication raden
exit
```

The following describes each line in the above configuration:

- The `aaa authentication login "rad" radius` command creates a login authentication list called "rad" that contains the method radius. If this method returns an error, the user will fail to login.
- The `aaa authentication enable "raden" radius` command creates an enable authentication list called "raden" that contains the method radius. If this method fails, then the user will fail to execute the enable command.
- The `radius-server host 1.2.3.4` command is the first step in defining a RADIUS server at IP address 1.2.3.4. The result of this command is to place the user in radius-server mode to allow further configuration of the server.
- The `key "secret"` command defines the shared secret. This must be the same as the shared secret defined on the RADIUS server.
- The `line telnet` command enters the configuration mode for the telnet line.
- The `login authentication rad` command assigns the rad login authentication method list to be used for users accessing the switch via telnet.
- The `enable authentication raden` command assigns the raden enable authentication method list to be used for users executing the enable command when accessing the switch via telnet.

Authorization Examples

Authorization allows the administrator to control which services a user is allowed to access. Some of the things that can be controlled with authorization include the user's initial privilege level and which commands the user is allowed to execute. When authorization fails, the user is denied access to the switch, even though the user has passed authentication.

The following examples assume that the configuration used in the previous examples has already been applied.

Local Authorization Example—Direct Login to Privileged EXEC Mode

Apply the following configuration to use the local user database for authorization, such that a user can enter privileged EXEC mode directly:

```
aaa authorization exec "locex" local
line telnet
authorization exec locex
exit
```

With the users that were previously configured, the guest user will still log into user EXEC mode, since the guest user only has privilege level 1 (the default). The admin user will be able to login directly to privileged EXEC mode since his privilege level was configured as 15.

TACACS+ Authorization Example—Direct Login to Privileged EXEC Mode

Apply the following configuration to use TACACS+ for authorization, such that a user can enter privileged EXEC mode directly:

```
aaa authorization exec "tacex" tacacs
line telnet
authorization exec tacex
exit
```

Configure the TACACS+ server so that the shell service is enabled and the priv-lvl attribute is sent when user authorization is performed. For example:

shell:priv-lvl=15

The following describes each line in the above configuration:

- The `aaa authorization exec "tacex" tacacs` command creates an exec authorization method list called tacex which contains the method tacacs.
- The `authorization exec tacex` command assigns the tacex exec authorization method list to be used for users accessing the switch via telnet.

Notes:

- If the privilege level is zero (that is, blocked), then authorization will fail and the user will be denied access to the switch.
- If the privilege level is higher than one, the user will be placed directly in Privileged EXEC mode. Note that all commands in Privileged EXEC mode require privilege level 15, so assigning a user a lower privilege level will be of no value.
- A privilege level greater than 15 is invalid and treated as if privilege level zero had been supplied.
- The shell service must be enabled on the TACACS+ server. If this service is not enabled, authorization will fail and the user will be denied access to the switch.

TACACS+ Authorization Example—Administrative Profiles

The switch should use the same configuration as for the previous authorization example.

The TACACS+ server should be configured such that it will send the "roles" attribute. For example:

```
shell:roles=router-admin
```

The above example attribute will give the user access to the commands permitted by the router-admin profile.

NOTE: If the priv-lvl attribute is also supplied, the user can also be placed directly into privileged EXEC mode.

TACACS+ Authorization Example—Custom Administrative Profile

This example creates a custom profile that allows the user to control user access to the switch by configuring a administrative profile that only allows access to AAA related commands. Use the following commands to create the administrative profile:

```
admin-profile aaa
rule 99 permit command "^show aaa .*"
rule 98 permit command "^show authentication .*"
rule 97 permit command "^show authorization .*"
rule 96 permit command "^show accounting .*"
rule 95 permit command "^show tacacs .*"
rule 94 permit command "^aaa .*"
rule 93 permit command "^line .*"
rule 92 permit command "^login .*"
rule 91 permit command "^authorization .*"
rule 90 permit command "^accounting .*"
rule 89 permit command "^configure .*"
rule 88 permit command "^password .*"
rule 87 permit command "^username .*"
rule 86 permit command "^show user.*"
rule 85 permit command "^radius-server .*"
rule 84 permit command "^tacacs-server .*"
rule 83 permit mode radius-auth-config
rule 82 permit mode radius-acct-config
rule 81 permit mode tacacs-config
exit
```

The following describes each line in the above configuration:

- The `admin-profile aaa` command will create an administrative profile call aaa and place the user in admin-profile-config mode.
- Each `rule` *`number`* `permit command` *`regex`* command allows any command that matches the regular expression.
- Each `rule` *`number`* `permit mode` *`mode-name`* command allows all commands in the named mode.
- The command rules use regular expressions as implemented by Henry Spencer's regex library (the POSIX 1003.2 compliant version). In the regular expressions used in this example, the caret (^) matches the null

string at the beginning of a line, the period (.) matches any single character, and the asterisk (*) repeats the previous match zero or more times.

- To assign this profile to a user, configure the TACACS+ server so that it sends the following "roles" attribute for the user:

  ```
  shell:roles=aaa
  ```

  If it is desired to also permit the user access to network-operator commands (basically, all the command in User EXEC mode), then the "roles" attribute would be configured as follows:

  ```
  shell:roles=aaa,network-operator
  ```

TACACS+ Authorization Example—Per-command Authorization

An alternative method for command authorization is to use the TACACS+ feature of per-command authorization. With this feature, every time the user enters a command, a request is sent to the TACACS+ server to ask if the user is permitted to execute that command. Exec authorization does not need to be configured to use per-command authorization.

Apply the following configuration to use TACACS+ to authorize commands:

```
aaa authorization commands "taccmd" tacacs
line telnet
authorization commands taccmd
exit
```

The following describes each line in the above configuration:

- The `aaa authorization commands "taccmd" tacacs` command creates a command authorization method list called taccmd that includes the method tacacs.
- The `authorization commands taccmd` command assigns the taccmd command authorization method list to be used for users accessing the switch via telnet.

The TACACS+ server must be configured with the commands that the user is allowed to execute. If the server is configured for command authorization as "None", then no commands will be authorized. If both administrative

profiles and per-command authorization are configured for a user, any command must be permitted by both the administrative profiles and by per-command authorization.

RADIUS Authorization Example—Direct Login to Privileged EXEC Mode

Apply the following configuration to use RADIUS for authorization, such that a user can enter privileged exec mode directly:

```
aaa authorization exec "rad" radius
line telnet
authorization exec rad
exit
```

Configure the RADIUS server so that the RADIUS attribute Service Type (6) is sent with value Administrative. Any value other than Administrative is interpreted as privilege level 1.

The following describes each line in the above configuration:

- The `aaa authorization exec "rad" radius` command creates an exec authorization method list called "rad" that contains the method radius.
- The `authorization exec rad` command assigns the rad exec authorization method list to be used for users accessing the switch via telnet.

Notes:

- If the privilege level is zero (that is, blocked), then authorization will fail and the user will be denied access to the switch.
- If the privilege level is higher than one, the user will be placed directly in Privileged EXEC mode. Note that all commands in Privileged EXEC mode require privilege level 15, so assigning a user a lower privilege level will be of no value.
- A privilege level greater than 15 is invalid and treated as if privilege level zero had been supplied.

RADIUS Authorization Example—Administrative Profiles

The switch should use the same configuration as in the previous authorization example.

The RADIUS server should be configured such that it will send the Cisco AV Pair attribute with the "roles" value. For example:

```
shell:roles=router-admin
```

The above example attribute gives the user access to the commands permitted by the router-admin profile.

Using RADIUS Servers to Control Management Access

The RADIUS client on the switch supports multiple RADIUS servers. When multiple authentication servers are configured, they can help provide redundancy. One server can be designated as the primary and the other(s) will function as backup server(s). The switch attempts to use the primary server first. if the primary server does not respond, the switch attempts to use the backup servers. A priority value can be configured to determine the order in which the backup servers are contacted.

How Does RADIUS Control Management Access?

Many networks use a RADIUS server to maintain a centralized user database that contains per-user authentication information. RADIUS servers provide a centralized authentication method for:

- Telnet Access
- Web Access
- Console to Switch Access
- Access Control Port (802.1X)

Like TACACS+, RADIUS access control utilizes a database of user information on a remote server. Making use of a single database of accessible information—as in an Authentication Server—can greatly simplify the authentication and management of users in a large network. One such type of Authentication Server supports the Remote Authentication Dial In User Service (RADIUS) protocol as defined by RFC 2865.

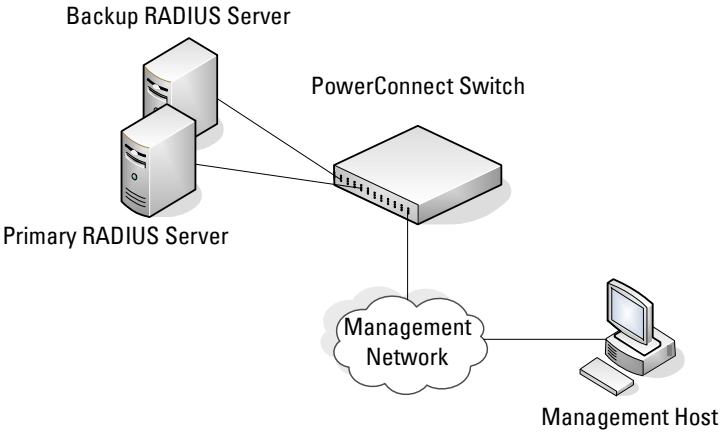
For authenticating users prior to access, the RADIUS standard has become the protocol of choice by administrators of large accessible networks. To accomplish the authentication in a secure manner, the RADIUS client and RADIUS server must both be configured with the same shared password or

"secret". This "secret" is used to generate one-way encrypted authenticators that are present in all RADIUS packets. The "secret" is never transmitted over the network.

RADIUS conforms to a secure communications client/server model using UDP as a transport protocol. It is extremely flexible, supporting a variety of methods to authenticate and statistically track users. RADIUS is also extensible, allowing for new methods of authentication to be added without disrupting existing functionality.

As a user attempts to connect to the switch management interface, the switch first detects the contact and prompts the user for a name and password. The switch encrypts the supplied information, and a RADIUS client transports the request to a pre-configured RADIUS server.

Figure 10-1. RADIUS Topology



The server can authenticate the user itself or make use of a back-end device to ascertain authenticity. In either case a response may or may not be forthcoming to the client. If the server accepts the user, it returns a positive result with attributes containing configuration information. If the server rejects the user, it returns a negative result. If the server rejects the client or the shared *secrets* differ, the server returns no result. If the server requires additional verification from the user, it returns a challenge, and the request process begins again.

If you use a RADIUS server to authenticate users, you must configure user attributes in the user database on the RADIUS server. The user attributes include the user name, password, and privilege level.

NOTE: To set the privilege level, it is recommended to use the Service-Type attribute instead of the Cisco AV pair priv-lvl attribute.

Which RADIUS Attributes Does the Switch Support?

Table 10-5 lists the RADIUS attributes that the switch supports and indicates whether the 802.1X feature, user management feature, or Captive Portal feature supports the attribute. You can configure these attributes on the RADIUS server(s) when utilizing the switch RADIUS service.

Table 10-5. Supported RADIUS Attributes

Type	RADIUS Attribute Name	802.1X	User Manager	Captive Portal
1	USER-NAME	Yes	Yes	No
2	USER-PASSWORD	Yes	Yes	No
4	NAS-IP-ADDRESS	Yes	No	No
5	NAS-PORT	Yes	Yes	No
6	SERVICE-TYPE	No	Yes	No
11	FILTER-ID	Yes	No	No
12	FRAMED-MTU	Yes	No	No
15	LOGIN-SERVICE	No	Yes	
18	REPLY-MESSAGE	Yes	Yes	No
24	STATE	Yes	Yes	No
25	CLASS	Yes	No	No
26	VENDOR-SPECIFIC	No	Yes	Yes
27	SESSION-TIMEOUT	Yes	No	Yes
28	IDLE-TIMEOUT	No	No	Yes
29	TERMINATION-ACTION	Yes	No	No
30	CALLED-STATION-ID	Yes	No	No

Table 10-5. Supported RADIUS Attributes (Continued)

Type	RADIUS Attribute Name	802.1X	User Manager	Captive Portal
31	CALLING-STATION-ID	Yes	No	No
32	NAS-IDENTIFIER	Yes	Yes	No
40	ACCT-STATUS-TYPE	Set by RADIUS client for Accounting	Yes	No
42	ACCT-INPUT-OCTETS	Yes	No	No
43	ACCT-OUTPUT-OCTETS	Yes	No	No
44	ACCT-SESSION-ID	Set by RADIUS client for Accounting	Yes	No
46	ACCT-SESSION-TIME	Yes	Yes	No
49	ACCT-TERMINATECAUSE	Yes	No	No
52	ACCT-INPUTGIGAWORDS	Yes	No	No
53	ACCT-OUTPUTGIGAWORDS	Yes	No	No
61	NAS-PORT-TYPE	Yes	No	No
64	TUNNEL-TYPE	Yes	No	No
65	TUNNEL-MEDIUM-TYPE	Yes	No	No
79	EAP-MESSAGE	Yes	No	No
80	MESSAGEAUTHENTICATOR	Set by RADIUS client for Accounting	Yes	No
81	TUNNEL-PRIVATEGROUP-ID	Yes	No	No

How Are RADIUS Attributes Processed on the Switch?

The following attributes are processed in the RADIUS Access-Accept message received from a RADIUS server:

- NAS-PORT—ifIndex of the port to be authenticated.
- REPLY-MESSAGE—Trigger to respond to the Access-Accept message with an EAP notification.
- STATE-RADIUS—Server state. Transmitted in Access-Request and Accounting-Request messages.
- SESSION-TIMEOUT—Session timeout value for the session (in seconds). Used by both 802.1x and Captive Portal.
- TERMINATION-ACTION—Indication as to the action taken when the service is completed.
- EAP-MESSAGE—Contains an EAP message to be sent to the user. This is typically used for MAB clients.
- VENDOR-SPECIFIC—The following Cisco AV Pairs are supported:
  - shell:priv-lvl
  - shell:roles
- FILTER-ID—Name of the filter list for this user.
- TUNNEL-TYPE—Used to indicate that a VLAN is to be assigned to the user when set to tunnel type VLAN (13).
- TUNNEL-MEDIUM-TYPE—Used to indicate the tunnel medium type. Must be set to medium type 802 (6) to enable VLAN assignment.
- TUNNEL-PRIVATE-GROUP-ID—Used to indicate the VLAN to be assigned to the user. May be a string which matches a preconfigured VLAN name or a VLAN id. If a VLAN id is given, the string must only contain decimal digits.

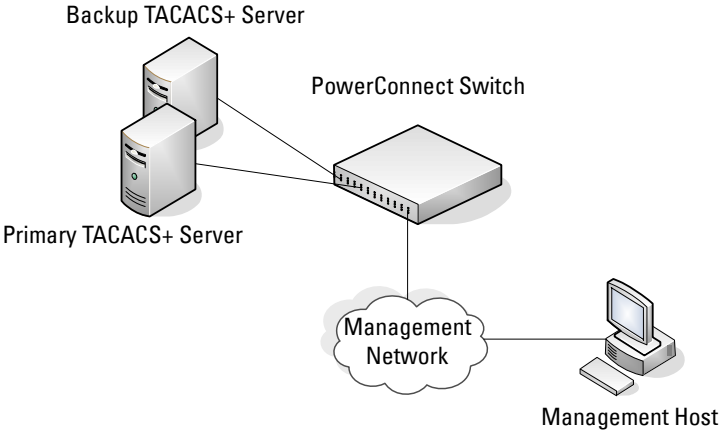
Using TACACS+ Servers to Control Management Access

TACACS+ (Terminal Access Controller Access Control System) provides access control for networked devices via one or more centralized servers. TACACS+ simplifies authentication by making use of a single database that can be shared by many clients on a large network. TACACS+ uses TCP to ensure reliable delivery and a shared key configured on the client and daemon server to encrypt all messages.

If you configure TACACS+ as the authentication method for user login and a user attempts to access the user interface on the switch, the switch prompts for the user login credentials and requests services from the TACACS+ client. The client then uses the configured list of servers for authentication, and provides results back to the switch.

Figure 10-2 shows an example of access management using TACACS+.

Figure 10-2. Basic TACACS+ Topology



You can configure the TACACS+ server list with one or more hosts defined via their network IP address. You can also assign each a priority to determine the order in which the TACACS+ client will contact them. TACACS+ contacts the server when a connection attempt fails or times out for a higher priority server.

You can configure each server host with a specific connection type, port, timeout, and shared key, or you can use global configuration for the key and timeout.

The TACACS+ server can do the authentication itself, or redirect the request to another back-end device. All sensitive information is encrypted and the shared secret is never passed over the network; it is used only to encrypt the data.

Which TACACS+ Attributes Does the Switch Support?

Table 10-6 lists the TACACS+ attributes that the switch supports and indicates whether the authorization or accounting service supports sending or receiving the attribute. The authentication service does not use attributes. You can configure these attributes on the TACACS+ server(s) when utilizing the switch TACACS+ service.

Table 10-6. Supported TACACS+ Attributes

Attribute Name	Exec Authorization	Command Authorization	Accounting
cmd	both (optional)	sent	sent
cmd-arg		sent	
elapsed-time			sent
priv-lvl	received		
protocol			sent
roles	both (optional)		
service=shell	both	sent	sent
start-time			sent
stop-time			sent

Default Configurations

Method Lists

The method lists shown in Table 10-7 are defined by default. They cannot be deleted, but they can be modified. Using the "no" command on these lists will return them to their default configuration.

Table 10-7. Default Method Lists

AAA Service (type)	List Name	List Methods
Authentication (login)	defaultList	none
Authentication (login)	networkList	local
Authentication (enable)	enableList	enable none
Authentication (enable)	enableNetList	enable
Authorization (exec)	dfltExecAuthList	none
Authorization (commands)	dfltCmdAuthList	none
Accounting (exec)	dfltExecList	tacacs (start-stop)
Accounting (commands)	dfltCmdList	tacacs (stop-only)

Access Lines (AAA)

Table 10-8 shows the method lists assigned to the various access lines by default.

Table 10-8. Default AAA Methods

AAA Service (type)	Console	Telnet	SSH
Authentication (login)	defaultList	networkList	networkList
Authentication (enable)	enableList	enableNetList	enableNetList
Authorization (exec)	dfltExecAuthList	dfltExecAuthList	dfltExecAuthList
Authorization (commands)	dfltCmdAuthList	dfltCmdAuthList	dfltCmdAuthList

Table 10-8. Default AAA Methods (Continued)

AAA Service (type)	Console	Telnet	SSH
Accounting (exec)	none	none	none
Accounting (commands)	none	none	none

Access Lines (Non-AAA)

Table 10-9 shows the default configuration of the access lines that do not use method lists.

Table 10-9. Default Configuration for Non-AAA Access Lines

Access Line	Authentication	Authorization
HTTP	local	n/a
HTTPS	local	n/a
802.1X	none	none

Administrative Profiles

The administrative profiles shown in Table 10-10 are system-defined and may not be deleted or altered. To see the rules in a profile, use the `show admin-profiles name` *`profile name`* command.

Table 10-10. Default Administrative Profiles

Name	Description
network-admin	Allows access to all commands.
network-security	Allows access to network security features such as 802.1X, Voice VLAN, Dynamic ARP Inspection and IP Source Guard.
router-admin	Allows access to Layer 3 features such as IPv4 Routing, IPv6 Routing, OSPF, RIP, etc.
multicast-admin	Allows access to multicast features at all layers, this includes L2, IPv4 and IPv6 multicast, IGMP, IGMP Snooping, etc.
dhcp-admin	Allows access to DHCP related features such as DHCP Server and DHCP Snooping.

Table 10-10. Default Administrative Profiles (Continued)

Name	Description
CP-admin	Allows access to the Captive Portal feature.
network-operator	Allows access to all User EXEC mode commands and show commands.

11

Monitoring and Logging System Information

This chapter provides information about the features you use to monitor the switch, including logging, cable tests, and email alerting. The topics covered in this chapter include:

- System Monitoring Overview
- Default Log Settings
- Monitoring System Information and Configuring Logging (Web)
- Monitoring System Information and Configuring Logging (CLI)
- Logging Configuration Examples

System Monitoring Overview

What System Information Is Monitored?

The CLI and web-based interfaces provide information about physical aspects of the switch, such as system health and cable diagnostics, as well as information about system events, such as management login history. The switch also reports system resource usage.

The system logging utility can monitor a variety of events, including the following:

- System events — System state changes and errors that range in severity from Emergency to Debug
- Audit events — Attempts to login or logout from the switch and attempts to perform any operations with files on the flash drive
- CLI commands — Commands executed from the CLI
- Web page visits — Pages viewed by using OpenManage Switch Administrator
- SNMP events — SNMP **set** operations

Why Is System Information Needed?

The information the switch provides can help you troubleshoot issues that might be affecting system performance. The cable diagnostics test help you troubleshoot problems with the physical connections to the switch. Auditing access to the switch and the activities an administrator performed while managing the switch can help provide security and accountability.

Where Are Log Messages Sent?

The messages the switch generates in response to events, faults, errors, and configuration changes can be recorded in several locations. By default, these messages are stored locally on the switch in the RAM (cache). This collection of log files is called the RAM log or buffered log. When the RAM log file reaches the configured maximum size, the oldest message is deleted from the RAM when a new message is added. If the system restarts, all messages are cleared.

In addition to the RAM log, you can specify that log files are sent to the following sources:

- Console — If you are connected to the switch CLI through the console port, messages display to the screen as they are generated.
- Log file — Messages sent to the log file are saved in the flash memory and are not cleared when the system restarts.
- Remote server — Messages can be sent to a remote log server for viewing and storage.
- Email — Messages can be sent to one or more email addresses. You must configure information about the network Simple Mail Transport Protocol SMTP) server for email to be successfully sent from the switch.

What Are the Severity Levels?

For each local or remote log file, you can specify the severity of the messages to log. Each severity level is identified by a name and a number. Table 11-1 provides information about the severity levels.

Table 11-1. Log Message Severity

Severity Keyword	Severity Level	Description
emergencies	0	The switch is unusable.
alerts	1	Action must be taken immediately.
critical	2	The switch is experiencing critical conditions.
errors	3	The switch is experiencing error conditions.
warnings	4	The switch is experiencing warning conditions.
notification	5	The switch is experiencing normal but significant conditions.
informational	6	The switch is providing non-critical information.
debugging	7	The switch is providing debug-level information.

When you specify the severity level, messages with that severity level and higher are sent to the log file. For example, if you specify the severity level as critical, messages with a severity level of alert and emergency are also logged. When you specify the severity level in a CLI command, you can use the keyword or the numerical level.

What Are the System Startup and Operation Logs?

Two types of log files exist in flash (persistent) memory:

- The first log type is the system startup log. The system startup log stores the first 32 messages received after system reboot. The log file stops when it is full.
- The second log type is the system operation log. The system operation log stores the last 1000 messages received during system operation. The oldest messages are overwritten when the file is full.

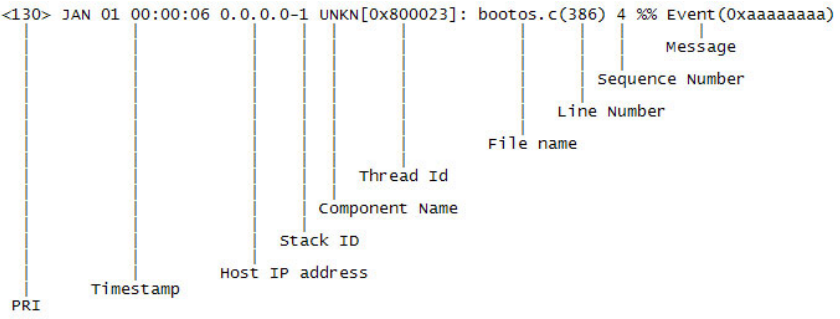
A message is only logged in one file. On system startup, if the Log file is enabled, the startup log stores messages up to its limit. Then the operation log begins to store the messages.

To view the log messages in the system startup and operational log files, you must download the log files to an administrative host. For more information about downloading files, see "Managing Images and Files" on page 307.

What Is the Log Message Format?

The first part of the log message up to the first left bracket is fixed by the Syslog standard (RFC 3164). The second part up to the two percent signs is standardized for all Dell PowerConnect logs. The variable text of the log message follows. The log message is limited to 96 bytes.

Each log message uses the following format:



- PRI—This consists of the facility code (see RFC 3164) multiplied by 8 and added to the severity. The log messages use the local7 facility code (23). This implies that a message of severity 0 will have a priority of 184 and a message of severity 7 will have a priority of 191.
- Timestamp—This is the system up time. For systems that use SNTP, this is UTC. When time zones are enabled, local time will be used.
- Host IP address—This is the IP address of the local system.
- Stack ID —This is the assigned stack ID. For the PowerConnect 8000-series and 8100-series switches, the Stack ID number is always 1.
- Component name—The component name for the logging component. Component “UNKN” is substituted for components that do not identify themselves to the logging component.
- Thread ID—The thread ID of the logging component.
- File name —The name of the file containing the invoking macro.

- Line number —The line number which contains the invoking macro.
- Sequence number —The message sequence number for this stack component. Sequence numbers may be skipped because of filtering but are always monotonically increasing on a per-stack member basis.
- Message — Contains the text of the log message.

What Factors Should Be Considered When Configuring Logging?

Dell recommends that network administrators deploy a syslog server in their network and configure all switches to log messages to the syslog server.

Default Log Settings

System logging is enabled, and messages are sent to the console (severity level: warning and above), and RAM log (severity level: informational and above). Switch auditing, CLI command logging, Web logging, and SNMP logging are disabled. No messages are sent to the log file that is stored in flash, and no remote log servers are defined.

Email alerting is disabled, and no recipient email address is configured. Additionally, no mail server is defined. If you add a mail server, by default, no authentication or security protocols are configured, and the switch uses TCP port 25 for SMTP.

After you enable email alerting and configure the mail server and recipient email address, log messages with a severity level of emergency and alert are sent immediately with each log message in a separate mail. The email subject is “Urgent Log Messages.” Log messages with a severity level of critical, error, and warning are sent periodically in a single email. The email subject is “Non Urgent Log Messages.” Messages with a severity level of notice and below are not sent in an email.

Monitoring System Information and Configuring Logging (Web)

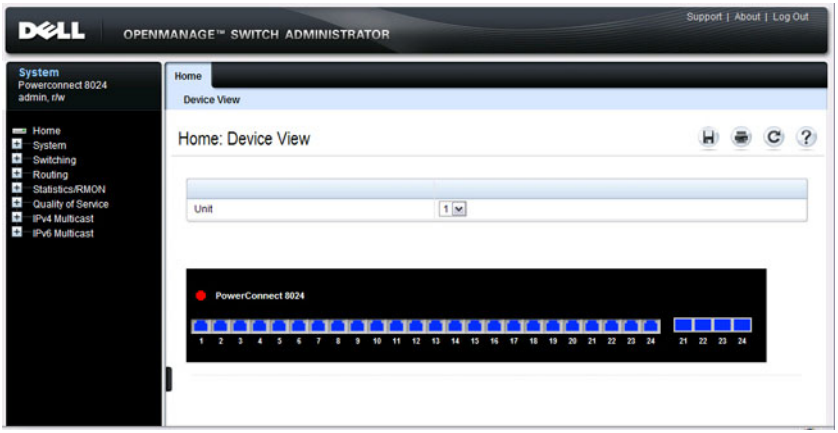
This section provides information about the OpenManage Switch Administrator pages to use to monitor system information and configure logging on the PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

Device Information

The **Device Information** page displays after you successfully log on to the switch by using the Dell OpenManage Switch Administrator. This page is a virtual representation of the switch front panel. Use the **Device Information** page to view information about the port status or system status. Click on a port to access the **Port Configuration** page for the selected port.

To display the **Device Information** page, click **Home** in the navigation panel.

Figure 11-1. Device Information



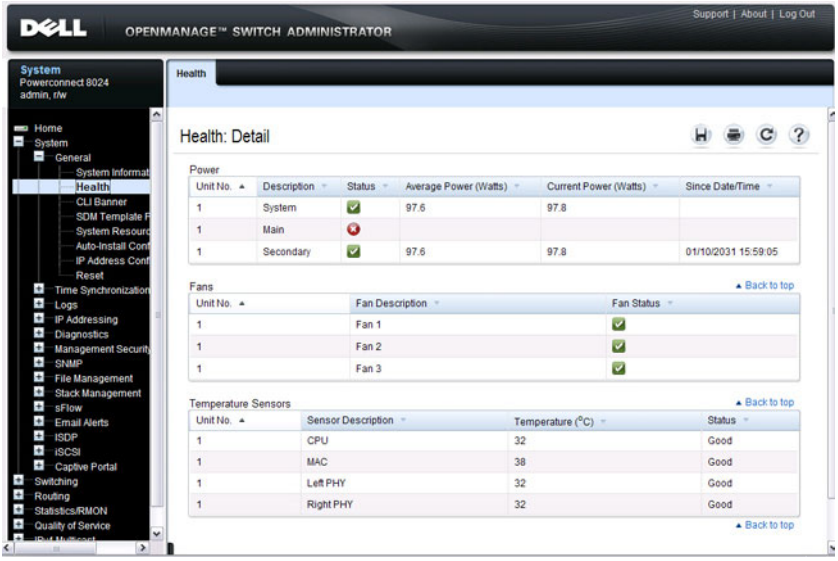
For more information about the device view features, see "Understanding the Device View" on page 106.

System Health

Use the **Health** page to view status information about the switch power and ventilation sources.

To display the **Health** page, click **System** → **General** → **Health** in the navigation panel.

Figure 11-2. Health

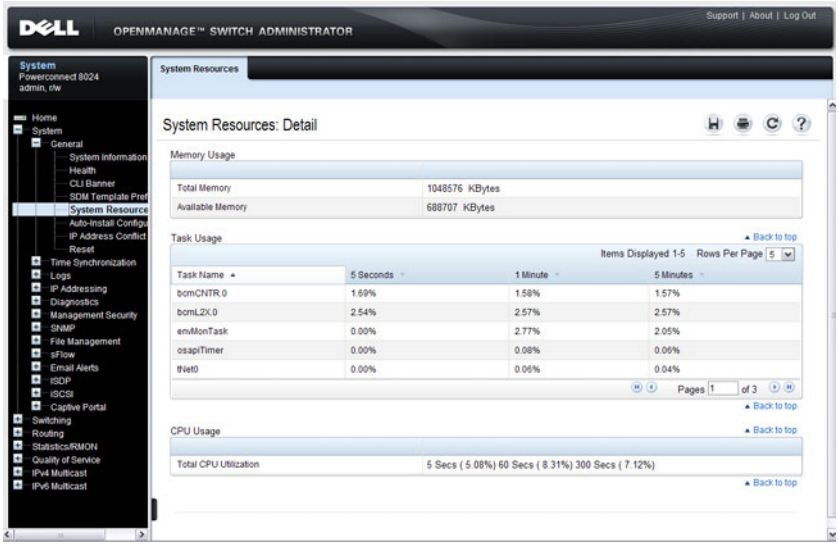


System Resources

Use the **System Resources** page to view information about memory usage and task utilization.

To display the **System Resources** page, click **System** → **General** → **System Resources** in the navigation panel.

Figure 11-3. System Resources

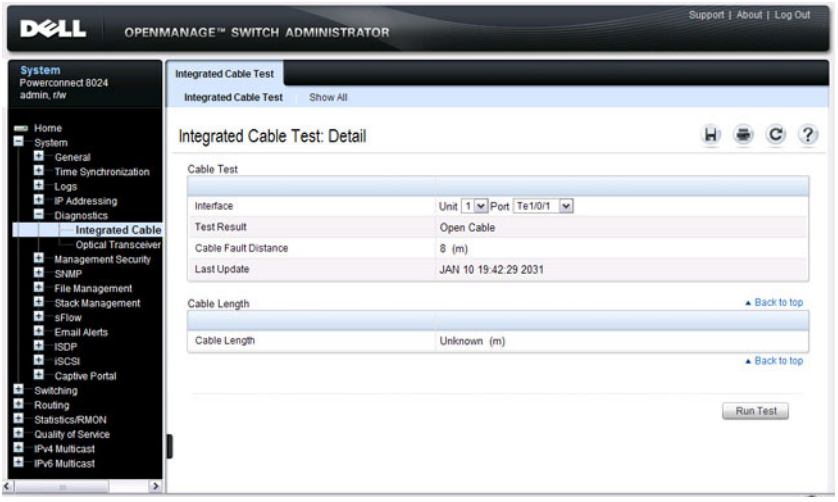


Integrated Cable Test for Copper Cables

Use the **Integrated Cable Test for Copper Cables** page to perform tests on copper cables. Cable testing provides information about where errors occurred in the cable, the last time a cable test was performed, and the type of cable error which occurred. The tests use Time Domain Reflectometry (TDR) technology to test the quality and characteristics of a copper cable attached to a port. Cables up to 120 meters long can be tested. Cables are tested when the ports are in the down state, with the exception of the Approximated Cable Length test.

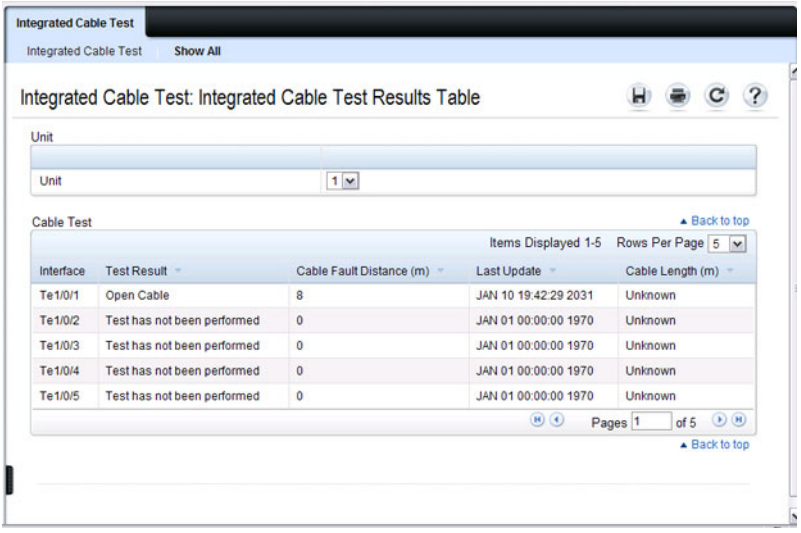
To display the **Integrated Cable Test for Copper Cables** page, click **System → Diagnostics → Integrated Cable Test** in the navigation panel.

Figure 11-4. Integrated Cable Test for Copper Cables



To view a summary of all integrated cable tests performed, click the **Show All** link.

Figure 11-5. Integrated Cable Test Summary



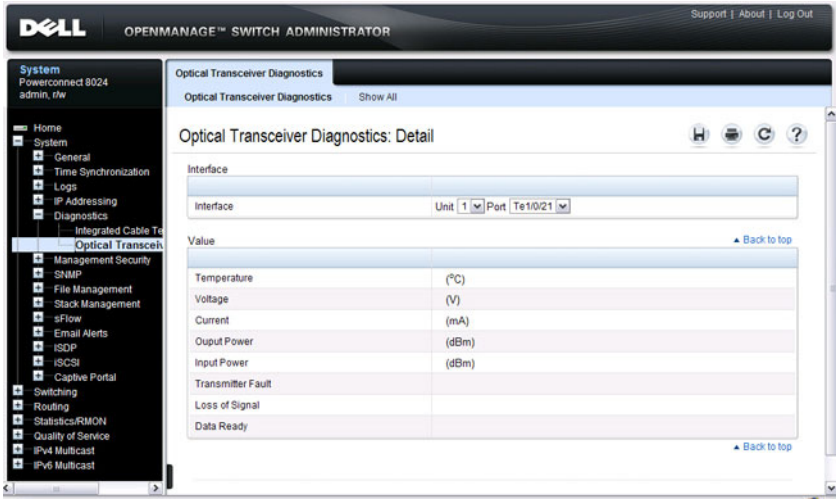
Optical Transceiver Diagnostics

Use the **Optical Transceiver Diagnostics** page to perform tests on Fiber Optic cables.

To display the **Optical Transceiver Diagnostics** page, click **System** → **Diagnostics** → **Optical Transceiver Diagnostics** in the navigation panel.

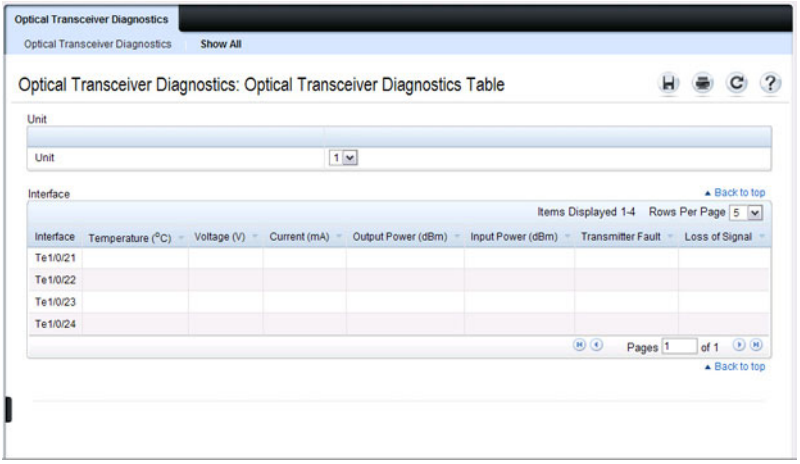
NOTE: Optical transceiver diagnostics can be performed only when the link is present.

Figure 11-6. Optical Transceiver Diagnostics



To view a summary of all optical transceiver diagnostics tests performed, click the **Show All** link.

Figure 11-7. Optical Transceiver Diagnostics Summary



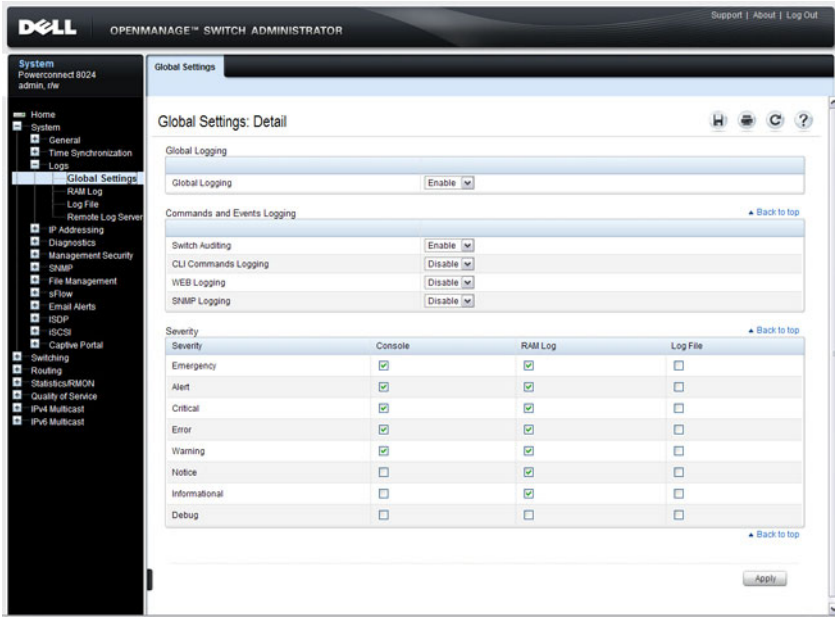
Log Global Settings

Use the **Global Settings** page to enable logging globally, to enable other types of logging. You can also specify the severity of messages that are logged to the console, RAM log, and flash-based log file.

The **Severity** table lists log messages from the highest severity (Emergency) to the lowest (Debug). When you select a severity level, all higher levels are automatically selected. To prevent log messages from being sent to the console, RAM log, or flash log file, clear all check boxes in the **Severity** column.

To display the **Global Settings** page, click **System** → **Logs** → **Global Settings** in the navigation panel.

Figure 11-8. Global Settings

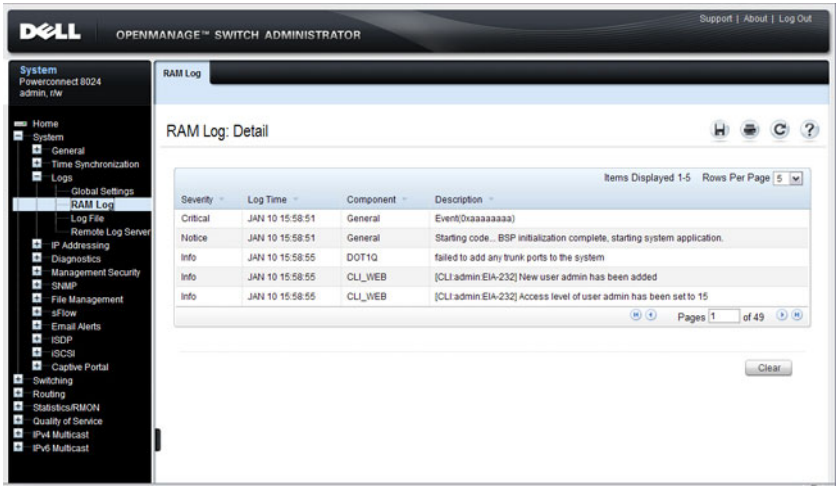


RAM Log

Use the **RAM Log** page to view information about specific RAM (cache) log entries, including the time the log was entered, the log severity, and a description of the log.

To display the **RAM Log**, click **System → Logs → RAM Log** in the navigation panel.

Figure 11-9. RAM Log Table

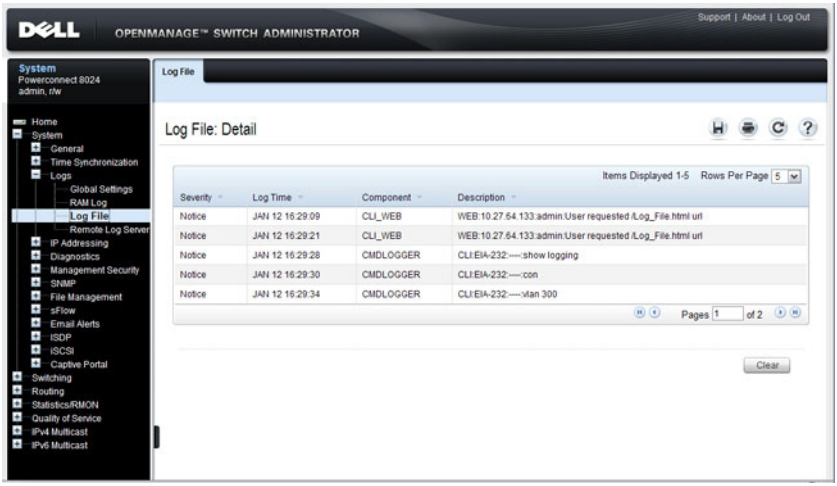


Log File

The **Log File** contains information about specific log entries, including the time the log was entered, the log severity, and a description of the log.

To display the **Log File**, click **System** → **Logs** → **Log File** in the navigation panel.

Figure 11-10. Log File

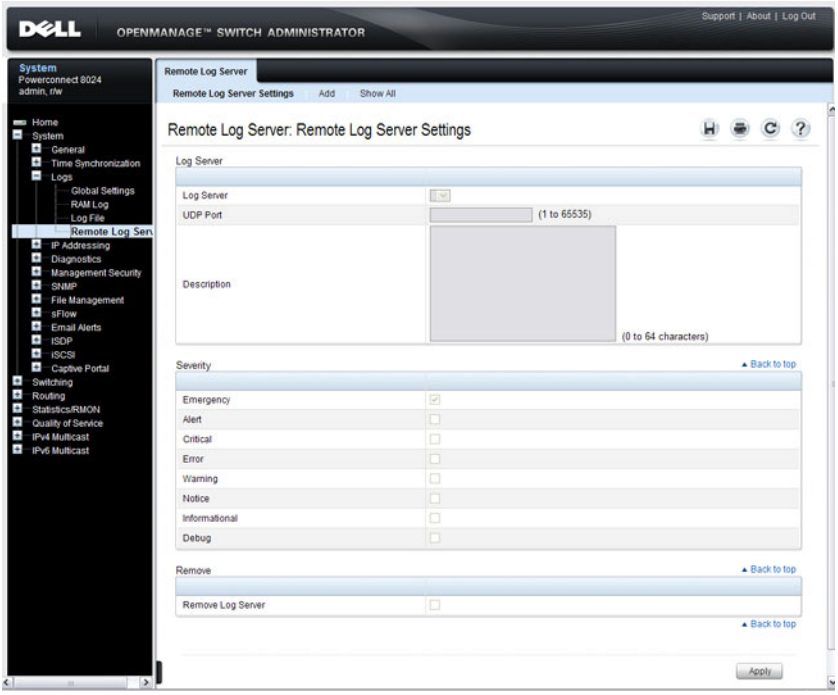


Remote Log Server

Use the **Remote Log Server** page to view and configure the available log servers, to define new log servers, and to set the severity of the log events sent to the server.

To display the **Remote Log Server** page, click **System** → **Logs** → **Remote Log Server**.

Figure 11-11. Remote Log Server

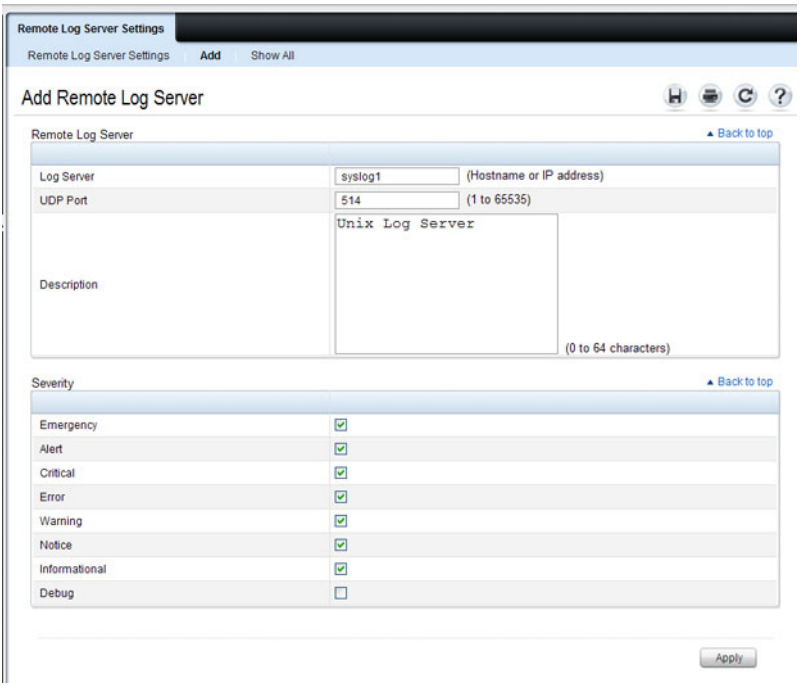


Adding a New Remote Log Server

To add a log server:

1. Open the **Remote Log Server** page.
2. Click **Add** to display the **Add Remote Log Server** page.
3. Specify the IP address or hostname of the remote server.
4. Define the **UDP Port** and **Description** fields.

Figure 11-12. Add Remote Log Server



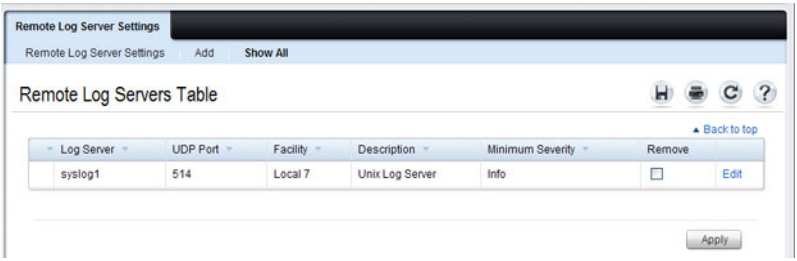
5 Select the severity of the messages to send to the remote server.

NOTE: When you select a severity level, all higher severity levels are automatically selected.

6 Click **Apply**.

Click the **Show All** link to view or remove remote log servers configured on the system.

Figure 11-13. Show All Log Servers

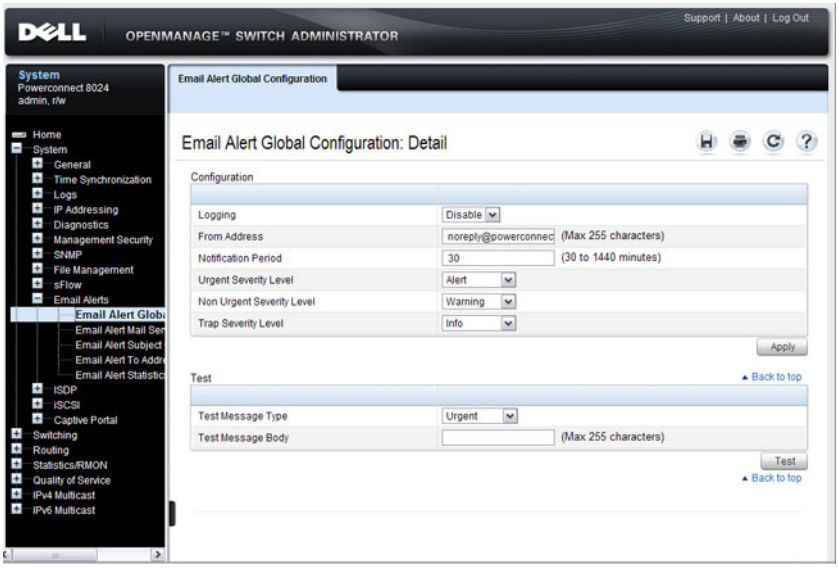


Email Alert Global Configuration

Use the **Email Alert Global Configuration** page to enable the email alerting feature and configure global settings so that system log messages can be sent to from the switch to one or more email accounts.

To display the **Email Alert Global Configuration** page, click **System** → **Email Alerts** → **Email Alert Global Configuration** in the navigation panel.

Figure 11-14. Email Alert Global Configuration

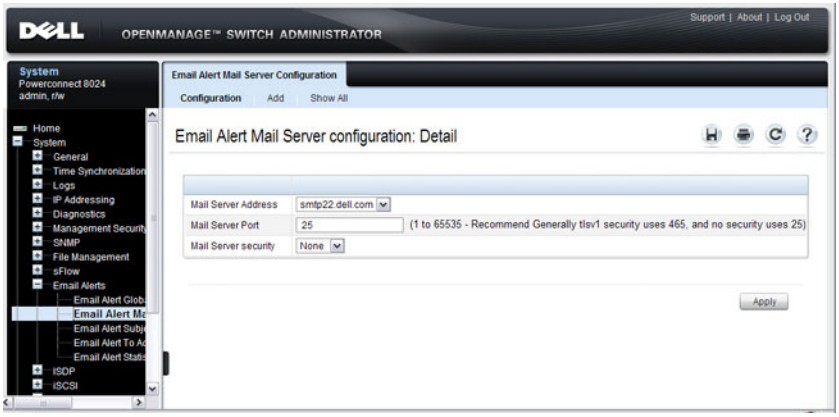


Email Alert Mail Server Configuration

Use the **Email Alert Mail Server Configuration** page to configure information about the mail server the switch uses for sending email alert messages.

To display the **Email Alert Mail Server Configuration** page, click **System** → **Email Alerts** → **Email Alert Mail Server Configuration** in the navigation panel.

Figure 11-15. Email Alert Mail Server Configuration

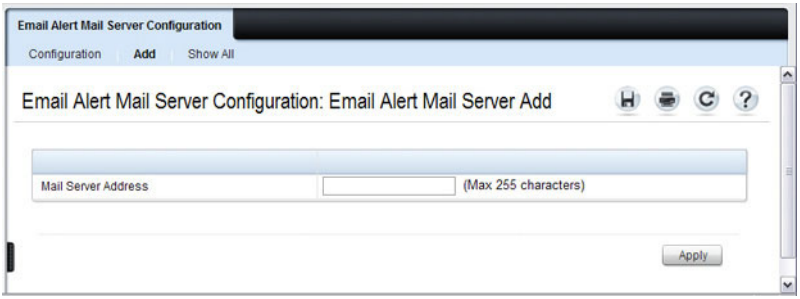


Adding a Mail Server

To add a mail server:

1. Open the **Email Alert Mail Server Configuration** page.
2. Click **Add** to display the **Email Alert Mail Server Add** page.
3. Specify the hostname of the mail server.

Figure 11-16. Add Mail Server

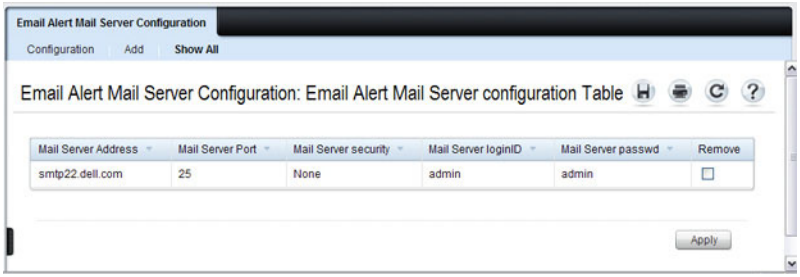


4 Click **Apply**.

5 If desired, click **Configuration** to return to the **Email Alert Mail Server Configuration** page to specify port and security settings for the mail server.

Click the **Show All** link to view or remove mail servers configured on the switch.

Figure 11-17. Show All Mali Servers

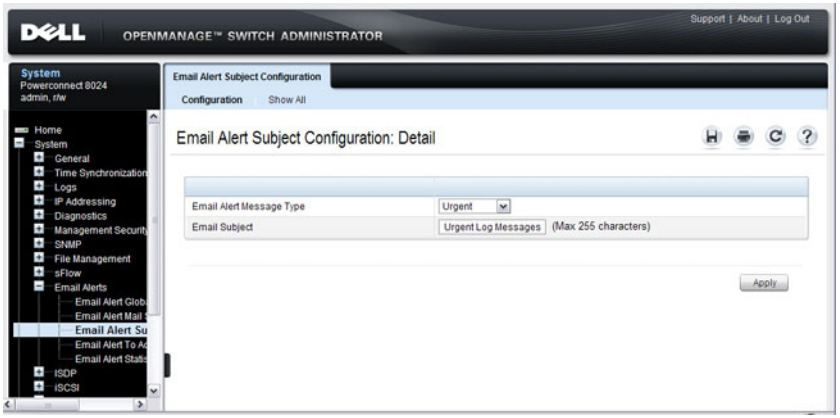


Email Alert Subject Configuration

Use the **Email Alert Subject Configuration** page to configure the subject line for email alerts that are sent by the switch. You can customize the subject for the message severity and entry status.

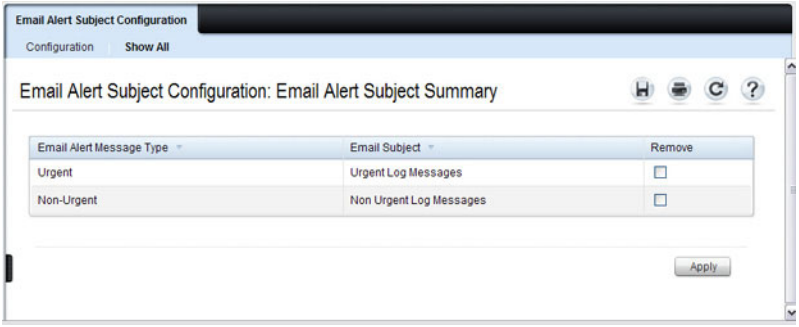
To display the **Email Alert Subject Configuration** page, click **System → Email Alerts → Email Alert Subject Configuration** in the navigation panel.

Figure 11-18. Email Alert Subject Configuration



To view all configured email alert subjects, click the **Show All** link.

Figure 11-19. View Email Alert Subjects

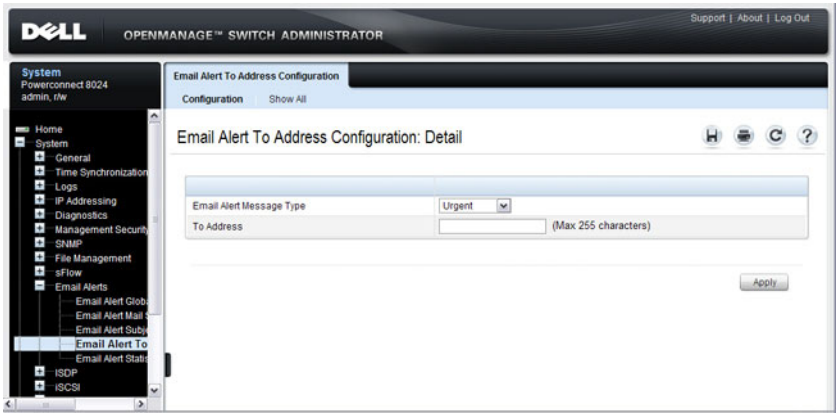


Email Alert To Address Configuration

Use the **Email Alert To Address Configuration** page to specify where the email alerts are sent. You can configure multiple recipients and associate different message severity levels with different recipient addresses.

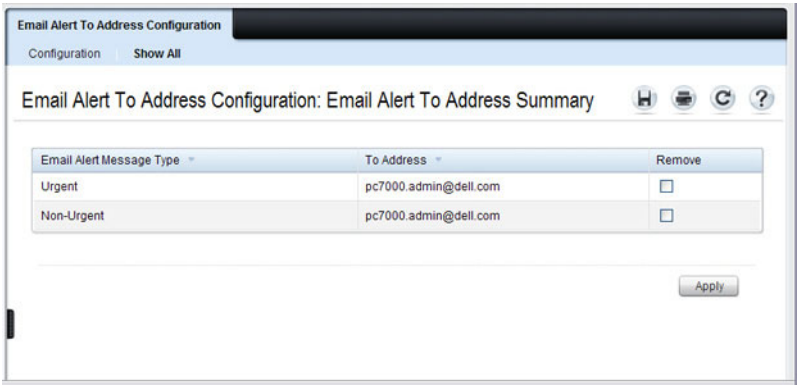
To display the **Email Alert To Address Configuration** page, click **System** → **Email Alerts** → **Email Alert To Address Configuration** in the navigation panel.

Figure 11-20. Email Alert To Address Configuration



To view configured recipients, click the **Show All** link.

Figure 11-21. View Email Alert To Address Configuration

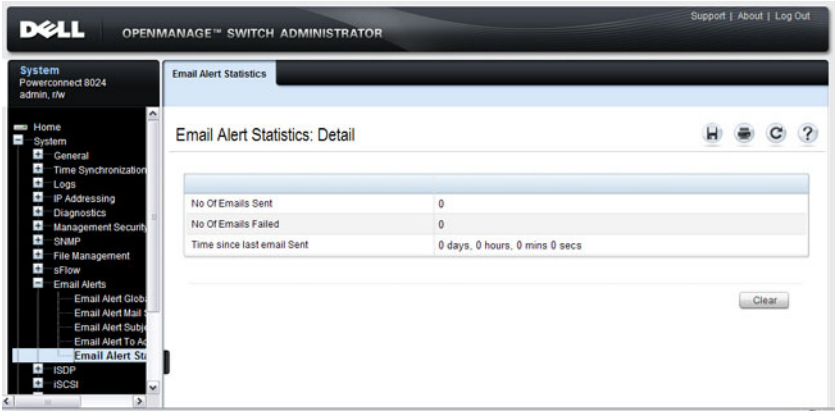


Email Alert Statistics

Use the **Email Alert Statistics** page to view the number of emails that were successfully and unsuccessfully sent, and when emails were sent.

To display the **Email Alert Statistics** page, click **System** → **Email Alerts** → **Email Alert Statistics** in the navigation panel.

Figure 11-22. Email Alert Statistics



Monitoring System Information and Configuring Logging (CLI)

This section provides information about the commands you use to configure information you use to monitor the PowerConnect 8000-series and 8100-series switches. For more information about these commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Viewing System Information

Beginning in Privileged EXEC mode, use the following commands to view system health and resource information.

Command	Purpose
show system	Display various system information.
show system power	Displays the power supply status.
show system temperature	Displays the system temperature and fan status.
show memory cpu	Displays the total and available RAM space on the switch.
show process cpu	Displays the CPU utilization for each process currently running on the switch.

Running Cable Diagnostics

Beginning in Privileged EXEC mode, use the following commands to run the cable diagnostic tests.

NOTE: Cable diagnostics may give misleading results if green mode is enabled on the port. Disable green mode prior to running any cable diagnostics.

Command	Purpose
test copper-port tdr *interface*	Perform the Time Domain Reflectometry (TDR) test to diagnose the quality and characteristics of a copper cable attached to the specified port. **CAUTION: Issuing the test copper-port tdr command will bring the interface down.** The interface is specified in unit/slot/port format.
show copper-ports tdr [*interface*]	Display diagnostic information about all ports or a specified port.
show fiber-ports optical-transceiver [*interface*]	Display the optical transceiver diagnostics for all ports. Include the *interface* option to show information for the specified port.

Configuring Local Logging

Beginning in Privileged EXEC mode, use the following commands to configure the type of messages that are logged and where the messages are logged locally.

Command	Purpose
configure	Enter Global Configuration mode.
logging on	Globally enables logging.
logging audit	Enable switch auditing.
logging cli-command	Enable CLI command logging
logging web-sessions	Enable logging of the switch management Web page visits.
logging snmp	Enable logging of SNMP **set** commands.

Command	Purpose
logging {buffered|console| file} [*severity*]	Enable logging to the specified file. Optionally, you can define a logging discriminator to help filter log messages and set the severity of the messages to log. • **buffered** — Enables logging to the RAM file (cache). If the switch resets, the buffered logs are cleared. • **console** — Enables logging to the screen when you are connected to the CLI through the console port. • **file** — Enables logging to the startup and operational log files on the flash. • **discriminator** *disc-name* — (Optional) Include a message discriminator to help filter log messages. The *disc-name* can contain up to eight alphanumeric characters. Spaces are not permitted. • *severity* — (Optional) Enter the number or name of the desired severity level. For information about severity levels, see Table 11-1.
logging facility *facility-type*	Set the facility for logging messages. Permitted *facility-type* values are local0, local1, local2, local3, local4, local5, local 6, local7
CTRL + Z	Exit to Privileged EXEC mode.
show logging	Displays the state of logging and the syslog messages stored in the internal buffer.
show logging file	View information about the flash (persistent) log file.
clear logging	Use to clear messages from the logging buffer.

Configuring Remote Logging

Beginning in Privileged EXEC mode, use the following commands to define a remote server to which the switch sends log messages.

Command	Purpose
configure	Enter Global Configuration mode.
logging {*ip-address*| *hostname*}	Define a remote log server and enter the configuration mode for the specified log server.
description *description*	Describe the log server. Use up to 64 characters. If the description includes spaces, surround it with quotation marks.
level *severity*	Specify the severity level of the logs that should be sent to the remote log server. For information about severity levels, see Table 11-1.
port *udp-port*	Specify the UDP port to use for sending log messages. The range is 1 to 65535, and the default is 514.
CTRL + Z	Exit to Privileged EXEC mode.
show syslog-servers	Verify the remote log server configuration.

Configuring Mail Server Settings

Beginning in Privileged EXEC mode, use the following commands to configure information about the mail server (SMTP host) on the network that will initially receive the email alerts from the switch and relay them to the correct recipient.

Command	Purpose
configure	Enter Global Configuration mode.
mail-server *ip-address*	Specify the IP address of the SMTP server on the network and enter the configuration mode for the mail server.
security {tlsv1|none}	(Optional) Specify the security protocol to use with the mail server.
port {25|465}	Configure the TCP port to use for SMTP, which can be 25 (SMTP) or 465 (SMTP over SSL).
username *username*	If the SMTP server requires authentication, specify the username to use for the switch. The same username and password settings must be configured on the SMTP host.
password *password*	If the SMTP server requires authentication from clients, specify the password to associate with the switch username.
CTRL + Z	Exit to Privileged EXEC mode.
show mail-server all config	View mail server configuration information for all configured mail servers.

Configuring Email Alerts for Log Messages

Beginning in Privileged EXEC mode, use the following commands to configure email alerts so that log messages are sent to the specified address.

Command	Purpose
configure	Enter Global Configuration mode.
logging email [*severity*]	Enable email alerting and determine which non-critical log messages should be emailed. Including the *severity* value sets the lowest severity for which log messages are emailed. These messages are collected and sent in a single email at the configured log duration. *severity* — (Optional) Enter the number or name of the severity level for non-critical messages. Log messages at or above this severity level are emailed. For information about severity levels, see Table 11-1. Log messages below the specified level are not emailed.
logging email urgent {*severity* | **none**}	Determine which log messages are critical and should be sent in a single email as soon as they are generated. *severity* — (Optional) Enter the number or name of the severity level for critical messages. For information about severity levels, see Table 11-1.
logging email logtime *minutes*	Specify how often to send the non-critical email alerts that have been collected. . The valid range is 30 - 1440 minutes.
logging email message-type {urgent | non-urgent | both} to-addr *email-address*	Specify the email address of the recipient for log messages.
logging email from-addr *email-address*	Specify the email address of the sender, which is the switch.
logging email message-type {urgent | non-urgent | both} subject *subject*	Specify the text that will appear in the subject line of email alerts sent by the switch.

Command	Purpose
logging email test message-type {urgent | non-urgent | both} message-body *body*	Send a test email to the configured recipient to verify that the feature is properly configured.
CTRL + Z	Exit to Privileged EXEC mode.
show logging email config	View the configured settings for email alerts.
show logging email statistics	View information about the number of emails sent and the time they were sent.
clear logging email statistics	Clear the email alerting statistics.

Logging Configuration Examples

This section contains the following examples:

- Configuring Local and Remote Logging
- Configuring Email Alerting

Configuring Local and Remote Logging

This example shows how to enable switch auditing and CLI command logging. Log messages with a severity level of Notification (level 5) and above are sent to the RAM (buffered) log. Emergency, Critical, and Alert (level 2) log messages are written to the log file on the flash drive. All log messages are displayed on the console and sent to a remote syslog server.

To configure the switch:

1. Enable switch auditing and CLI command logging.

```
console#configure
console(config)#logging audit
console(config)#logging cli-command
```

2. Specify where the logs are sent locally and what severity level of message is to be logged. You can specify the severity as the level number, as shown in the first two commands, or as the keyword, shown in the third command.

```
console(config)#logging buffered 5
console(config)#logging file 2
console(config)#logging console debugging
```

3. Define the remote log server.

```
console(config)#logging 192.168.2.10
console(Config-logging)#description "Syslog
Server"
console(Config-logging)#level debug
console(Config-logging)#exit
console(config)#exit
```

4 Verify the remote log server configuration.

```
console#show syslog-servers

IP Address/Hostname        Port    Severity        Description
-------------------------- ------  --------------  ----------
192.168.2.10               514     debugging       Syslog Server
```

5 Verify the local logging configuration and view the log messages stored in the buffer (RAM log).

```
console#show logging

Logging is enabled
Console Logging: level debugging. Console
Messages: 748 Dropped.
Buffer Logging: level notifications. Buffer
Messages: 79 Logged,
File Logging: level critical. File Messages: 973
Dropped.
CLI Command Logging : enabled
Switch Auditing : enabled
Web Session Logging : disabled
SNMP Set Command Logging : disabled
Syslog server  192.168.2.10 logging: debug.
Messages: 0 dropped
412 Messages dropped due to lack of resources.
Buffer Log:
<186> FEB 02 05:53:03 0.0.0.0-1 UNKN[1073741088]:
bootos.c(232) 1 %% Event(0xaaaaaaaa)
<189> FEB 02 05:53:03 0.0.0.0-1 UNKN[1073741088]:
bootos.c(248) 2 %% Starting code... BSP
initialization complete, starting application.

--More-- or (q)uit
```

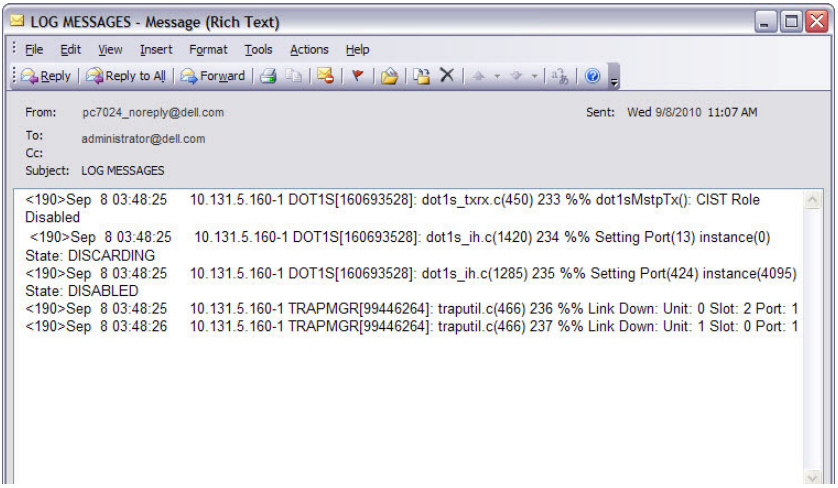
Configuring Email Alerting

The commands in this example define the SMTP server to use for sending email alerts. The mail server does not require authentication and uses the standard TCP port for SMTP, port 25, which are the default values. Only

Emergency messages (severity level 0) will be sent immediately as individual emails, and messages with a severity of alert, critical, and error (levels 1-3) will be sent in a single email every 120 minutes. Warning, notice, info, and debug messages are not sent in an email.

The email the administrator will in the inbox has a format similar to the following:

Figure 11-23. Email Alert Message Format



For emergency-level messages, the subject is LOG MESSAGE - EMERGENCY. For messages with a severity level of alert, critical, and error, the subject is LOG MESSAGE.

To configure the switch:

1. Specify the mail server to use for sending messages.

```
console#configure
console(config)#mail-server ip-address 192.168.2.34
```

2. Configure the username and password for the switch must use to authenticate with the mail server.

```
console(Mail-Server)#username switch8024
console(Mail-Server)#password password8024
console(Mail-Server)#exit
```

3 Configure emergencies and alerts to be sent immediately, and all other messages to be sent in a single email every 120 minutes.

```
console(config)#logging email error
console(config)#logging email urgent emergency
console(config)#logging email logtime 120
```

4 Specify the email address of the sender (the switch).

```
console(config)#logging email from-addr
pc8024_noreply@dell.com
```

5 Specify the address where email alerts should be sent.

```
console(config)#logging email message-type both
to-addr administrator@dell.com
```

6 Specify the text that will appear in the email alert Subject line.

```
console(config)#logging email message-type urgent
subject "LOG MESSAGES - EMERGENCY"
console(config)#logging email message-type non-
urgent subject "LOG MESSAGES"
```

7 Verify the configuration.

```
console#show mail-server all config

Mail Servers Configuration:

No of mail servers configured.................. 1

Email Alert Mail Server Address................ 192.168.2.34
Email Alert Mail Server Port................... 25
Email Alert SecurityProtocol................... none
Email Alert Username........................... switch8024
Email Alert Password........................... password8024
```

```
console#show logging email config

Email Alert Logging............................ enabled
Email Alert From Address.......................
pc8024_noreply@dell.com
Email Alert Urgent Severity Level.............. 0
Email Alert Non Urgent Severity Level.......... 3
Email Alert Trap Severity Level................ 6
Email Alert Notification Period................ 120 min
```

```
Email Alert To Address Table:

For Msg Type.....................1

Address1.........................administrator@dell.com

For Msg Type.....................2

Address1.........................administrator@dell.com

Email Alert Subject  Table      :

For Msg Type 1, subject is.....LOG MESSAGES - EMERGENCY

For Msg Type 2, subject is.....LOG MESSAGE
```

12

Managing General System Settings

This chapter describes how to set system information, such as the hostname, and time settings, and how to select the Switch Database Management (SDM) template to use on the switch.

The topics covered in this chapter include:

- System Settings Overview
- Default General System Information
- Configuring General System Settings (Web)
- Configuring System Settings (CLI)
- General System Settings Configuration Examples

System Settings Overview

The system settings include the information described in Table 12-1. This information helps identify the switch.

Table 12-1. System Information

Feature	Description
System Name	The switch name (host name). If you change the system name, the CLI prompt changes from `console` to the system name.
System contact	Identifies the person to contact for information regarding the switch.
System location	Identifies the physical location of the switch.
Asset tag	Uniquely identifies the switch. Some organizations use asset tags to identify, control, and track each piece of equipment.
CLI Banner	Displays a message upon connecting to the switch or logging on to the switch by using the CLI.
SDM Template	Determines the maximum resources a switch or router can use for various features. For more information, see "What Are SDM Templates?" on page 239

The switch can obtain the time from a Simple Network Time Protocol (SNTP) server, or you can set the time manually. Table 12-2 describes the settings that help the switch keep track of time.

Table 12-2. Time Settings

Feature	Description
SNTP	Controls whether the switch obtains its system time from an SNTP server and whether communication with the SNTP server requires authentication and encryption. You can configure information for up to eight SNTP servers. The SNTP client on the switch can accept updates from both IPv4 and IPv6 SNTP servers.
Real time clock (RTC)	If SNTP is disabled, you can manually enter the system time and date.
Time Zone	Allows you to specify the offset from Coordinated Universal Time (UTC), which is also known as Greenwich Mean Time (GMT).
Summer Time	In some regions, the time shifts by one hour in the fall and spring. In the United States, this is called daylight saving time.

Why Does System Information Need to Be Configured?

Configuring system information is optional. However, it can be helpful in providing administrative information about the switch. For example, if you manage several standalone PowerConnect 8000-series and 8100-series switches and have Telnet sessions open with several different switches, the system name can help you quickly identify the switch because the host name replaces `console` as the CLI command prompt.

The Banner can provide information about the switch status. For example, if multiple users connect to the switch, the message of the day (MOTD) banner might alert everyone who connects to the switch about a scheduled switch image upgrade.

What Are SDM Templates?

An SDM template is a description of the maximum resources a switch or router can use for various features. Different SDM templates allow different combinations of scaling factors, enabling different allocations of resources depending on how the device is used. In other words, SDM templates enable you to reallocate system resources to support a different mix of features based on your network requirements.

PowerConnect 8000-series and 8100-series switches support the following three templates:

- Dual IPv4 and IPv6 (default)
- IPv4 Routing
- IPv4 Data Center

Table 12-3 describes the parameters that are scaled for each template and the per-template maximum value of the parameter.

Table 12-3. SDM Template Parameters and Values

Parameter	Dual IPv4/IPv6	IPv4 Only	IPv4 Data Center
ARP entries	6144	6144	6144
IPv4 unicast routes	8160	12256	8160
IPv6 Neighbor Discovery Protocol (NDP) entries	2560	0	0
IPv6 unicast routes	4096	0	0
ECMP next hops	4	4	16
IPv4 multicast routes	1536	2048	2048
IPv6 multicast routes	512	0	0

SDM Template Configuration Guidelines

When you configure the switch to use an SDM template that is not currently in use, you must reload the switch for the configuration to take effect.

If the IPv4 Routing or IPv4 Data Center template is currently in use and you attempt to configure IPv6 routing features without first selecting the Dual IPv4-IPv6 Routing template, the IPv6 commands do not take effect. IPv6 features are not available when an IPv4-only template is active.

Why is the System Time Needed?

The switch uses the system clock to provide time stamps on log messages. Additionally, some **show** commands include the time in the command output. For example, the **show users login-history** command includes a Login Time field. The system clock provides the information for the Login Time field.

How Does SNTP Work?

SNTP assures accurate switch clock time synchronization. Time synchronization is performed by a network SNTP server.

Time sources are established by Stratums. Stratums define the accuracy of the reference clock. The higher the stratum (where zero is the highest), the more accurate the clock. The switch is at a stratum that is one lower than its time source. For example, if the SNTP server in an internal network is a Stratum 3 device, the switch is a Stratum 4 device.

You can configure the switch to request the time from an SNTP server on the network, or you can allow the switch to receive SNTP broadcasts.

Requesting the time from a unicast SNTP server is more secure. Use this method if you know the IP address of the SNTP server on your network. If you allow the switch to receive SNTP broadcasts, any clock synchronization information is accepted, even if it has not been requested by the device. This method is less secure than polling a specified SNTP server.

To increase security, you can require authentication between the configured SNTP server and the SNTP client on the switch. Authentication is provided by Message Digest 5 (MD5). MD5 verifies the integrity of the communication and authenticates the origin of the communication.

Default General System Information

By default, no system information or time information is configured, and the SNTP client is disabled. The default SDM Template applied to the switch is the Dual IPv4-IPv6 template.

Configuring General System Settings (Web)

This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring general system settings on the PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

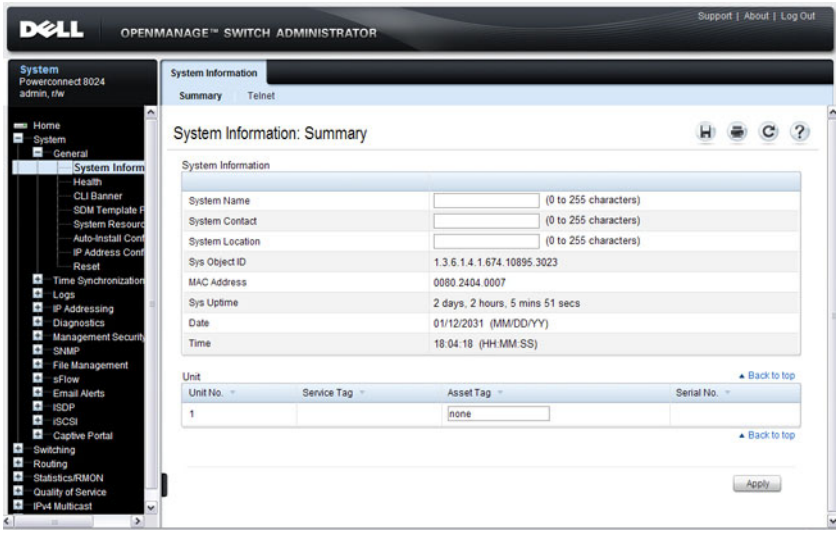
System Information

Use the **System Information** page to configure the system name, contact name, location, and asset tag.

NOTE: From the **System Information** page, you can also initiate a Telnet session to the switch.

To display the **System Information** page, click **System** → **General** → **System Information** in the navigation panel.

Figure 12-1. System Information



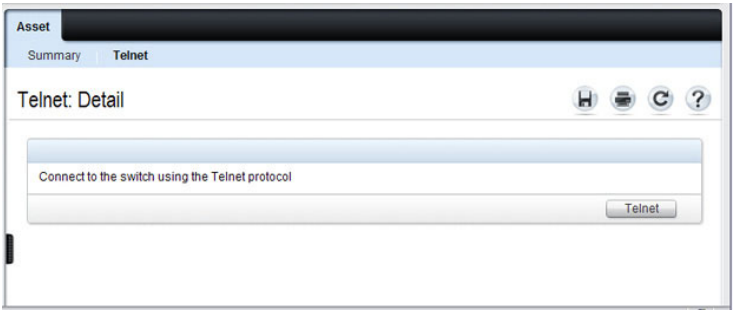
Initiating a Telnet Session from the Web Interface

NOTE: The Telnet client feature does not work with Microsoft Windows Internet Explorer 7 and later versions. Initiating this feature from any browser running on a Linux operating system is not supported.

To launch a Telnet session:

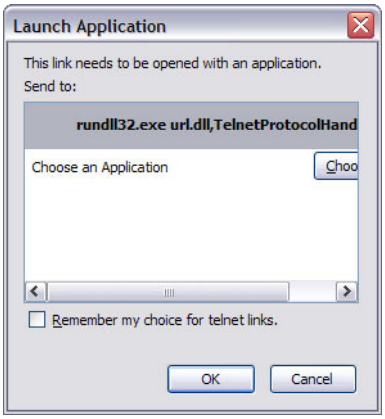
1. From the **System → General → System Information** page, click the Telnet link.
2. Click the **Telnet** button.

Figure 12-2. Telnet



3. Select the Telnet client, and click **OK**.

Figure 12-3. Select Telnet Client



The selected Telnet client launches and connects to the switch CLI.

Figure 12-4. Telnet Session

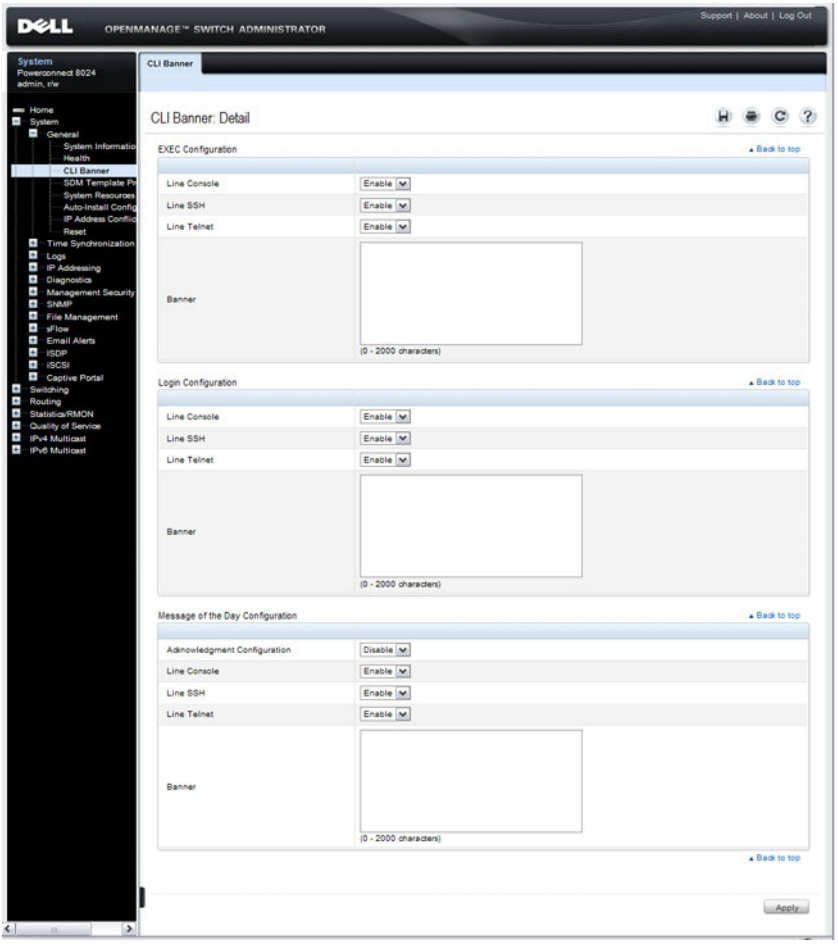


CLI Banner

Use the **CLI Banner** page to configure a message for the switch to display when a user connects to the switch by using the CLI. You can configure different banners for various CLI modes and access methods.

To display the **CLI Banner** page, click **System → General → CLI Banner** in the navigation panel.

Figure 12-5. CLI Banner

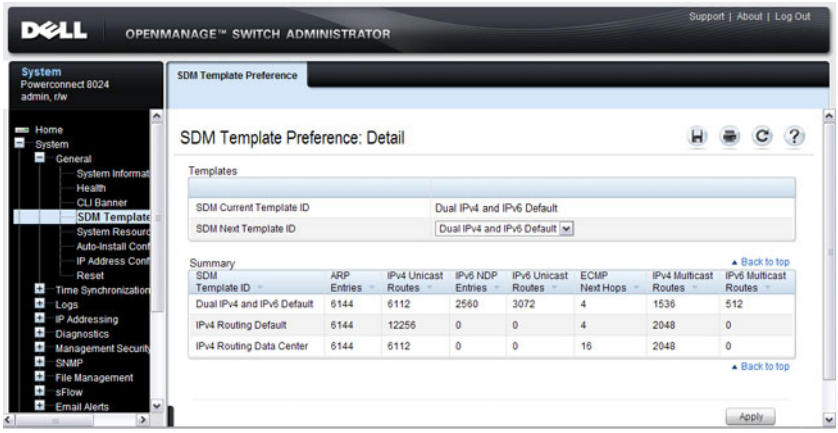


SDM Template Preference

Use the **SDM Template Preference** page to view information about template resource settings and to select the template that the switch uses. If you select a new SDM template for the switch to use, you must reboot the switch before the template is applied.

To display the **SDM Template Preference** page, click **System** → **General** → **SDM Template Preference** in the navigation panel.

Figure 12-6. SDM Template Preference

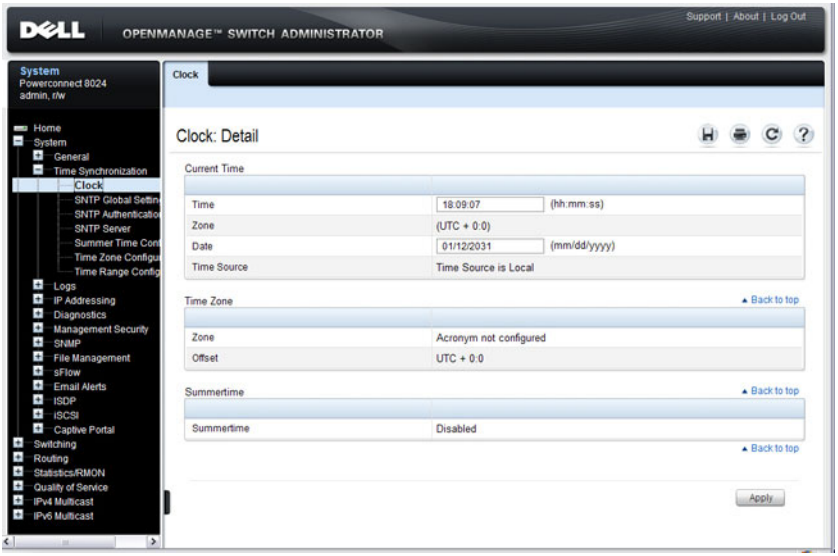


Clock

If you do not obtain the system time from an SNTP server, you can manually set the date and time on the switch on the **Clock** page. The **Clock** page also displays information about the time settings configured on the switch.

To display the **Clock** page, click **System** → **Time Synchronization** → **Clock** in the navigation panel.

Figure 12-7. Clock



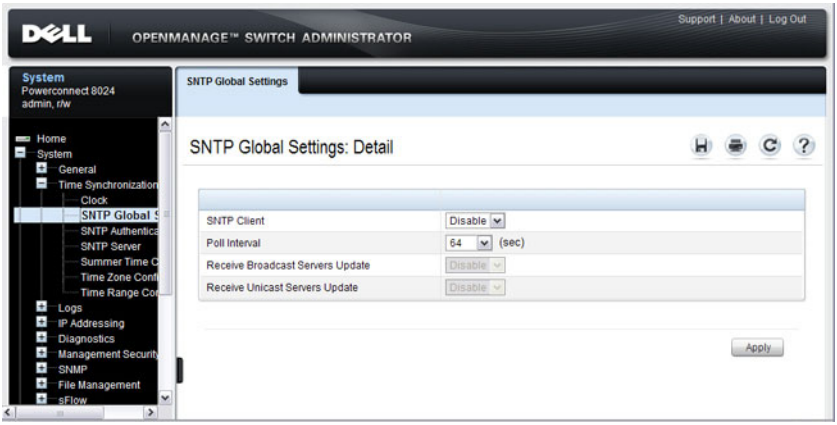
NOTE: The system time cannot be set manually if the SNTP client is enabled. Use the **SNTP Global Settings** page to enable or disable the SNTP client.

SNTP Global Settings

Use the **SNTP Global Settings** page to enable or disable the SNTP client, configure whether and how often the client sends SNTP requests, and determine whether the switch can receive SNTP broadcasts.

To display the SNTP Global Settings page, click **System** → **Time Synchronization** → **SNTP Global Settings** in the navigation panel.

Figure 12-8. SNTP Global Settings



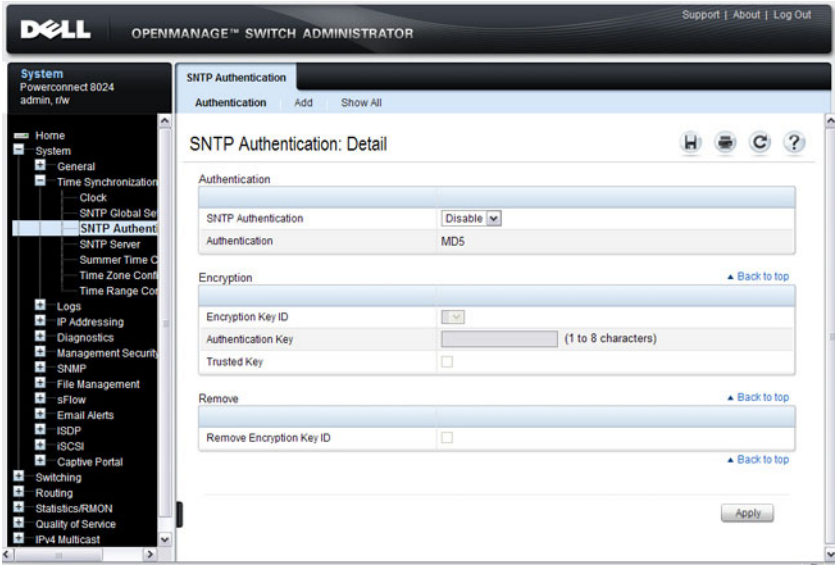
SNTP Authentication

Use the **SNTP Authentication** page to enable or disable SNTP authentication, to modify the authentication key for a selected encryption key ID, to designate the selected authentication key as a trusted key, and to remove the selected encryption key ID.

NOTE: The SNTP server must be configured with the same authentication information to allow time synchronization to take place between the two devices.

Click **System** → **Time Synchronization** → **SNTP Authentication** in the navigation panel to display the **SNTP Authentication** page.

Figure 12-9. SNTP Authentication



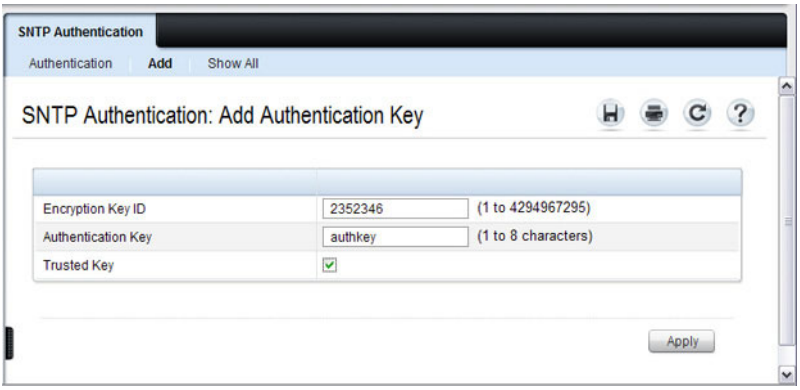
Adding an SNTP Authentication Key

To configure SNTP authentication:

1. Open the **SNTP Authentication** page.
2. Click the **Add** link.

The **Add Authentication Key** page displays:

Figure 12-10. Add Authentication Key



3 Enter a numerical encryption key ID and an authentication key in the appropriate fields.

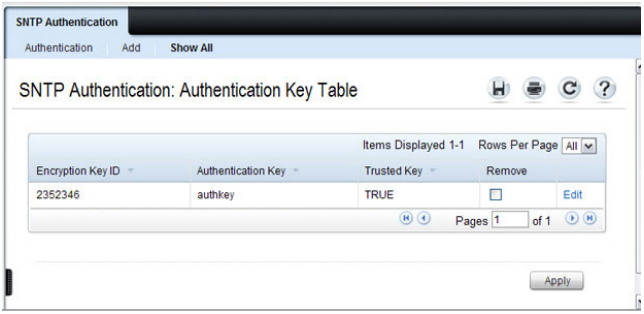
4 If the key is to be used to authenticate a unicast SNTP server, select the **Trusted Key** check box. If the check box is clear, the key is untrusted and cannot be used for authentication.

5 Click **Apply**.

The SNTP authentication key is added, and the device is updated.

To view all configured authentication keys, click the **Show All** link. The **Authentication Key Table** displays. You can also use the **Authentication Key Table** to remove or edit existing keys.

Figure 12-11. Authentication Key Table

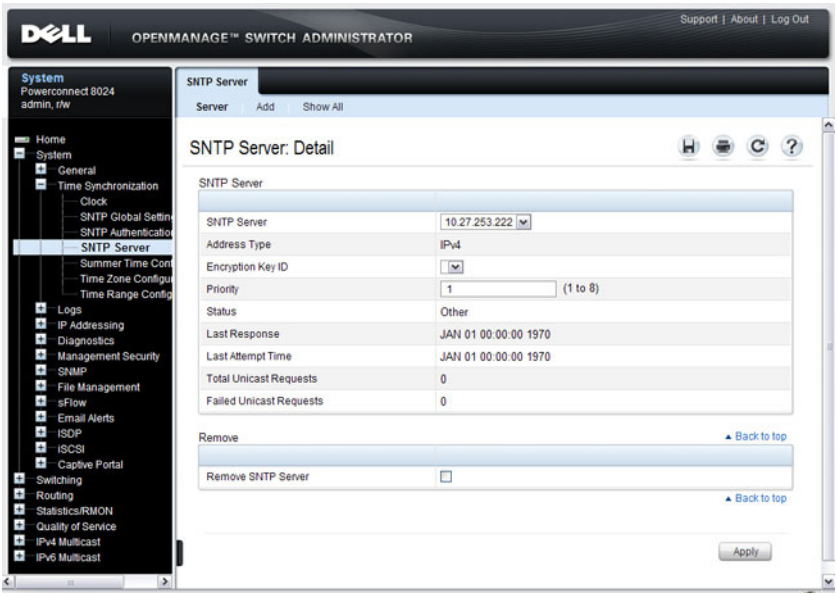


SNTP Server

Use the **SNTP Server** page to view and modify information about SNTP servers, and to add new SNTP servers that the switch can use for time synchronization. The switch can accept time information from both IPv4 and IPv6 SNTP servers.

To display the **SNTP Server** page, click **System → Time Synchronization → SNTP Server** in the navigation panel. If no servers have been configured, the fields in the following image are not displayed.

Figure 12-12. SNTP Servers



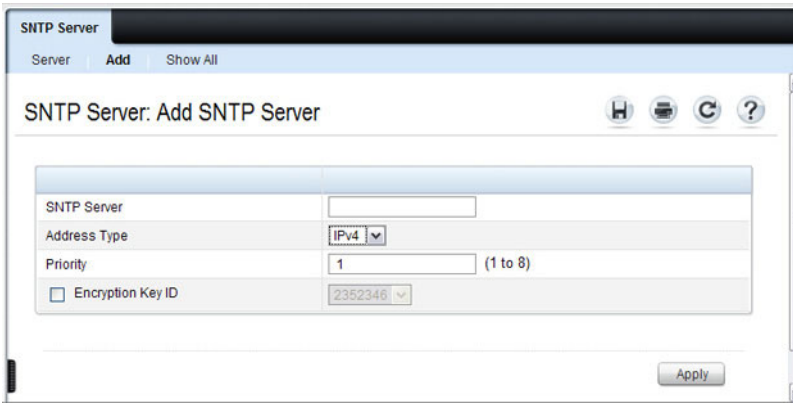
Defining a New SNTP Server

To add an SNTP server:

1. Open the **SNTP Servers** page.
2. Click **Add**.

   The **Add SNTP Server** page displays.

Figure 12-13. Add SNTP Server



3 In the **SNTP Server** field, enter the IP address or host name for the new SNTP server.

4 Specify whether the information entered in the **SNTP Server** field is an IPv4 address, IPv6 address, or a hostname (DNS).

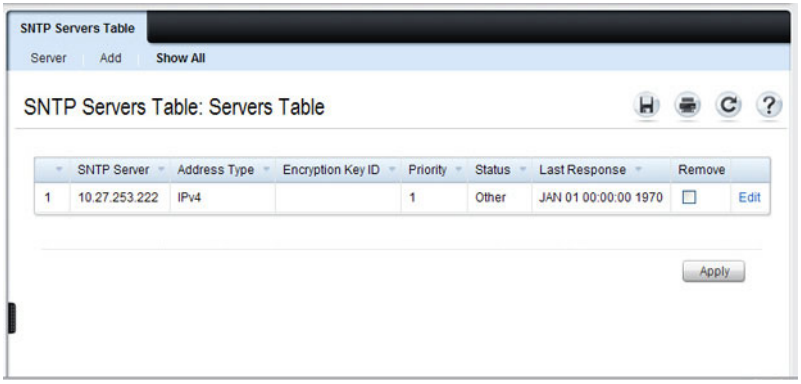
5 If you require authentication between the SNTP client on the switch and the SNTP server, select the Encryption Key ID check box, and then select the key ID to use.

To define a new encryption key, see "Adding an SNTP Authentication Key" on page 248.

NOTE: The SNTP server must be configured with the same authentication information to allow time synchronization to take place between the two devices.

To view all configured SNTP servers, click the **Show All** link. The SNTP **Server Table** displays. You can also use the SNTP **Server Table** page to remove or edit existing SNTP servers.

Figure 12-14. SNTP Servers Table

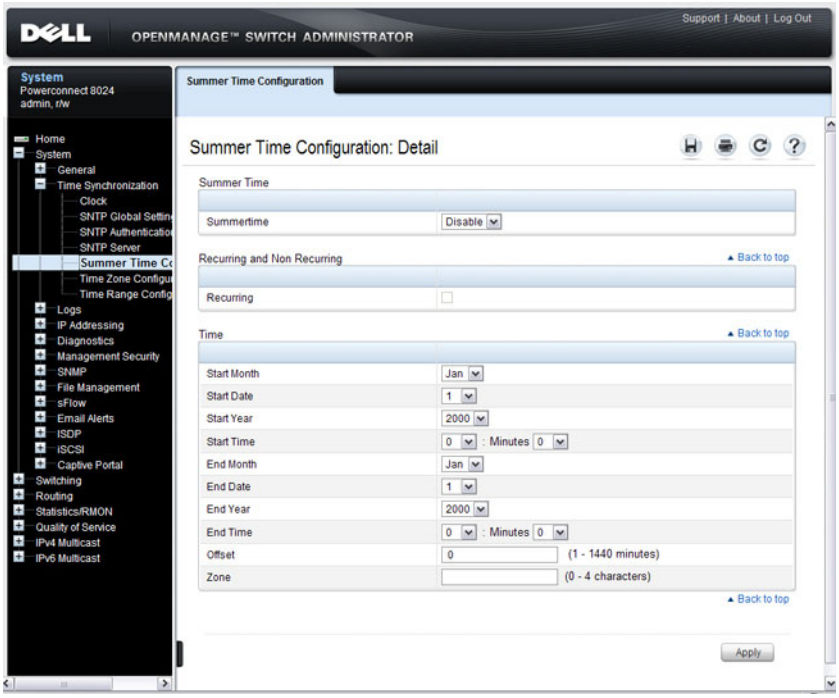


Summer Time Configuration

Use the **Summer Time Configuration** page to configure summer time (daylight saving time) settings.

To display the **Summer Time Configuration** page, click **System** → **Time Synchronization** → **Summer Time Configuration** in the navigation panel.

Figure 12-15. Summer Time Configuration



NOTE: The fields on the **Summer Time Configuration** page change when you select or clear the Recurring check box.

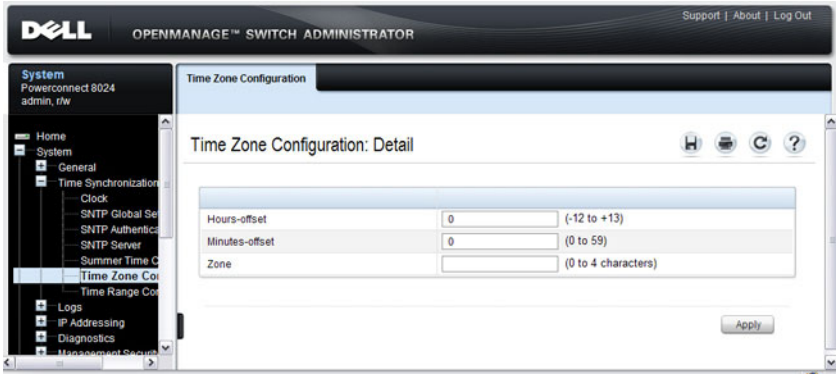
To use the preconfigured summer time settings for the United States or European Union, select the **Recurring** check box and specify USA or EU from the Location menu.

Time Zone Configuration

Use the **Time Zone Configuration** to configure time zone information, including the amount time the local time is offset from UTC and the acronym that represents the local time zone.

To display the **Time Zone Configuration** page, click **System** → **Time Synchronization** → **Time Zone Configuration** in the navigation panel.

Figure 12-16. Time Zone Configuration

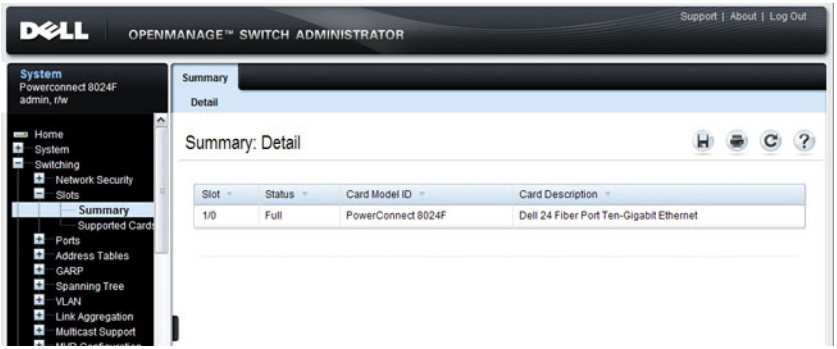


Slot Summary

Use the **Slot Summary** page to view information about the expansion slot status.

To display the **Slot Summary** page, click **Switching** → **Slots** → **Summary** in the navigation panel.

Figure 12-17. Slot Summary

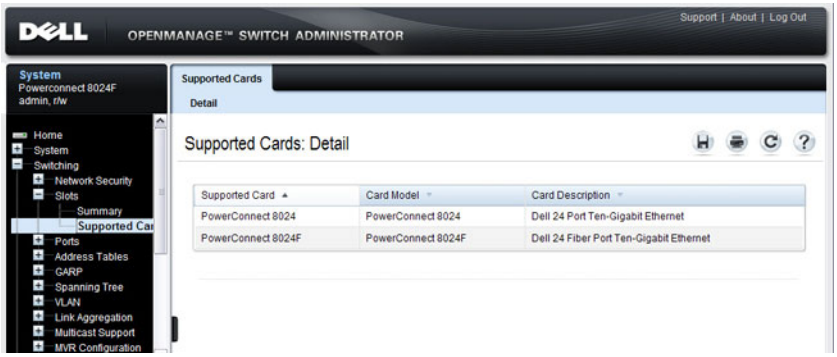


Supported Cards

Use the **Supported Cards** page to view information about the supported plug-in modules for the switch.

To display the **Supported Cards** page, click **Switching** → **Slots** → **Supported Cards** in the navigation panel.

Figure 12-18. Supported Cards



Configuring System Settings (CLI)

This section provides information about the commands you use to configure system information and time settings on the PowerConnect 8000-series and 8100-series switches. For more information about these commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Configuring System Information

Beginning in Privileged EXEC mode, use the following commands to configure system information.

Command	Purpose
configure	Enter Global Configuration mode.
hostname *name*	Configure the system name. The CLI prompt changes to the host name after you execute the command.
snmp-server contact *name*	Configure the name of the switch administrator. If the name contains a space, use quotation marks around the name.
snmp-server location *location*	Configure the switch location.
asset-tag *tag*	Configure the asset tag for the switch.
CTRL + Z	Exit to Privileged EXEC mode.
show system [id]	Display system information. Include the **id** keyword to display additional system information.

Configuring the Banner

Beginning in Privileged EXEC mode, use the following commands to configure the MOTD, login, or User EXEC banner. The switch supports the following banner messages:

- MOTD—Displays when a user connects to the switch.
- Login—Displays after the MOTD banner and before the login prompt.
- Exec—Displays immediately after the user logs on to the switch.

Command	Purpose
configure	Enter Global Configuration mode.
banner {motd|login|exec} *text*	Configure the banner message that displays when you connect to the switch (motd and login) or enter User EXEC mode (exec). Use quotation marks around a message if it includes spaces.
line {telnet|ssh|console}	Enter the terminal line configuration mode for Telnet, SSH, or the console.
motd-banner	Specify that the configured MOTD banner displays. To prevent the banner from displaying, enter **no motd-banner**.
exec-banner	Specify that the configured exec banner displays. To prevent the banner from displaying, enter **no exec-banner**.
login-banner	Specify that the configured login banner displays. To prevent the banner from displaying, enter **no login-banner**.
CTRL + Z	Exit to Privileged EXEC mode.
show banner	Display the banner status on all line terminals.

Managing the SDM Template

Beginning in Privileged EXEC mode, use the following commands to set the SDM template preference and to view information about the available SDM templates.

Command	Purpose
configure	Enter Global Configuration mode.
sdm prefer {dual-ipv4-and-ipv6 default| ipv4-routing {data-center | default}}	Select the SDM template to apply to the switch after the next boot.
CTRL + Z	Exit to Privileged EXEC mode.
show sdm prefer [*template*]	View information about the SDM template the switch is currently using. Use the *template* variable to view the parameters for the specified template.

Configuring SNTP Authentication and an SNTP Server

Beginning in Privileged EXEC mode, use the following commands to require the SNTP client to use authentication when communicating with the SNTP server. The commands also show how to configure an SNTP server.

Requiring authentication is optional. However, if you configure authentication on the switch SNTP client, the SNTP server must be configured with the same authentication information to allow time synchronization to take place between the two devices.

Command	Purpose
configure	Enter Global Configuration mode.
sntp authentication-key *key_id* **md5** *key_word*	Define an authentication key for SNTP. The variables are: • *key_id*— The encryption key ID, which is a number from 1–4294967295. • *key_word*—The authentication key, which is a string of up to eight characters.

Command	Purpose
sntp trusted-key *key_id*	Specify the authentication key the SNTP server must include in SNTP packets that it sends to the switch. The *key_id* number must be an encryption key ID defined in the previous step.
sntp authenticate	Require authentication for communication with the SNTP server. A trusted key must be configured before this command is executed.
sntp server {*ip_address* | *hostname*} [**priority** *priority*] [**key** *key_id*]	Define the SNTP server. • *ip_address*—The IP address (or host name) of the SNTP server to poll. The IP address can be an IPv4 or IPv6 address. • *priority*—(Optional) If multiple SNTP servers are defined, this number determines which server the switch polls first. The priority is 1–8, where 1 is the highest priority. If you do not specify a priority, the servers are polled in the order that they are entered. • *key_id*—(Optional) Enter an authentication key to use. The key must be previously defined by the **sntp authentication-key** command.
sntp {unicast|broadcast} client enable	This command enables the SNTP client and allows the switch to poll configured unicast SNTP servers for updates or receive broadcasts from any SNTP server.
sntp client poll timer *seconds*	Specify how often the SNTP client requests SNTP packets from the configured server(s). *seconds*—The poll interval can be 64, 128, 256, 512, or 1024 seconds.
CTRL + Z	Exit to Privileged EXEC mode.
show sntp configuration	Verify the SNTP configuration.
show sntp status	View information about the SNTP updates.

Setting the System Time and Date Manually

Beginning in Privileged EXEC mode, use the following commands to configure the time and date, time zone, and summer time settings.

Command	Purpose
clock set {*mm/dd/yyyy hh:mm:ss*} | {*hh:mm:ss mm/dd/yyyy*	Configure the time and date. You can enter the time first and then the date, or the date and then the time. • *hh:mm:ss* —Time in hours (24-hour format, from 01-24), minutes (00-59), and seconds (00-59). • *mm/dd/yyyy* — Two digit month (1-12), two-digit date of the month (01-31), and four-digit year.
clock timezone hours-offset *hours-offset* [**minutes** *minutes-offset*] [**zone** *acronym*]	Configure the time zone settings. • *hours-offset* — Hours difference from UTC. (Range: –12 to +13) • *minutes-offset* — Minutes difference from UTC. (Range: 0–59) • *acronym* — The acronym for the time zone. (Range: Up to four characters)
clock summer-time recurring {**usa** | **eu** | {*week day month hh:mm week day month hh:mm*}} [**offset** *offset*] [**zone** *acronym*]	Use this command if the summer time starts and ends every year based on a set pattern. For switches located in the United States or European Union, use the **usa** or **eu** keywords to use the preconfigured values. Otherwise, configure the start and end times by using the following values: • *week* — Week of the month. (Range: 1–5, first, last) • *day* — Day of the week. (The first three letters by name) • *month* — Month. (The first three letters by name; jan, for example.) • *hh:mm* — Time in 24-hour format in hours and minutes. (Range: hh: 0–23, mm: 0–59) • *offset* — Number of minutes to add during the summertime. (Range:1–1440) • *acronym* — The acronym for the time zone to be displayed when summertime is in effect. (Up to four characters)

Command	Purpose
clock summer-time date {*date month* | *month date*} *year hh:mm* {*date month* | *month date*} *year hh:mm* [**offset** *offset*] [**zone** *acronym*]	Use this command if the summer time does not start and end every year according to a recurring pattern. You can enter the month and then the date, or the date and then the month. • *date*— Day of the month. (Range: 1-31.) • *month* — Month. (Range: The first three letters by name) • *hh:mm* — Time in 24-hour format in hours and minutes. (Range: hh: 0–23, mm: 0–59) • *offset* — Number of minutes to add during the summertime. (Range:1–1440) • *acronym* — The acronym for the time zone to be displayed when summertime is in effect. (Range: Up to four characters)
CTRL + Z	Exit to Privileged EXEC mode.
show clock [**detail**]	View information about the time. Include the **detail** keyword to view information about the time zone and summer time.

Viewing Slot Information

Use the following commands to view information about Slot 0 and its support.

Command	Purpose
show slot	Display status information about the expansion slots.
show supported cardtype	Display information about the modules the switch supports.

General System Settings Configuration Examples

This section contains the following examples:

- Configuring System and Banner Information
- Configuring SNTP
- Configuring the Time Manually

Configuring System and Banner Information

In this example, an administrator configures the following system information:

- System name: PC8024
- System contact: Jane Doe
- System location: RTP100
- Asset tag: 006429

The administrator then configures the MOTD banner to alert other switch administrators of an upcoming event.

To configure the switch:

1. Configure the hosts name.

```
console#configure
console(config)#hostname PC8024
```

2. Configure the contact, location, and asset tag. Notice that the prompt changed to the host name.

```
PC8024(config)#snmp-server contact “Jane Doe”
PC8024(config)#snmp-server location RTP100
PC8024(config)#asset-tag 006429
```

3. Configure the message that displays when a user connects to the switch.

```
(config)#banner motd This switch connects users in
cubicles C121-C139.”
(config)#exit
```

4 View system information to verify the configuration.

```
PC8024#show system
System Description: Dell Ethernet Switch
System Up Time: 0 days, 19h:36m:36s
System Contact: Jane Doe
System Name: PC8024
System Location: RTP100
Burned In MAC Address: 001E.C9AA.AA07
System Object ID: 1.3.6.1.4.1.674.10895.3035
System Model ID: PC8024
Machine Type: PowerConnect PC8024
Temperature Sensors:

Temperature Sensors:

Unit     Description   Temperature    Status
                        (Celsius)
----     -----------   -----------    ------
1        CPU           33             Good
1        MAC           39             Good
1        Left PHY      32             Good
1        Right PHY     33             Good

Fans:

Unit     Description    Status
----     -----------    ------
1       Fan 1          OK
1       Fan 2          OK
1       Fan 3          OK
```

```
Power Supplies:

Unit  Description    Status    Average    Current    Since
                               Power      Power      Date/Time
                              (Watts)    (Watts)
---- ----------    --------  ----------  -------- ------------
1     System       OK          5.0        97.8
1     Main         Failure
1     Secondary    OK          97.6       97.8      01/10/2031
                                                     15:59:05
```

5 View additional information about the system.

```
PC8024#show system id

Service Tag: 0000000
Chassis Service Tag:
Serial Number: TW282987BK0002
Asset Tag: 111222
Unit Service tag Chassis Serv tag  Serial number Asset tag
---- ------------ ---------------- ------------- ----------
1    0000000                      TW282987BK0002 111222
```

6 Initiate a new Telnet session to verify the MOTD.

Figure 12-19. Verify MOTD



Configuring SNTP

The commands in this example configure the switch to poll an SNTP server to synchronize the time. Additionally, the SNTP sessions between the client and server must be authenticated.

To configure the switch:

1. Configure the authentication information. The SNTP server must be configured with the same authentication key and ID.

```
console#configure
console(config)#sntp authentication-key 23456465
md5 sntpkey
console(config)#sntp trusted-key 23456465
console(config)#sntp authenticate
```

2. Specify the IP address of the SNTP server to poll and include the authentication key. This command automatically enables polling and sets the priority to 1.

```
console(config)#sntp server 192.168.10.30 key
23456465
console(config)#sntp unicast client enable
```

3. Verify the configuration.

```
console#show sntp configuration

Polling interval: 512 seconds
MD5 Authentication keys: 23456465
Authentication is required for synchronization.
Trusted keys: 23456465
Unicast clients: Enable

Unicast servers:
Server          Key         Polling   Priority
------------    ----------- --------- --------
192.168.10.30   23456465    Enabled   1
```

4 View the SNTP status on the switch.

```
console#show sntp status

Client Mode:           Unicast
Last Update Time:      MAR 01 09:12:43 2010

Unicast servers:
Server          Status       Last response
--------------- ------------ ---------------------
192.168.10.30   Other        09:12:43 Mar 1 2011
```

Configuring the Time Manually

The commands in this example manually set the system time and date. The time zone is set to Eastern Standard Time (EST), which has an offset of -5 hours. Summer time is enabled and uses the preconfigured United States settings.

To configure the switch:

1. Configure the time zone offset and acronym.

```
console#configure
console(config)#clock timezone -5 zone EST
```

2. Configure the summer time (daylight saving time) to use the preconfigured settings for the United States.

```
console(config)#clock summer-time recurring us
```

3. Set the local time and date.

```
console(config)#clock set 16:13.06 03/01/2010
```

4. Verify the time settings.

```
console#show clock detail

00:27:19 EST(UTC-5:00) Feb 3 2039
No time source

 Time zone:
 Acronym is EST
 Offset is UTC-5:00

 Summertime:
 Acronym not configured
 Recurring every year (USA)
 Begins on second Sunday of Mar at 02:00
 Ends on first Sunday of Nov at 02:00
 Offset is +60 minutes
```

13

Configuring SNMP

The topics covered in this chapter include:

- SNMP Overview
- Default SNMP Values
- Configuring SNMP (Web)
- Configuring SNMP (CLI)
- SNMP Configuration Examples

SNMP Overview

Simple Network Management Protocol (SNMP) provides a method for managing network devices. The PowerConnect 8000-series and 8100-series switches support SNMP version 1, SNMP version 2, and SNMP version 3.

What Is SNMP?

SNMP is a standard protocol that enables remote monitoring and management of a device through communication between an SNMP manager and an SNMP agent on the remote device. The SNMP manager is typically part of a Network Management System (NMS) that runs on an administrative host. The switch software includes Management Information Base (MIB) objects that the SNMP agent queries and modifies. The switch uses standard public MIBs and private MIBs.

A MIB acts as a structured road map for managed objects. A managed object is any feature or setting that can be configured or monitored on the switch. An Object Identifier (OID) is the unique number assigned to an object defined in a MIB. An OID is written as a sequence of subidentifiers in decimal notation.

The SNMP agent maintains a list of variables that are used to manage the switch. The variables are defined in the MIB. The MIB presents the variables controlled by the agent. The SNMP agent defines the MIB specification format, as well as the format used to access the information over the network. Access rights to the SNMP agent are controlled by access strings.

SNMP v3 also applies access control and a new traps mechanism to SNMPv1 and SNMPv2 PDUs. In addition, the User Security Model (USM) is defined for SNMPv3 and includes:

- Authentication — Provides data integrity and data origin authentication.
- Privacy — Protects against disclosure of message content. Cipher-Bock-Chaining (CBC) is used for encryption. Either authentication is enabled on an SNMP message, or both authentication and privacy are enabled on an SNMP message. However privacy cannot be enabled without authentication.
- Timeliness — Protects against message delay or message redundancy. The SNMP agent compares incoming message to the message time information.
- Key Management — Defines key generation, key updates, and key use.

Authentication or Privacy Keys are modified in the SNMPv3 User Security Model (USM).

What Are SNMP Traps?

SNMP is frequently used to monitor systems for fault conditions such as temperature violations, link failures, and so on. Management applications can monitor for these conditions by polling the appropriate OIDs with the `get` command and analyzing the returned data. This method has its drawbacks. If it is done frequently, significant amounts of network bandwidth can be consumed. If it is done infrequently, the response to the fault condition may not occur in a timely fashion. SNMP traps avoid these limitations of the polling method.

An SNMP trap is an asynchronous event indicating that something significant has occurred. This is analogous to a pager receiving an important message, except that he SNMP trap frequently contains all the information needed to diagnose a fault.

You can configure various features on the switch to generate SNMP traps that inform the NMS about events or problems that occur on the switch. Traps generated by the switch can also be viewed locally by using the web-based interface or CLI.

Why Is SNMP Needed?

Some network administrators prefer to use SNMP as the switch management interface. Settings that you view and configure by using the web-based Dell OpenManage Switch Administrator and the CLI are also available by using SNMP.

If you do not use NMS software to manage or monitor other devices on your network, it might not be necessary to configure SNMP on the switch.

Default SNMP Values

By default, SNMPv2 is automatically enabled on the device. SNMPv1 and SNMPv3 are disabled. To enable SNMPv3, you must define a local engine ID for the device. The local engineID is by default set to the switch MAC address. This local engineID must be defined so that it is unique within the network.

Table 13-1 summarizes the default values for SNMP.

Table 13-1. SNMP Defaults

Parameter	Default Value
SNMPv1	Disabled
SNMPv2	Enabled
SNMPv3	Disabled
SNMP traps	Enabled
SNMP trap receiver	None configured
Switch traps	Enabled
QoS traps	Enabled
Multicast traps	Disabled
Captive Portal traps	Disabled

Table 13-1. SNMP Defaults

Parameter	Default Value
OSPF traps	Disabled

Table 13-2 describes the two views that are defined by default.

Table 13-2. SNMP Default Views

View Name	OID Subtree	View Type
Default	iso	Included
	snmpVacmMIB	Excluded
	usmUser	Excluded
	snmpCommunityTable	Excluded
DefaultSuper	iso	Included

By default, three groups are defined. Table 13-3 describes the groups. The Read, Write, and Notify values define the preconfigured views that are associated with the groups.

Table 13-3. SNMP Default Groups

Group Name	Security Level	Read	Write	Notify
DefaultRead	No Auth No Priv	Default	–	Default
DefaultWrite	No Auth No Priv	Default	Default	Default
DefaultSuper	No Auth No Priv	DefaultSuper	DefaultSuper	DefaultSuper

Configuring SNMP (Web)

This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring the SNMP agent on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

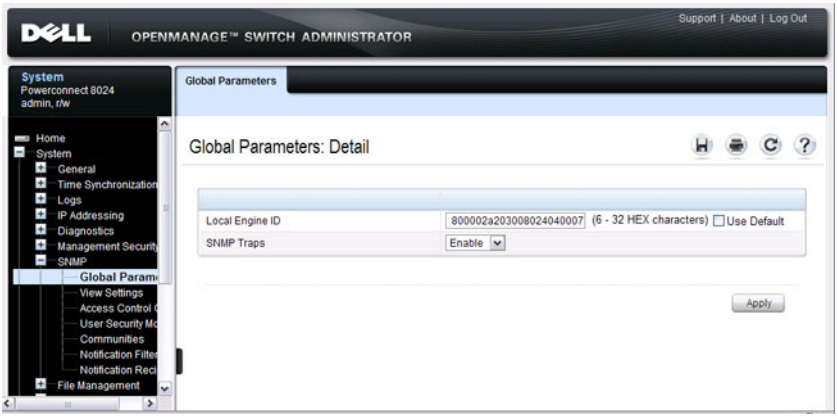
NOTE: For some features, the control to enable or disable traps is available from a configuration page for that feature and not from the **Trap Manager** pages that this chapter describes.

SNMP Global Parameters

Use the **Global Parameters** page to enable SNMP and Authentication notifications.

To display the **Global Parameters** page, click **System** → **SNMP** → **Global Parameters** in the navigation panel.

Figure 13-1. SNMP Global Parameters

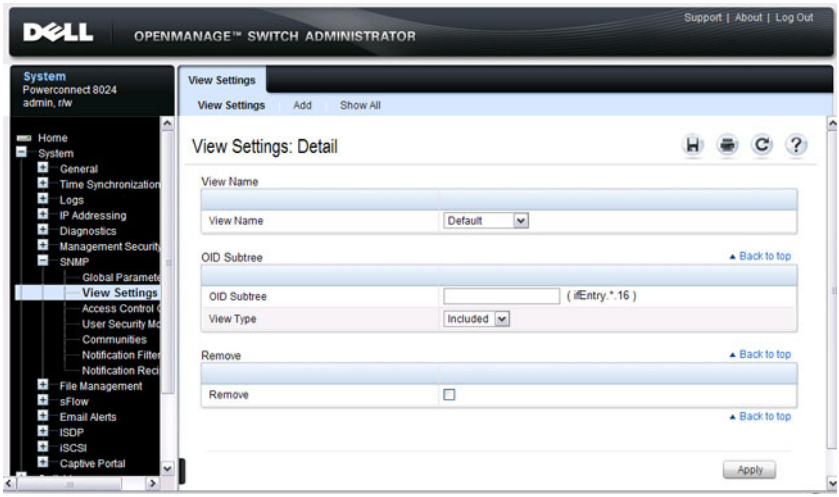


SNMP View Settings

Use the SNMP **View Settings** page to create views that define which features of the device are accessible and which are blocked. You can create a view that includes or excludes OIDs corresponding to interfaces.

To display the **View Settings** page, click **System** → **SNMP** → **View Settings** in the navigation panel.

Figure 13-2. SNMP View Settings



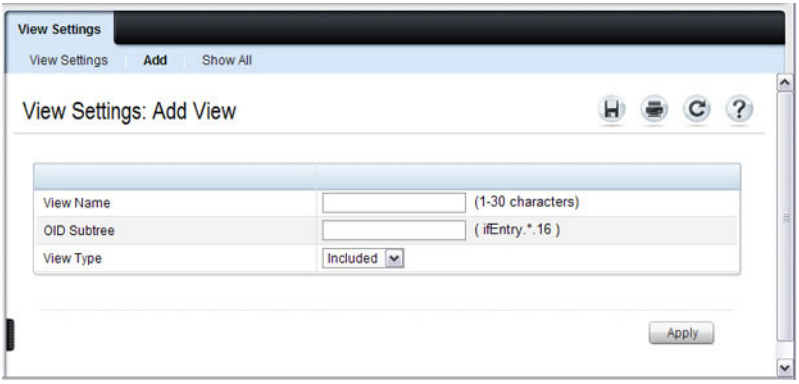
Adding an SNMP View

To add a view:

1. Open the **View Settings** page.
2. Click **Add**.

   The **Add View** page displays:

Figure 13-3. Add View



3 Specify a name for the view and a valid SNMP OID string.

4 Select the view type.

5 Click **Apply**.

The SNMP view is added, and the device is updated.

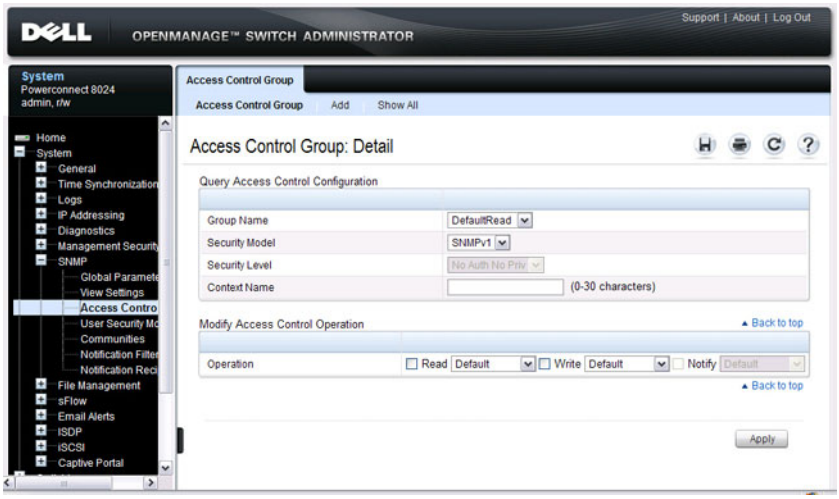
Click **Show All** to view information about configured SNMP Views.

Access Control Group

Use the **Access Control Group** page to view information for creating SNMP groups, and to assign SNMP access privileges. Groups allow network managers to assign access rights to specific device features or features aspects.

To display the **Access Control Group** page, click **System** → **SNMP** → **Access Control** in the navigation panel.

Figure 13-4. SNMP Access Control Group



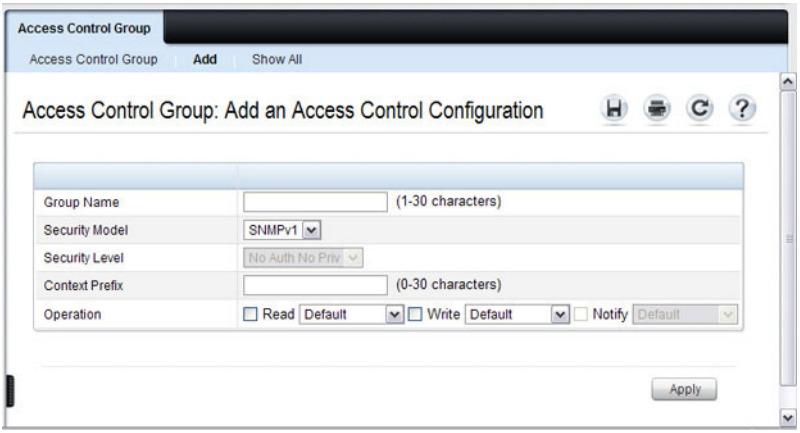
Adding an SNMP Group

To add a group:

1. Open the **Access Control Configuration** page.
2. Click **Add**.

   The **Add an Access Control Configuration** page displays:

Figure 13-5. Add Access Control Group



3 Specify a name for the group.

4 Select a security model and level

5 Define the context prefix and the operation.

6 Click **Apply** to update the switch.

Click **Show All** to view information about existing access control configurations.

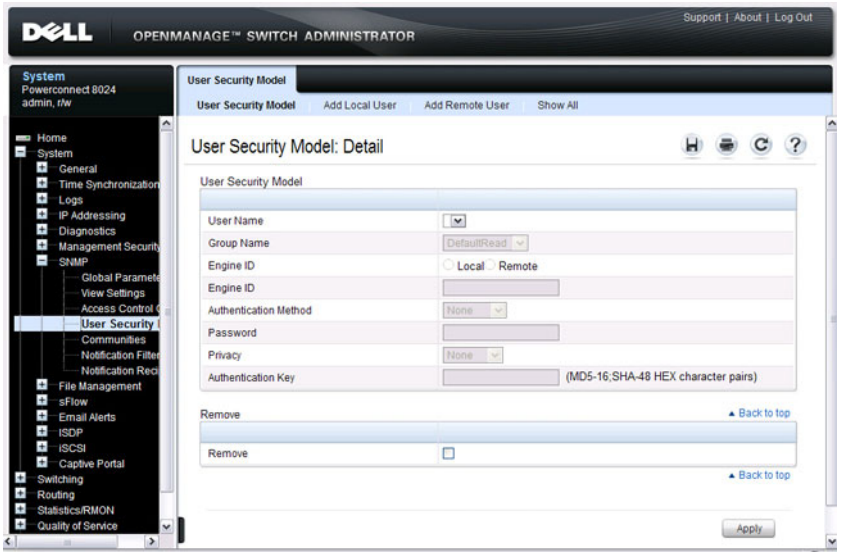
SNMPv3 User Security Model (USM)

Use the **User Security Model** page to assign system users to SNMP groups and to define the user authentication method.

NOTE: You can also use the Local User Database page under **Management Security** to configure SNMPv3 settings for users. For more information, see "Configuring Authentication, Authorization, and Accounting" on page 177.

To display the **User Security Model** page, click **System** → **SNMP** → **User Security Model** in the navigation panel.

Figure 13-6. SNMPv3 User Security Model



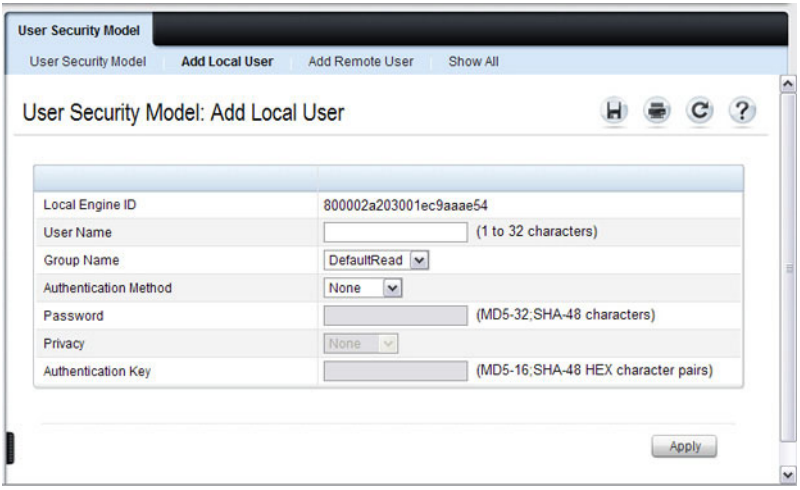
Adding Local SNMPv3 Users to a USM

To add local users:

1. Open the **User Security Model** page.
2. Click **Add Local User**.

   The **Add Local User** page displays:

Figure 13-7. Add Local Users



3 Define the relevant fields.

4 Click **Apply** to update the switch.

Click **Show All** to view the User Security Model Table, which contains information about configured Local and Remote Users.

Adding Remote SNMPv3 Users to a USM

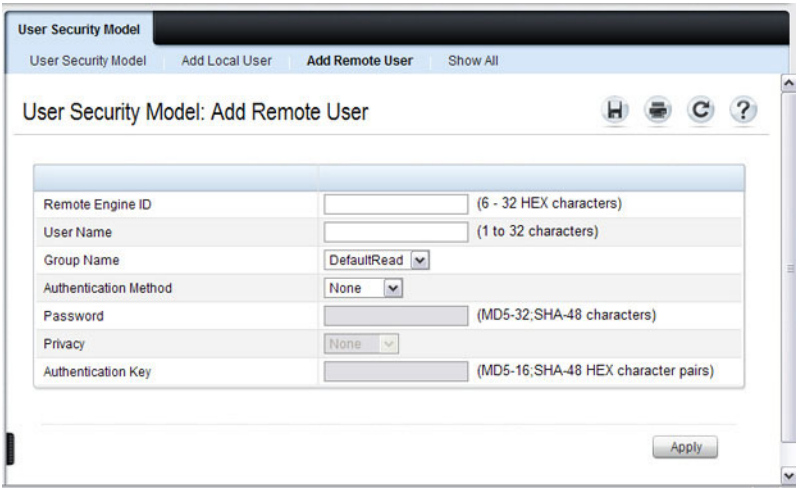
To add remote users:

1 Open the **SNMPv3 User Security Model** page.

2 Click **Add Remote User**.

The **Add Remote User** page displays:

Figure 13-8. Add Remote Users



3 Define the relevant fields.

4 Click **Apply** to update the switch.

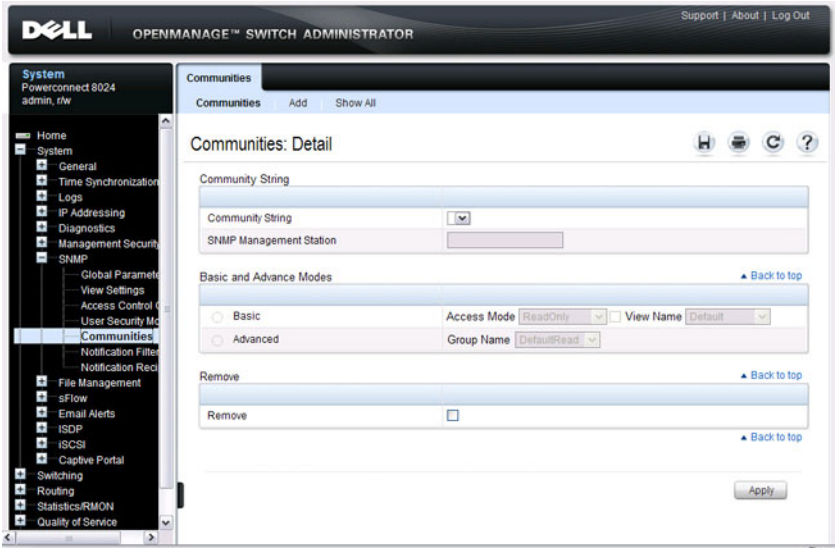
Click **Show All** to view the User Security Model Table, which contains information about configured Local and Remote Users.

Communities

Access rights for SNMPv1 and SNMPv2 are managed by defining communities **Communities** page. When the community names are changed, access rights are also changed. SNMP Communities are defined only for SNMP v1 and SNMP v2.

To display the **Communities** page, click **System** → **SNMP** → **Communities** in the navigation panel.

Figure 13-9. SNMP Communities



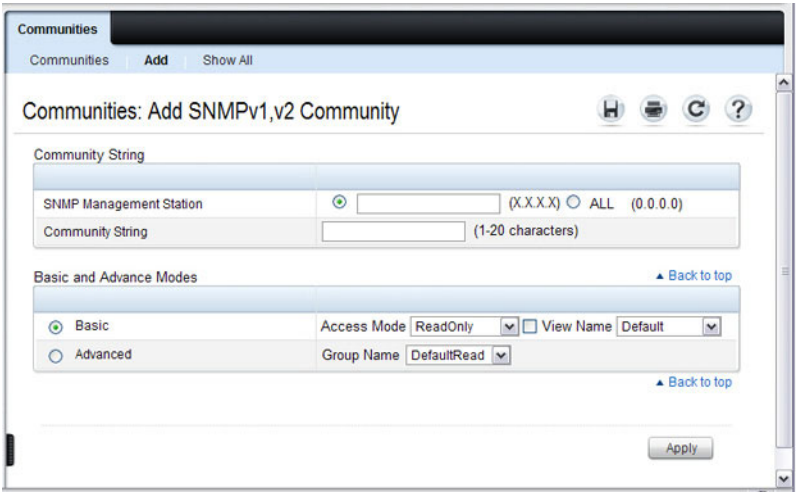
Adding SNMP Communities

To add a community:

1. Open the **Communities** page.
2. Click **Add**.

   The **Add SNMPv1,2 Community** page displays:

Figure 13-10. Add SNMPv1,2 Community



3 Specify the IP address of an SNMP management station and the community string to act as a password that will authenticate the management station to the SNMP agent on the switch.

4 Select the access mode.

5 Click **Apply** to update the switch.

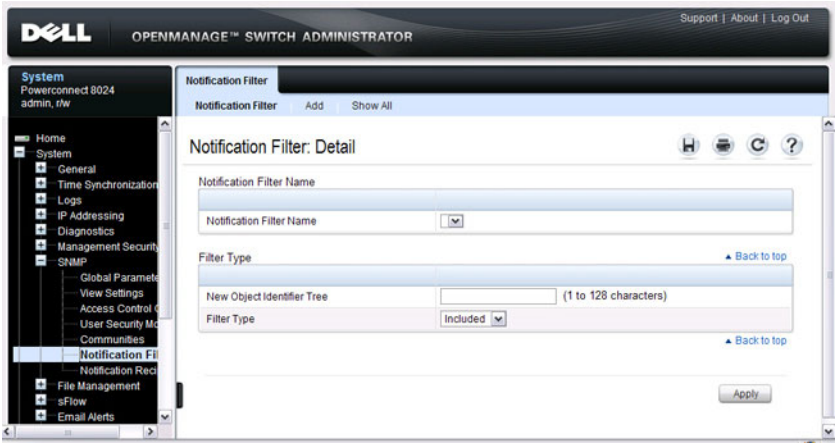
Click **Show All** to view the communities that have already been configured.

Notification Filter

Use the **Notification Filter** page to set filtering traps based on OIDs. Each OID is linked to a device feature or a feature aspect. The **Notification Filter** page also allows you to filter notifications.

To display the **Notification Filter** page, click **System** → **SNMP** → **Notification Filters** in the navigation panel.

Figure 13-11. SNMP Notification Filter



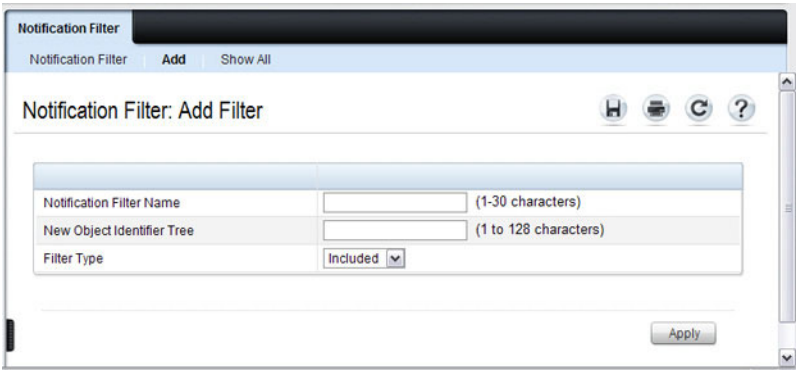
Adding a Notification Filter

To add a filter:

1. Open the **Notification Filter** page.
2. Click **Add**.

   The **Add Filter** page displays:

Figure 13-12. Add Notification Filter



3 Specify the name of the filter, the OID for the filter.

4 Choose whether to send (include) traps or informs to the trap recipient or prevent the switch from sending (exclude) the traps or informs.

5 Click **Apply** to update the switch.

Click **Show All** to view information about the filters that have already been configured.

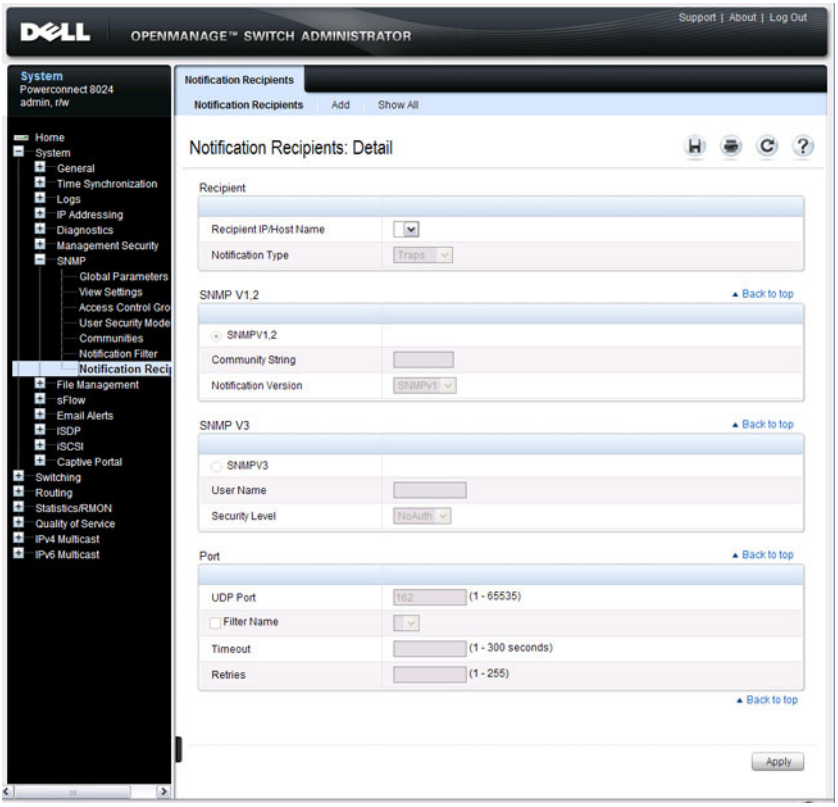
Notification Recipients

Use the **Notification Recipients** page to view information for defining filters that determine whether traps are sent to specific users, and the trap type sent. SNMP notification filters provide the following services:

- Identifying Management Trap Targets
- Trap Filtering
- Selecting Trap Generation Parameters
- Providing Access Control Checks

To display the **Notification Recipients** page, click **System** → **SNMP** → **Notification Recipient** in the navigation panel.

Figure 13-13. SNMP Notification Recipient



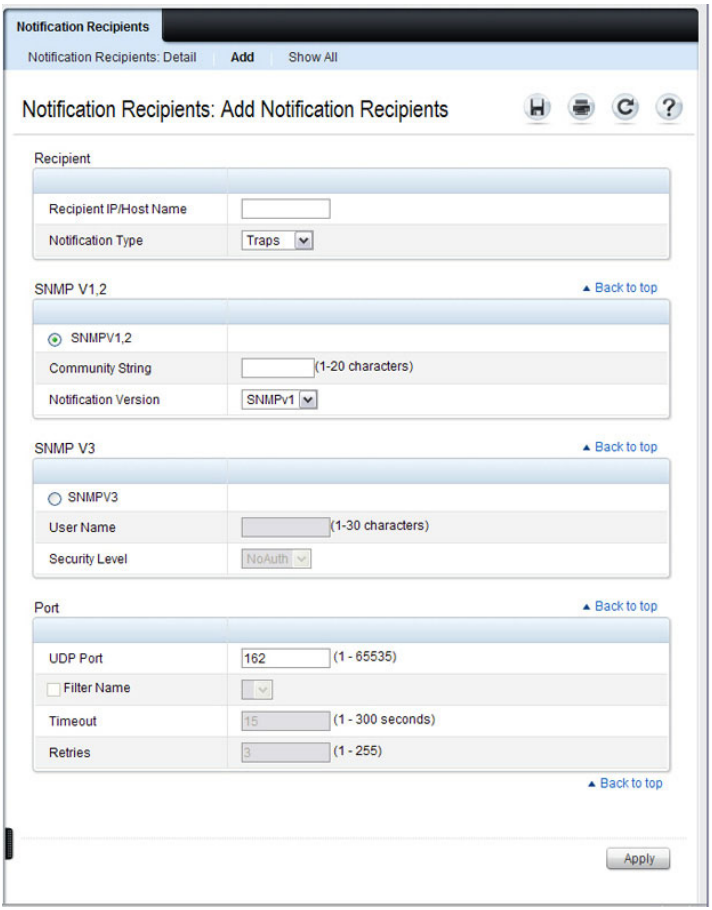
Adding a Notification Recipient

To add a recipient:

1. Open the **Notification Recipient** page.
2. Click **Add**.

   The **Add Recipient** page displays:

Figure 13-14. Add Notification Recipient



3 Specify the IP address or hostname of the host to receive notifications.

4 Select whether to send traps or informs to the specified recipient

5 Define the relevant fields for the SNMP version you use.

6 Configure information about the port on the recipient.

7 Click **Apply** to update the switch.

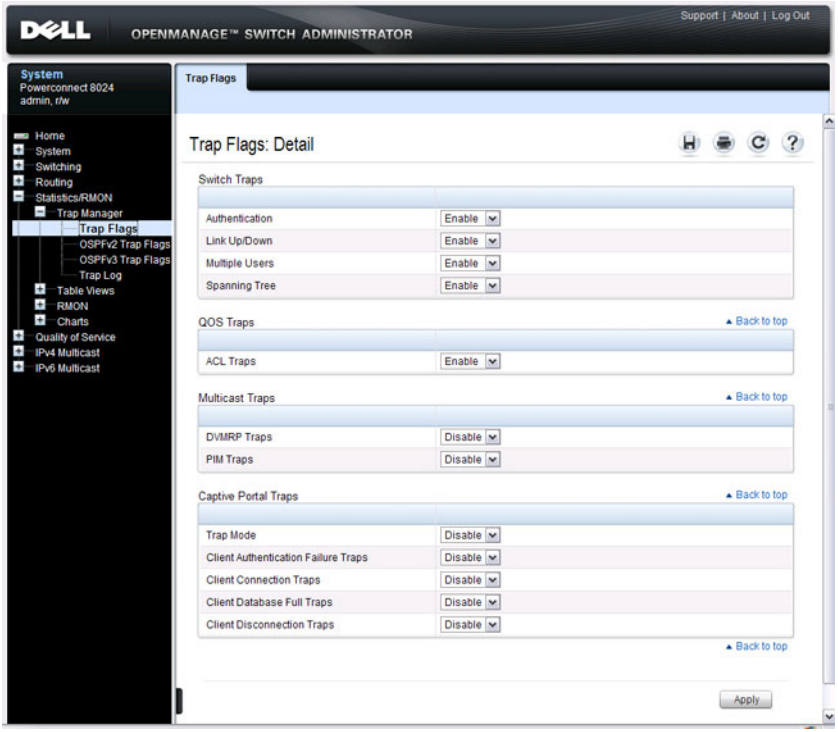
Click **Show All** to view information about the recipients that have already been configured.

Trap Flags

The **Trap Flags** page is used to specify which traps you want to enable or disable. When the condition identified by an active trap is encountered by the switch, a trap message is sent to any enabled SNMP Trap Receivers, and a message is written to the trap log.

To access the **Trap Flags** page, click **Statistics/RMON** → **Trap Manager** → **Trap Flags** in the navigation panel.

Figure 13-15. Trap Flags

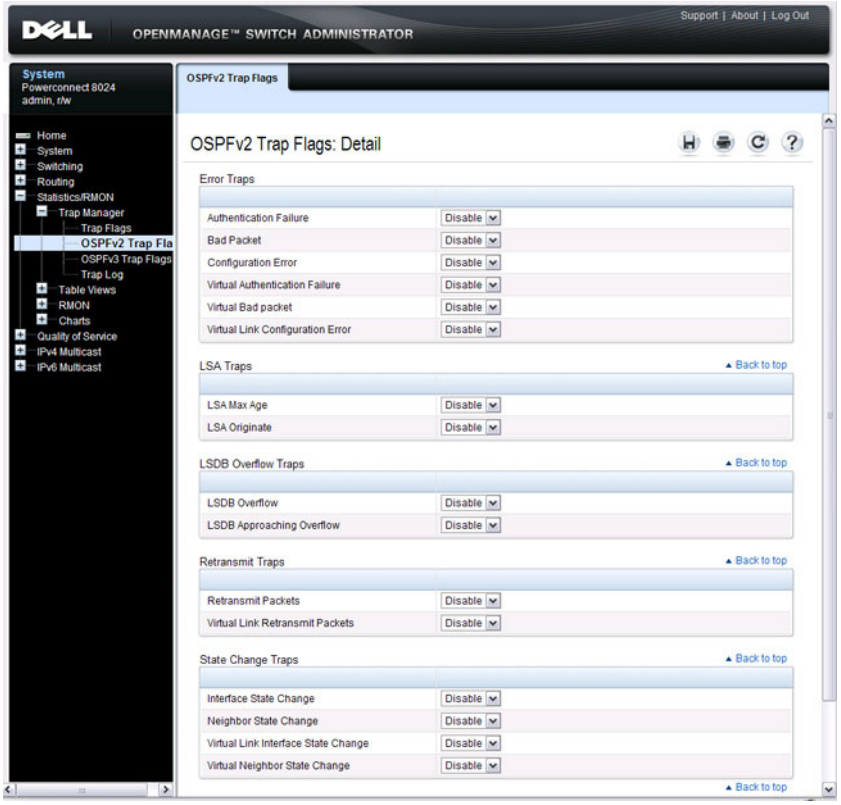


OSPFv2 Trap Flags

The **OSPFv2 Trap Flags** page is used to specify which OSPFv2 traps you want to enable or disable. When the condition identified by an active trap is encountered by the switch, a trap message is sent to any enabled SNMP Trap Receivers, and a message is written to the trap log.

To access the **OSPFv2 Trap Flags** page, click **Statistics/RMON → Trap Manager → OSPFv2 Trap Flags** in the navigation panel.

Figure 13-16. OSPFv2 Trap Flags

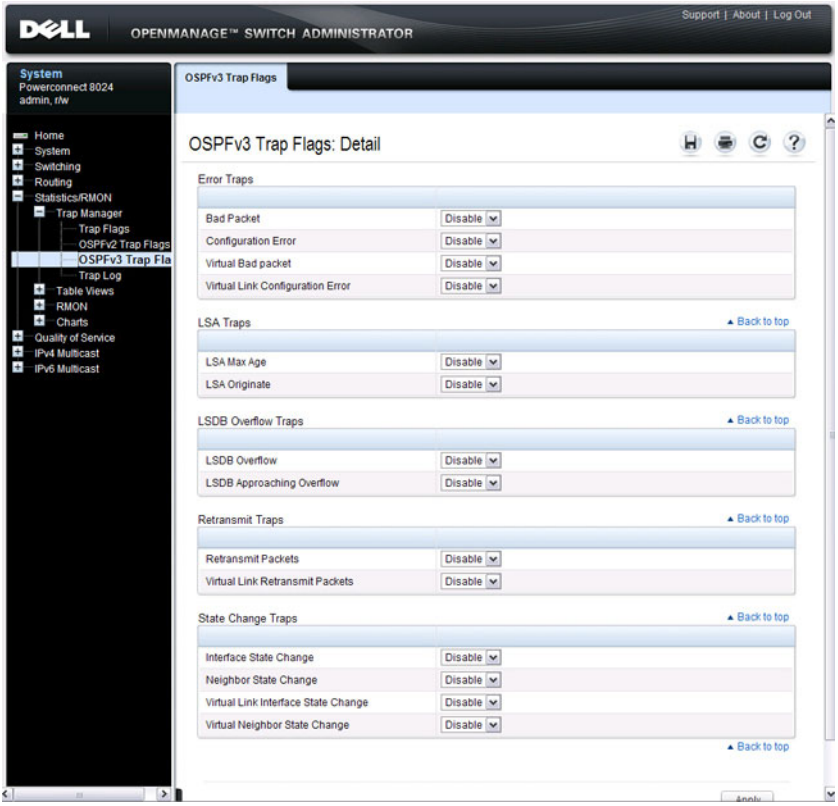


OSPFv3 Trap Flags

The **OSPFv3 Trap Flags** page is used to specify which OSPFv3 traps you want to enable or disable. When the condition identified by an active trap is encountered by the switch, a trap message is sent to any enabled SNMP Trap Receivers, and a message is written to the trap log.

To access the **OSPFv3 Trap Flags** page, click **Statistics/RMON → Trap Manager → OSPFv3 Trap Flags** in the navigation panel.

Figure 13-17. OSPFv3 Trap Flags

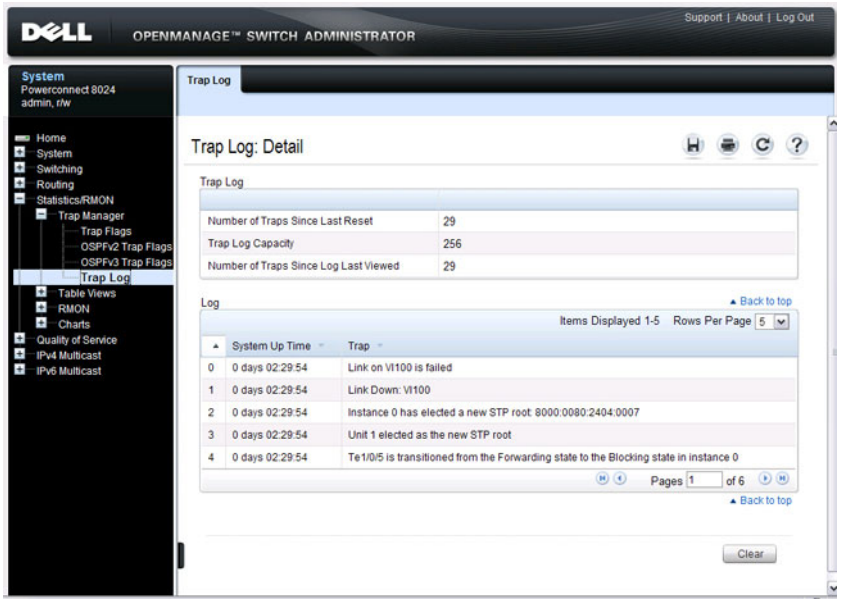


Trap Log

The **Trap Log** page is used to view entries that have been written to the trap log.

To access the **Trap Log** page, click **Statistics/RMON → Trap Manager → Trap Log** in the navigation panel.

Figure 13-18. Trap Logs



Click **Clear** to delete all entries from the trap log.

Configuring SNMP (CLI)

This section provides information about the commands you use to manage and view SNMP features on the switch. For more information about these commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Configuring the SNMPv3 Engine ID

To use SNMPv3, the switch must have engine ID. You can specify your own ID or use the default string that is generated using the MAC address of the switch. If the SNMPv3 engine ID is deleted, or if the configuration file is erased, then SNMPv3 cannot be used. Since the EngineID should be unique within an administrative domain, Dell recommends that you use the default keyword to configure the Engine ID.

Changing the value of SNMP EngineID has important side-effects. A user's password (entered on the command line) is converted to an MD5 or SHA security digest. This digest is based on both the password and the local engine ID. The command line password is then destroyed, as required by RFC 2274. Because of this deletion, if the local value of engineID changes, the security digests of SNMPv3 users will be invalid, and the users will have to be reconfigured.

Beginning in Privileged EXEC mode, use the following commands to configure an engine ID for SNMP.

Command	Purpose
configure	Enter Global Configuration mode
snmp-server engineID local {*engineid-string* | **default**}	Configure the SNMPv3 Engine ID. • engineid-string — The character string that identifies the engine ID. The engine ID is a concatenated hexadecimal string. Each byte in hexadecimal character strings is two hexadecimal digits. Each byte can be separated by a period or colon. (Range: 6-32 characters) • default — The engineID is created automatically, based on the device MAC address.
exit	Exit to Privileged EXEC mode.
show snmp engineid	View the local SNMP engine ID.

Configuring SNMP Views, Groups, and Users

Beginning in Privileged EXEC mode, use the following commands to define SNMP views, and SNMP groups, and local and remote SNMPv3 users.

Command	Purpose
configure	Enter Global Configuration mode
snmp-server view *view-name oid-tree* **{included | excluded}**	Configure the SNMP view. When you configure groups, users, and communities, you can specify a view to associate with the group, user, or community • *view-name* — Specifies the name of the view. (Range: 1-30 characters.) • *oid-tree* — Specifies the object identifier of the ASN.1 subtree to be included or excluded from the view. To identify the subtree, specify a text string consisting of numbers, such as `1.3.6.2.4`, or a word, such as `system`. Replace a single subidentifier with the asterisk (*) wildcard to specify a subtree family; for example `1.3.*.4`. • **included** — Indicates that the view type is included. • **excluded** — Indicates that the view type is excluded.

Command	Purpose
snmp-server group *groupname* {**v1** | **v2** | **v3** {**noauth** | **auth** | **priv**} [**notify** *view-name*]} [**context** *view-name*] [**read** *view-name*] [**write** *view-name*]	Specify the identity string of the receiver and set the receiver timeout value. • *groupname* — Specifies the name of the group. (Range: 1-30 characters.) • **v1** — Indicates the SNMP Version 1 security model. • **v2** — Indicates the SNMP Version 2 security model. • **v3** — Indicates the SNMP Version 3 security model. • **noauth** — Indicates no authentication of a packet. Applicable only to the SNMP Version 3 security model. • **auth** — Indicates authentication of a packet without encrypting it. Applicable only to the SNMP Version 3 security model. • **priv** — Indicates authentication of a packet with encryption. Applicable only to the SNMP Version 3 security model. • *view-name* — Specifies the view (defined in the previous step) to use for the context, notification, read, and write privileges for the group.

Command	Purpose
snmp-server user *username groupname* [**remote** *engineid-string*] [{**auth-md5** *password* | **auth-sha** *password* | **auth-md5-key** *md5-key* | **auth-sha-key** *sha-key*} [**priv-des** *password* | **priv-des-key** *des-key*]]	Configure a new SNMPv3 user. • *username* — Specifies the name of the user on the host that connects to the agent. (Range: 1-30 characters.) • *groupname* — Specifies the name of the group to which the user belongs. (Range: 1-30 characters.) • *engineid-string* — Specifies the engine ID of the remote SNMP entity to which the user belongs. The engine ID is a concatenated hexadecimal string. Each byte in the hexadecimal character string is two hexadecimal digits. The remote engine id designates the remote management station, and should be defined to enable the device to receive acknowledgements to "informs." (Range: 5-32 characters.) • **auth-md5** — The HMAC-MD5-96 authentication level. • **auth-sha** — The HMAC-SHA-96 authentication level. • *password* — A password. (Range: 1 to 32 characters.) • auth-**md5-key** — The HMAC-MD5-96 authentication level. Enter a pregenerated MD5 key. • auth-**sha-key** — The HMAC-SHA-96 authentication level. Enter a pregenerated SHA key. • *md5-key* — Character string—length 32 hex characters. • *sha-key* — Character string—length 48 characters. • **priv-des** — The CBC-DES Symmetric Encryption privacy level. Enter a password. • **priv-des-key** — The CBC-DES Symmetric Encryption privacy level. The user should enter a pregenerated MD5 or SHA key depending on the authentication level selected. • *des-key* — The pregenerated DES encryption key. Length is determined by authentication method selected—32 hex characters if MD5 Authentication is selected, 48 hex characters if SHA Authentication is selected.
exit	Exit to Privileged EXEC mode.
show snmp views	View SNMP view configuration information.

Command	Purpose
show snmp group [*group_name*]	View SNMP group configuration information.
show snmp user [*user_name*]	View SNMP user configuration information.

Configuring Communities

Beginning in Privileged EXEC mode, use the following commands to configure access rights for SNMPv1 and SNMPv2.

Command	Purpose
configure	Enter Global Configuration mode
snmp-server community *string* [**ro** | **rw** | **su**] [**view** *view-name*] [**ipaddress** *ip_address*]	Configure the community string and specify access criteria for the community. • *community-string* — Acts as a password and is used to authenticate the SNMP management station to the switch. The string must also be defined on the NMS in order for the NMS to access the SNMP agent on the switch (Range: 1-20 characters) • **ro** — Indicates read-only access • **rw** — Indicates read-write access. • *view-name* — Specifies the name of a previously defined MIB view. • *ip_address* — Specifies the IP address of the management station. If no IP address is specified, all management stations are permitted

Command	Purpose
snmp-server community-group *community-string group-name* [**ipaddress** *ip-address*]	Map the internal security name for SNMP v1 and SNMP v2 security models to the group name. • *community-string* — Community string that acts like a password and permits access to the SNMP protocol. (Range: 1-20 characters) • *group-name* — Name of a previously defined group. The group defines the objects available to the community. (Range: 1-30 characters) • *ip-address* — Management station IP address. Default is all IP addresses.
exit	Exit to Privileged EXEC mode.
show snmp	View SNMP settings and verify the configuration

Configuring SNMP Notifications (Traps and Informs)

Beginning in Privileged EXEC mode, use the following commands to allow the switch to send SNMP traps and to configure which traps are sent.

Command	Purpose
configure	Enter Global Configuration mode
snmp-server enable traps [**acl** | **all** | **auto-copy-sw** | **captive-portal** *cp-type* | **dot1q** | **dvrmp** | **link** | **maclock** | **multiple-users** | **ospf** *ospftype* | **ospfv3** *ospfv3type* | **pim** | **poe** | **snmp authentication** | **spanning-tree** | **vrrp**]	Specify the traps to enable. The captive portal, OSPF and OSPFv3 traps include several different traps that can be enabled. For more information, use the CLI command help or see the CLI Command Reference.
snmp-server filter *filter-name oid-tree* {**included** | **excluded**}	Configure a filter for SNMP traps and informs based on OIDs. Each OID is linked to a device feature or a feature aspect. • *filter-name* — Specifies the label for the filter record that is being updated or created. The name is used to reference the record. (Range: 1-30 characters.) • *oid-tree* — Specifies the object identifier of the ASN.1 subtree to be included or excluded from the view. To identify the subtree, specify a text string consisting of numbers, such as `1.3.6.2.4`, or a word, such as `system`. Replace a single subidentifier with the asterisk (*) wildcard to specify a subtree family; for example, `1.3.*.4`. • **included** — Indicates that the filter type is included. • **excluded** — Indicates that the filter type is excluded.

Command	Purpose
snmp-server host *host-addr* [**informs** [**timeout** *seconds*] [**retries** *retries*] | **traps version** {1 | 2}]] *community-string* [**udp-port** *port*] [**filter** *filtername*]	For SNMPv1 and SNMPv2, configure the system to receive SNMP traps or informs. • *host-addr* — Specifies the IP address of the host (targeted recipient) or the name of the host. (Range:1-158 characters). • **informs** — Indicates that SNMPv2 informs are sent to this host • **timeout** *seconds* — Number of seconds to wait for an acknowledgment before resending informs. The default is 15 seconds. (Range: 1-300 characters.) • *retries* — Maximum number of times to resend an inform request. The default is 3 attempts. • **traps** — Indicates that SNMP traps are sent to this host – **version 1** — Indicates that SNMPv1 traps will be used – **version 2** — Indicates that SNMPv2 traps will be used • *community-string* — Specifies a password-like community string sent with the notification operation. (Range: 1-20 characters) • *port* — UDP port of the host to use. The default is 162. (Range: 1-65535 characters.) • *filtername* — A string that is the name of the filter that defines the filter for this host. If unspecified, does not filter anything (Range: 1-30 characters.)

Command	Purpose
snmp-server v3-host {*ip-address* | *hostname*} *username* {**traps** | **informs**} [**noauth** | **auth** | **priv**] [**timeout** *seconds*] [**retries** *retries*] [**udpport** *port*] [**filter** *filtername*]	For SNMPv3, configure the system to receive SNMP traps or informs. • *ip-address* — Specifies the IP address of the host (targeted recipient). • *hostname* — Specifies the name of the host. (Range: 1-158 characters.) • *username* — Specifies user name used to generate the notification. (Range: 1-25 characters.) • **traps** — Indicates that SNMP traps are sent to this host. • **informs** — Indicates that SNMPv2 informs are sent to this host. • **noauth** — Specifies sending of a packet without authentication. • **auth** — Specifies authentication of a packet without encrypting it • **priv** — Specifies authentication and encryption of a packet. • *seconds* — Number of seconds to wait for an acknowledgment before resending informs. This is not allowed for hosts configured to send traps. The default is 15 seconds. (Range: 1-300 seconds.) • *retries* — Maximum number of times to resend an inform request. This is not allowed for hosts configured to send traps. The default is 3 attempts. (Range: 0-255 retries.) • *port* — UDP port of the host to use. The default is 162. (Range: 1-65535.) • *filter-name* — Specifies the optional filter (defined with the **snmp-server filter** command) to use for the host. (Range: 1-30 characters.)
exit	Exit to Privileged EXEC mode.
show trapflags	View the status of the configurable SNMP traps.

SNMP Configuration Examples

This section contains the following examples:

- Configuring SNMPv1 and SNMPv2
- Configuring SNMPv3

Configuring SNMPv1 and SNMPv2

This example shows how to complete a basic SNMPv1/v2 configuration. The commands enable read-only access from any host to all objects on the switch using the community string *public*, and enable read-write access from any host to all objects on the switch using the community string *private*.

This example also shows how to allow the switch to generate traps for all features that produce traps. The traps are sent to the host with an IP address of 192.168.3.65 using the community string *public*.

To configure the switch:

1. Configure the public community string.

```
console#configure
console(config)#snmp-server community public ro
```

2. Configure the private community string.

```
console(config)#snmp-server community private rw
```

3. Enable all traps and specify the IP address of the host where the traps should be sent.

```
console(config)#snmp-server enable traps all
console(config)#snmp-server host 192.168.3.65
public
console(config)#exit
```

4. View the current SNMP configuration on the switch.

```
console#show snmp

Community-String     Community-Access View Name IP Address
-------------------- ---------------- --------- -------
private              Read/Write       Default    All
public               Read Only        Default    All
```

```
Community-String     Group Name           IP Address
-----------------    --------------       ------------
private              DefaultWrite         All
public               DefaultRead          All

Traps are enabled.
Authentication trap is enabled.
Version 1,2 notifications
Target Addr. Type  Community  Version UDP   Filter TO   Retries
                                      Port  Name   Sec
------------ ----  ---------  ----    ----- -----  ---  -------
192.168.3.65 Trap  public     1       162

Version 3 notifications
Target Addr. Type  Username   Security UDP   Filter TO   Retries
                              Level    Port  Name   Sec
------------ ----  ---------  -------  ----- -----  ---  -------

System Contact:
System Location:
```

Configuring SNMPv3

This example shows how to complete a basic SNMPv3 configuration. The commands create a view that includes objects from the *internet* MIB subtree (OID 1.3.6.1), which includes all objects on the switch.

The user named *admin* has read-write privileges to all objects within the view (in other words, all objects on the switch) after supplying the appropriate authentication credentials (secretkey).

To configure the switch:

1. Configure the view. *view_snmpv3* and specify the objects to include.

   ```
   console#configure
   console(config)#snmp-server view view_snmpv3
   internet included
   ```

2. Create the group *group_snmpv3* and allow read-write access to the view configured in the previous step.

   ```
   console(config)#snmp-server group group_snmpv3 v3
   auth read view_snmpv3 write view_snmpv3
   ```

3 Create the user *admin*, assign the user to the group, and specify the authentication credentials.

```
console(config)#snmp-server user admin
group_snmpv3 auth-md5 secretkey
```

4 Specify the IP address of the host where traps are to be sent. Packet authentication using MD5-SHA is enabled for the traps.

```
console(config)#snmp-server v3-host 192.168.3.35
admin traps auth
console(config)#exit
```

5 View the current SNMP configuration on the switch. The output includes the SNMPv1/2 configuration in the previous example.

```
console#show snmp

Community-String     Community-Access View Name IP Address
-------------------- ---------------- --------- -------
private              Read/Write       Default    All
public               Read Only        Default    All


Community-String     Group Name           IP Address
-----------------    --------------       ------------
private              DefaultWrite         All
public               DefaultRead          All

Traps are enabled.
Authentication trap is enabled.
Version 1,2 notifications
Target Addr. Type  Community  Version UDP   Filter TO   Retries
                                     Port  Name   Sec
------------ ---- ---------  ----    ----- -----  ---  -------
192.168.3.65 Trap  public     1       162

Version 3 notifications
Target Addr. Type  Username   Security UDP   Filter TO   Retries
                            Level    Port  Name   Sec
------------ ---- ---------  -------  ----- -----  ---  -------
192.168.3.35 Trap  admin      Auth-NoP 162          15   3

System Contact:
System Location:
```

```
console#show snmp views
Name                OID Tree                  Type
------------------  ------------------------  ------------
Default             iso                       Included
Default             snmpVacmMIB               Excluded
Default             usmUser                   Excluded
Default             snmpCommunityTable        Excluded
view_snmpv3         internet                  Included
DefaultSuper        iso                       Included
```

```
console#show snmp group

Name          Context  Model  Security Read     Views   Notify
              Prefix          Level            Write
------------ -------- ------ -------- -------- ------  -------
DefaultRead   ""       V1     NoAuth-  Default  ""      Default
                              NoPriv
DefaultRead   ""       V2     NoAuth-  Default  ""      Default
                              NoPriv
DefaultSuper  ""       V1     NoAuth-  DefaultSuDefault Default
                              NoPriv   per      Super   Super
DefaultSuper  ""       V2     NoAuth-  DefaultSuDefault Default
                              NoPriv   per      Super   Super
DefaultWrite  ""       V1     NoAuth-  Default  Default Default
                              NoPriv
DefaultWrite  ""       V2     NoAuth-  Default  Default Default
                              NoPriv
group_snmpv3  ""       V3     Auth-    view_snmpview_sn ""
                              NoPriv   v3       mpv3
```

```
console#show snmp user

Name       Group Name     Auth   Priv    Remote Engine ID
                          Meth   Meth
---------  -----------    -----  -----   ----------------
admin      group_snmpv3   MD5            800002a203001ec9aaaa07
```

14

Managing Images and Files

This chapter describes how to upload, download, and copy files, such as firmware images and configuration files, on the switch. The topics covered in this chapter include:

- Image and File Management Overview
- Managing Images and Files (Web)
- Managing Images and Files (CLI)
- File and Image Management Configuration Examples

NOTE: For information about the Auto Configuration feature that enables the switch to automatically upgrade the image or load a new configuration file during the boot process, see Automatically Updating the Image and Configuration.

Image and File Management Overview

What Files Can Be Managed?

PowerConnect 8000-series and 8100-series switches maintain several different types of files on the flash file system. Table 14-1 describes the files that you can manage. The table also lists the type of action you can take on the file, which is one or more of the following:

- Download the file to the switch from a remote system (or USB flash drive on the PowerConnect 8100-series devices).
- Upload the file from the switch to a remote system (or USB flash drive on the PowerConnect 8100-series devices).
- Copy the file from one location on the file system to another location.

Table 14-1. Files to Manage

File	Action	Description
image	Download Upload Copy	Firmware for the switch. The switch can maintain two images: the active image and the backup image.
startup-config	Download Upload Copy	Contains the software configuration that loads during the boot process.
running-config	Download Upload Copy	Contains the current switch configuration.
backup-config	Download Upload Copy	An additional configuration file that serves as a backup.
Configuration script	Download Upload	Text file with CLI commands. When you activate a script on the switch, the commands are executed and added to the running-config.
Log files	Upload	Provides various information about events that occur on the switch. For more information, see Monitoring and Logging System Information.
SSH key files	Download	Contains information to authenticate SSH sessions. The switch supports the following files for SSH: • SSH-1 RSA Key File • SSH-2 RSA Key File (PEM Encoded) • SSH-2 Digital Signature Algorithm (DSA) Key File (PEM Encoded)

Table 14-1. Files to Manage

File	Action	Description
SSL certificate files	Download	Contains information to encrypt, authenticate, and validate HTTPS sessions. The switch supports the following files for SSL: • SSL Trusted Root Certificate File (PEM Encoded) • SSL Server Certificate File (PEM Encoded) • SSL Diffie-Hellman Weak Encryption Parameter File (PEM Encoded) • SSL Diffie-Hellman Strong Encryption Parameter File (PEM Encoded)
IAS Users	Download	List of Internal Authentication Server (IAS) users for IEEE 802.1X authentication. For more information, see What is the Internal Authentication Server?

Why Is File Management Needed?

This section provides some reasons why you might choose to manage various files.

Image Files

The switch can store two firmware images, but only one is active. The other image file is a backup image. By default, the switch has only one image. You might copy an image or download an image to the switch for the following reasons:

- To create a backup image
- To upgrade the firmware as new images become available

Configuration Files

Configuration files contain the CLI commands that change the switch from its default configuration. The switch can maintain three separate configuration files: startup-config, running-config, and backup-config. The switch loads the startup-config file when the switch boots. Any configuration

changes that take place after the boot process completes are written to the running-config file. The backup-config file does not exist until you explicitly create one by copying an existing configuration file to the backup-config file or downloading a backup-config file to the switch.

You can also create configuration scripts, which are text files that contains CLI commands.

NOTE: You must use the CLI to manage configuration scripts. The configuration scripting feature is not available from the web interface.

When you apply (run) a configuration script on the switch, the commands in the script are executed in the order in which they are written as if you were typing them into the CLI. The commands that are executed in the configuration script are added to the running-config file.

You might upload a configuration file from the switch to a remote server for the following reasons:

- To create a backup copy
- To use the configuration file on another switch
- To manually edit the file

You might download a configuration file from a remote server to the switch for the following reasons:

- To restore a previous configuration
- To load the configuration copied from another switch
- To load the same configuration file on multiple switches

Use a text editor to open a configuration file and view or change its contents.

SSH/SSL Files

If you use OpenManage Switch Administrator to manage the switch over an HTTPS connection, you must copy the appropriate certificate files to the switch. If you use the CLI to manage the switch over an SSH connection, you must copy the appropriate key files to the switch.

What Methods Are Supported for File Management?

You can use any of the following protocols to download files from a remote system to the switch or to upload files from the switch to a remote system:

- TFTP
- SFTP
- SCP
- FTP
- HTTP (Web only)
- HTTPS (Web only)
- XMODEM (CLI only)

On the PowerConnect 8100-series switches, you can also copy files between the file system on the internal flash and a USB flash drive that is connected to the external USB port.

What Factors Should Be Considered When Managing Files?

Uploading and Downloading Files

To use TFTP, SFTP, SCP, or FTP for file management, you must provide the IP address of the remote system that is running the appropriate server (TFTP, SFTP, SCP or FTP). Make sure there is a route from the switch to the remote system. You can use the **ping** command from the CLI to verify that a route exists between the switch and the remote system.

If you are downloading a file from the remote system to the switch, be sure to provide the correct path to the file and the correct file name.

Managing Images

When you download a new image to the switch, it overwrites the backup image, if it exists. To use the new image, you must activate it and reload the switch. The image that was previously the active image becomes the backup image after the switch reloads. If you upgrade to a newer image and find that it is not compatible with your network, you can revert to the original image.

If you activate a new image and reload the switch, and the switch is unable to complete the boot process due to a corrupt image or other problem, you can use the boot menu to activate the backup image. You must be connected to

the switch through the console port to access the boot menu. The image files may contain firmware for the PHY processors on the switch. The PHY firmware may be updated to the firmware version supported by the switch firmware during the boot process or, in the case of switches that support the hot swap of cards, when the card is inserted into the switch.

Editing and Downloading Configuration Files

Each configuration file contains a list of executable CLI commands. The commands must be complete and in a logical order, as if you were entering them by using the switch CLI.

When you download a startup-config or backup-config file to the switch, the new file replaces the previous version. To change the running-config file, you execute CLI commands either by typing them into the CLI or by applying a configuration script with the **script apply** command. The startup-config and backup-config files can also be applied to the running-config by using the **script apply** command.

Creating and Applying Configuration Scripts

When you use configuration scripting, keep the following considerations and rules in mind:

- The application of scripts is partial if the script fails. For example, if the script executes four of ten commands and the script fails, the script stops at four, and the final six commands are not executed.
- Scripts cannot be modified or deleted while being applied.
- Validation of scripts checks for syntax errors only. It does not validate that the script will run.
- The file extension must be `.scr`.
- A maximum of seven scripts are allowed on the switch.
- The combined size of all script files on the switch cannot exceed 2 MB.
- The maximum number of configuration file command lines is 2000.

You can type single-line annotations in the configuration file to improve script readability. The exclamation point (!) character flags the beginning of a comment. The comment flag character can begin anywhere within a single

line, and all input following this character to the end of the line is ignored. Any line in the file that begins with the "!" character is recognized as a comment line and ignored by the parser.

The following example shows annotations within a file (commands are bold):

```
! Script file for displaying management access
show telnet !Displays the information about remote
connections
! Display information about direct connections
show serial
! End of the script file
```

Uploading Configuration Files by Using SNMP

When you use SNMP to upload a configuration file to a TFTP server, the agentTransferUploadFileName object must be set to the local filename, which is either startup-config or backup-config.

How Is the Running Configuration Saved?

Changes you make to the switch configuration while the switch is operating are written to the running-config. These changes are not automatically written to the startup-config. When you reload the switch, the startup-config file is loaded. If you reload the switch (or if the switch resets unexpectedly), any settings in the running-config that were not explicitly saved to the startup-config are lost. You must save the running-config to the startup-config to ensure that the settings you configure on the switch are saved across a switch reset.

To save the running-config to the startup-config by using the web-based interface, click  (the save icon), which is available at the top of each page.

To save the running-config to the startup-config from the CLI, use the **write** command.

Managing Images and Files (Web)

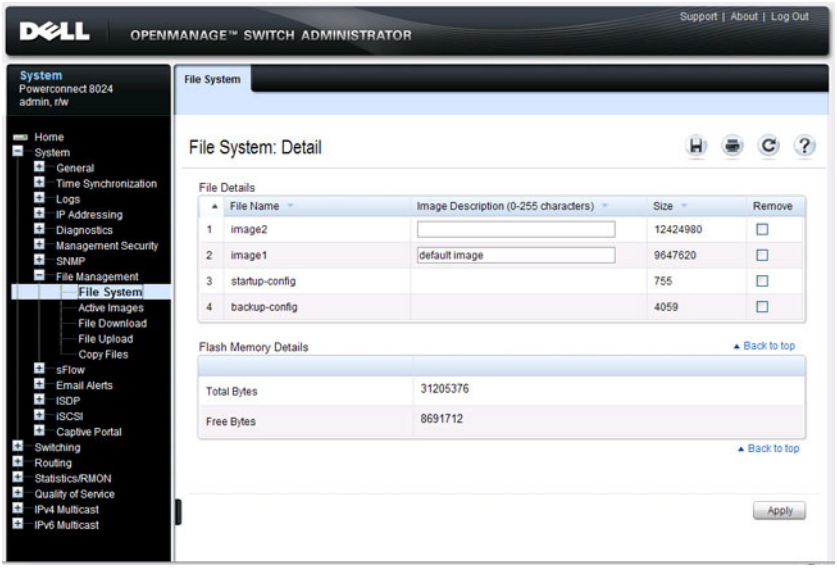
This section provides information about the OpenManage Switch Administrator pages to use to manage images and files on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

File System

Use the **File System** page to view a list of the files on the device and to modify the image file descriptions.

To display the **File System** page, click **System** → **File Management** → **File System** in the navigation panel.

Figure 14-1. File System



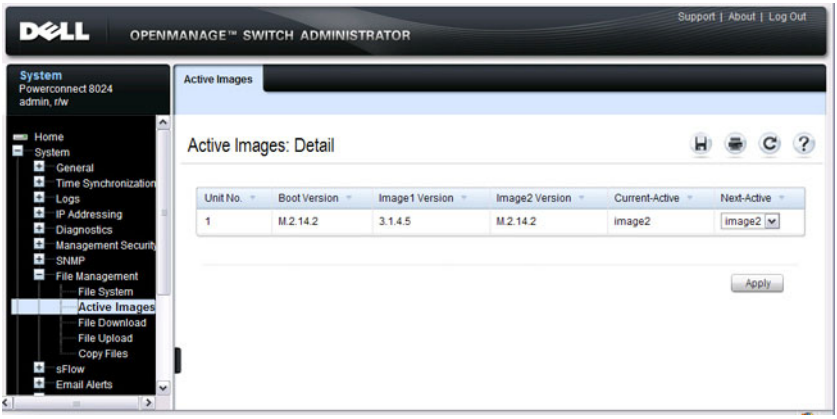
Active Images

Use the **Active Images** page to set the firmware image to use when the switch boots. If you change the boot image, it does not become the active image until you reset the switch.

NOTE: On the 8000-series switches, the images are named *image1* and *image2*. On the 8100-series switches, the images are named *active* and *backup*.

To display the **Active Images** page, click **System** → **File Management** → **Active Images** in the navigation panel.

Figure 14-2. Active Images

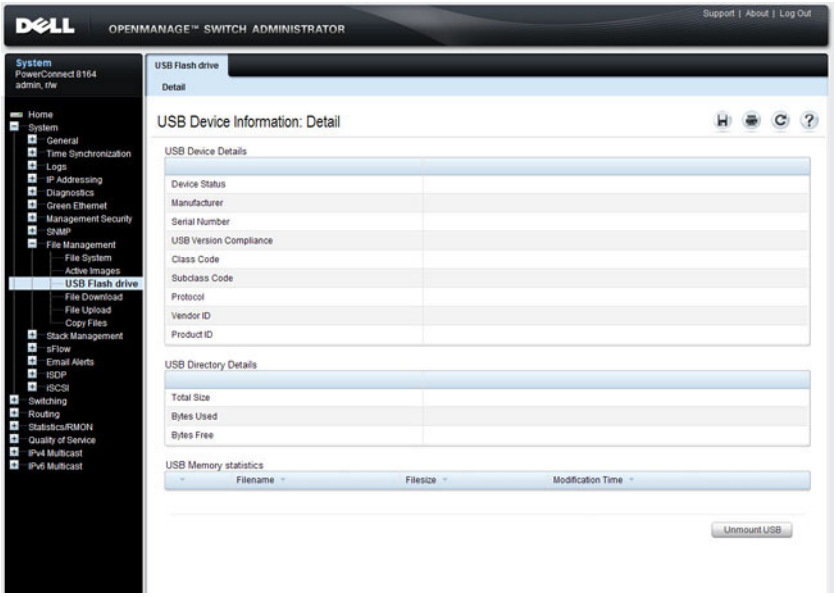


USB Flash Drive

Use the **USB Flash Drive** page (on PowerConnect 8100-series switches) to view information about a USB flash drive connected to the USB port on the front panel of the switch. The page also displays information about the files stored on the USB flash drive. To safely remove the USB flash drive from the USB port, click **Unmount USB** before removing the drive.

To display the **USB Flash Drive** page, click **System** → **File Management** → **USB Flash Drive** in the navigation panel.

Figure 14-3. USB Flash Drive

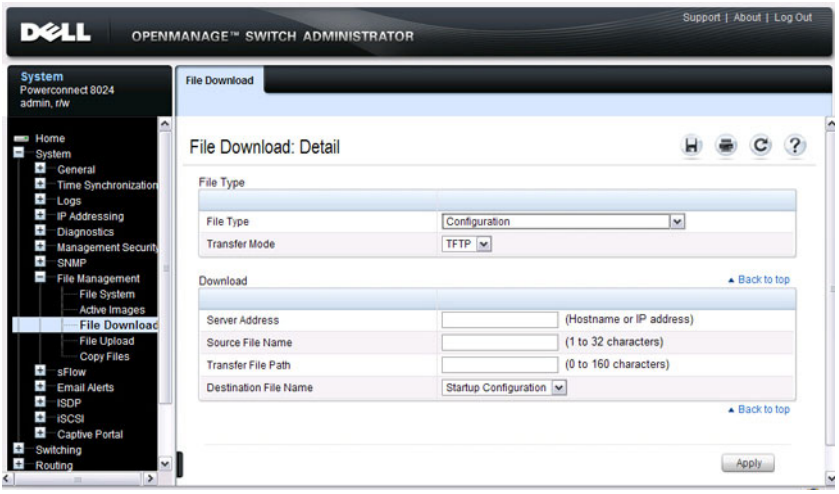


File Download

Use the **File Download** page to download image (binary) files, SSH and SSL certificates, IAS User files, and configuration (ASCII), files from a remote server to the switch.

To display the **File Download** page, click **System → File Management → File Download** in the navigation panel.

Figure 14-4. File Download



Downloading Files

To download a file to the switch:

1. Open the **File Download** page.
2. Select the type of file to download to the switch.
3. Select the transfer mode.

If you select a transfer mode that requires authentication, additional fields appear in the Download section. If you select HTTP as the download method, some of the fields are hidden.

NOTE: If you are using HTTPS to manage the switch, the download method will be HTTPS.

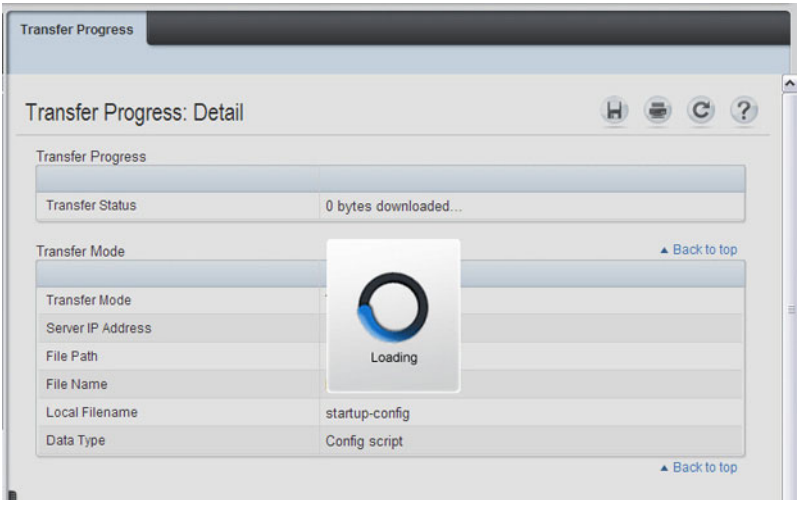
4 To download using HTTP, click **Browse** and select the file to download, then click **Apply**.

5 To download using any method other than HTTP, enter the IP address of the server that contains the file to download, the name of the file and the path on the server where it is located. For SFTP and SCP, provide the user name and password.

6 Click **Apply** to begin the download.

NOTE: After you start a file download, the page refreshes and a transfer status field appears to indicate the number of bytes transferred. The web interface is blocked until the file download is complete.

Figure 14-5. File Download in Progress



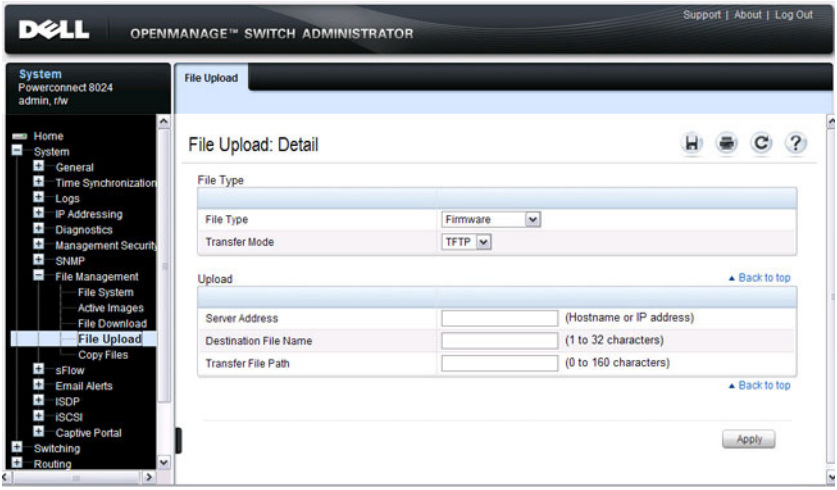
7 The file is downloaded to the switch.

File Upload

Use the **File Upload to Server** page to upload configuration (ASCII), image (binary), IAS user, operational log, and startup log files from the switch to a remote server.

To display the **File Upload to Server** page, click **System → File Management → File Upload** in the navigation panel.

Figure 14-6. File Upload



Uploading Files

To upload a file from the switch to a remote system:

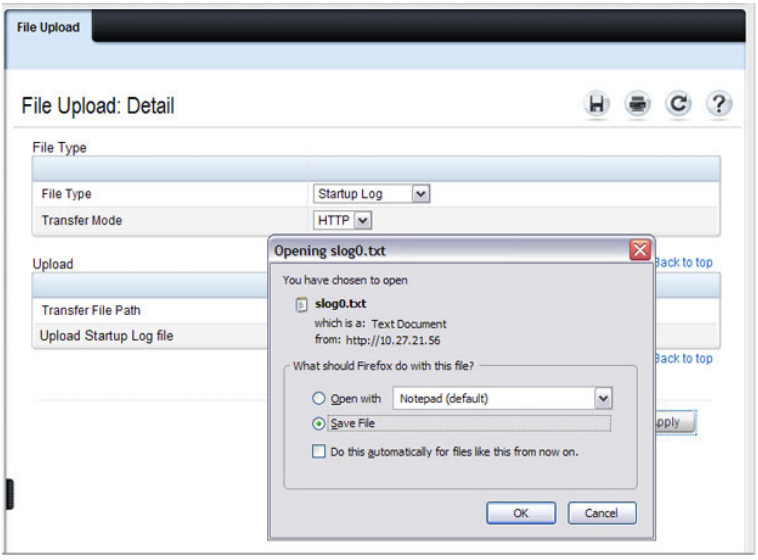
1. Open the **File Upload** page.
2. Select the type of file to download to the remote server.
3. Select the transfer mode.

   If you select a transfer mode that requires authentication, additional fields appear in the Upload section. If you select HTTP as the upload method, some of the fields are hidden.

NOTE: If you are using HTTPS to manage the switch, the download method will be HTTPS.

4 To upload by using HTTP, click **Apply**. A dialog box opens to allow you to open or save the file.

Figure 14-7. File Upload



5 To upload by using any method other than HTTP, enter the IP address of the server and specify a name for the file. For SFTP and SCP, provide the user name and password.

6 Click **Apply** to begin the upload.

NOTE: For some file uploads and methods, the page refreshes and a transfer status field appears to indicate the number of bytes transferred. The web interface is blocked until the file upload is complete.

7 The file is uploaded to the specified location on the remote server.

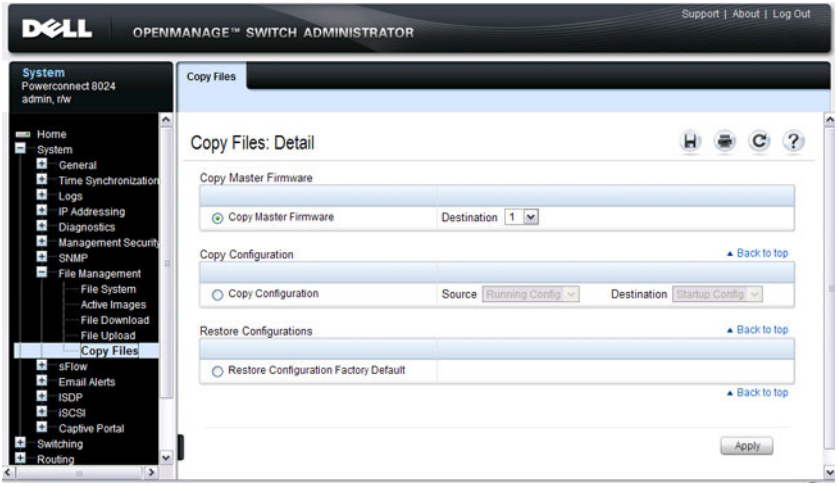
Copy Files

Use the **Copy Files** page to:

- Copy the active firmware image to the switch.
- Copy the running, startup, or backup configuration file to the startup or backup configuration file.
- Restore the running configuration to the factory default settings.

To display the **Copy Files** page, click **System** → **File Management** → **Copy Files** in the navigation panel.

Figure 14-8. Copy Files



Managing Images and Files (CLI)

This section provides information about the commands you use to upload, download, and copy files to and from the PowerConnect 8000-series and 8100-series switches. For more information about these commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals. It also describes the commands that control the Auto Configuration feature.

NOTE: Upload, download, and copy functions use the **copy** command. The basic syntax for the command is **copy** *source destination*. This section shows several different ways to use the **copy** command.

Downloading and Activating a New Image (TFTP)

Beginning in Privileged EXEC mode, use the following commands to download a new firmware image to the switch and to make it the active image. This example shows how to use TFTP to download the image.

Command	Purpose
copy tftp://{*ip-address*|*hostname*}/*path*/*file-name* **image**	Use TFTP to download the firmware image at the specified source to the non-active image. If the image file is in the TFTP file system root (download path), you do not need to specify the path in the command.
show bootvar	View the name of the image that is currently active (image1 or image2).
filedescr {**image1** | **image2**} *description*	Add a description to the image files.
boot system {**image1** | **image2**}	Set the image to use as the boot (active) image after the switch resets. Images on the PC8132/PC8164 are named *active* and *backup*.
reload	Reboot the switch to make the new image the active image. You are prompted to verify that you want to continue.

Managing Files in Internal Flash

Beginning in Privileged EXEC mode, use the following commands to copy, rename, delete and list the files in the internal flash.

Command	Purpose
dir	List the files in the flash file system.
rename *current_name new_name*	Rename a file in flash.
delete *filename*	Remove the specified file.
erase {**startup-config** | **backup-image** | **backup-config**}	Erase the startup configuration, the backup configuration or the backup image.
copy startup-config backup-config	Save the startup configuration to the backup configuration file.
copy running-config startup-config	Copy the current configuration to the startup configuration. This saves the current configuration to NVRAM.
show startup-config	View the contents of the startup-config file
show running-config	View the contents of the running-config file

Managing Files on a USB Flash Device (PowerConnect 8100-series switches only)

Beginning in Privileged EXEC mode, use the following commands to manage files that are on a USB device that is plugged into the USB flash port on the front panel of the switch.

Command	Purpose
show usb device	Display USB flash device details
dir usb	Display USB device contents and memory statistics
copy usb://*filename* {**backup-config** | **image** | **running-config** | **script** *filename*| **startup-config** | *filename*	Copy the specified file from the USB flash device to the specified file in internal flash.
unmount usb	Make the USB flash device inactive.

Uploading a Configuration File (SCP)

Beginning in Privileged EXEC mode, use the following commands to upload a configuration file from the switch to a remote system by using SCP.

Command	Purpose
copy *file* **scp:**//*user*@{*ip-address*|*hostname*}/*path*/*file-name*	Adds a description to an image file. The file can be one of the following files: • backup-config • image • operational-log • running-config • script *file-name* • startup-config • startup-log
Password entry	After you enter the **copy** command, the CLI prompts you for the password associated with the username.

Managing Configuration Scripts (SFTP)

Beginning in Privileged EXEC mode, use the following commands to download a configuration script from a remote system to the switch, validate the script, and activate it.



NOTE: The startup-config and backup-config files are essentially configuration scripts and can be validated and applied by using the commands in this section.

Command	Purpose
copy sftp://*user*@{*ip-address*|*hostname*}/*path*/*file-name* **script** *dest-name*	Downloads the specified script from the remote server to the switch.
Password entry	After you enter the **copy** command, the CLI prompts you for the password associated with the username.
script validate *script-name*	Checks the specified script for syntax errors. The script is automatically validated when you download it to the switch. You can validate again with this command.
script list	View the list of available scripts.
script activate *script-name*	Executes the commands within the script in order. The configuration changes in the script are applied to the running configuration.
script show *script-name*	View the contents of the specified script.

File and Image Management Configuration Examples

This section contains the following examples:

- Upgrading the Firmware
- Managing Configuration Scripts

Upgrading the Firmware

This example shows how to download a firmware image to the switch and activate it. The TFTP server in this example is PumpKIN, an open source TFTP server running on a Windows system.

- TFTP server IP address: 10.27.65.103
- File path: \image
- File name: dell_0308.stk

Use the following steps to prepare the download, and then download and upgrade the switch image.

1. Check the connectivity between the switch and the TFTP server.

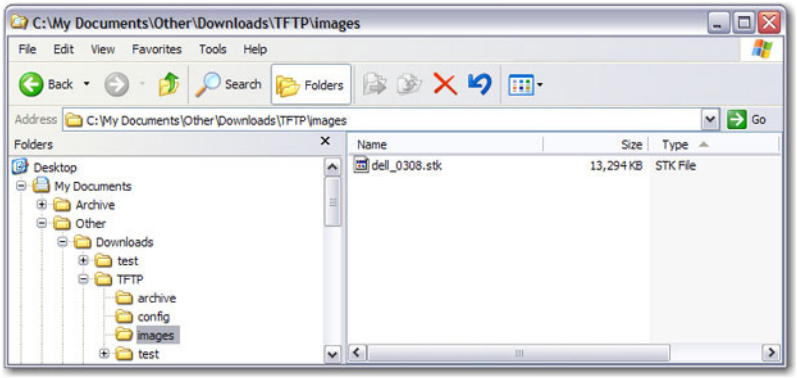
```
console#ping 10.27.65.103
Pinging 10.27.65.103 with 0 bytes of data:

Reply From 10.27.65.103: icmp_seq = 0. time <10 msec.
Reply From 10.27.65.103: icmp_seq = 1. time <10 msec.
Reply From 10.27.65.103: icmp_seq = 2. time <10 msec.
Reply From 10.27.65.103: icmp_seq = 3. time <10 msec.

----10.27.65.103 PING statistics----
4 packets transmitted, 4 packets received, 0% packet
loss
round-trip (msec) min/avg/max = <10/<10/<10
```

2. Copy the image file to the appropriate directory on the TFTP server. In this example, the TFTP root directory is `C:\My Documents\Other\Downloads\TFTP`, so the file path is `images`.

Figure 14-9. Image Path



3 View information about the current image.

```
console#show bootvar
Image Descriptions

 image1 :
 image2 :

Images currently available on Flash

------- ----------- ----------- -------------- --------------
unit    image1      image2      current-active next-active
------- ----------- ----------- -------------- --------------
1       2.23.11.17              image1         image1
```

4 Download the image to the switch. After you execute the **copy** command, you must verify that you want to start the download.

```
console#copy
tftp://10.27.65.103/images/dell_0308.stk image

Mode.............................................. TFTP
Set TFTP Server IP................................ 10.27.65.103
TFTP Path......................................... images/
TFTP Filename..................................... dell_0308.stk
Data Type......................................... Code
Destination Filename.............................. image
```

```
Management access will be blocked for the duration of the
transfer
Are you sure you want to start? (y/n)y
```

5 Activate the new image (image2) so that it becomes the active image after the switch resets.

```
console#boot system image2
Activating image image2..
```

6 View information about the current image.

```
console#show bootvar
Image Descriptions

 image1 :
 image2 :

Images currently available on Flash

------- ----------- ----------- -------------- --------------
unit    image1      image2      current-active next-active
------- ----------- ----------- -------------- --------------
1       2.23.11.17  3.8.18.37   image1         image2
```

7 Copy the running configuration to the startup configuration to save the current configuration to NVRAM.

```
console#copy running-config startup-config

This operation may take a few minutes.
Management interfaces will not be available during
this time.

Are you sure you want to save? (y/n)y

Configuration Saved!
```

8 Reset the switch to boot the system with the new image.

```
console#reload

Are you sure you want to continue? (y/n)y

Reloading all switches...
```

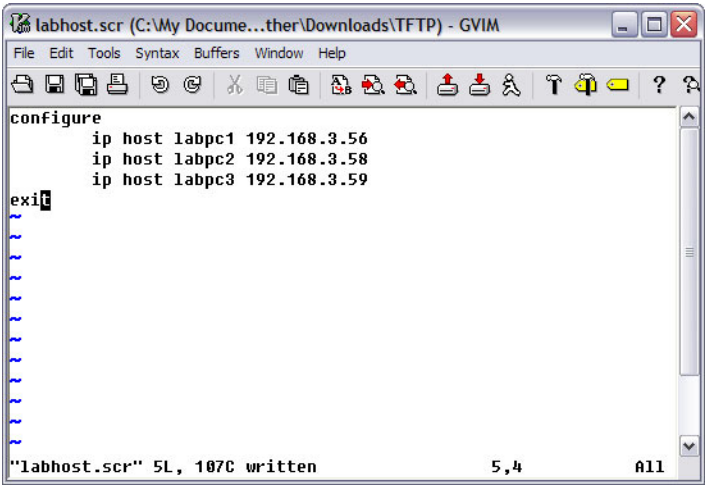
Managing Configuration Scripts

This example shows how to create a configuration script that adds three hostname-to-IP address mappings to the host table.

To configure the switch:

1 Open a text editor on an administrative computer and type the commands as if you were entering them by using the CLI.

Figure 14-10. Create Config Script



2 Save the file with an *.scr extension and copy it to the appropriate directory on your TFTP server.

3 Download the file from the TFTP server to the switch.

```
console#copy tftp://10.27.65.103/labhost.scr
script labhost.scr

Mode.......................................... TFTP
Set TFTP Server IP............................ 10.27.65.103
TFTP Path..................................... ./
TFTP Filename................................. labhost.scr
Data Type..................................... Config Script
Destination Filename.......................... labhost.scr
```

```
Management access will be blocked for the duration of the
transfer
```

4 After you confirm the download information and the script successfully downloads, it is automatically validated for correct syntax.

```
Are you sure you want to start? (y/n) y

135 bytes transferred

Validating configuration script...
configure
exit
configure
ip host labpc1 192.168.3.56

ip host labpc2 192.168.3.58

ip host labpc3 192.168.3.59

Configuration script validated.
File transfer operation completed successfully.
```

5 Run the script to execute the commands.

```
console#script apply labhost.scr

Are you sure you want to apply the configuration
script? (y/n)y

configure
exit
configure
ip host labpc1 192.168.3.56

ip host labpc2 192.168.3.58

ip host labpc3 192.168.3.59

Configuration script 'labhost.scr' applied.
```

6 Verify that the script was successfully applied.

```
console#show hosts

Host name: test
Name/address lookup is enabled
Name servers (Preference order): 192.168.3.20
Configured host name-to-address mapping:
Host                        Addresses
--------------------------  ------------------------
labpc1                      192.168.3.56
labpc2                      192.168.3.58
labpc3                      192.168.3.59
```

Managing Files by Using the USB Flash Drive (PowerConnect 8100-series switches only)

In this example, the administrator copies the backup image to a USB flash drive before overwriting the backup image on the switch with a new image. The administrator also makes a backup copy of the running-config by uploading it to a USB flash drive. After the backups are performed, the administrator downloads a new image from the USB flash drive to the switch to prepare for the upgrade.

This example assumes the new image is named new_img.stk and has already been copied from an administrative host onto the USB flash drive.

To configure the switch:

1 Insert the USB flash drive into the USB port on the front panel of the switch. The USB flash drive is automatically mounted.

2 Copy the backup image from the switch to the USB flash drive.

```
console#copy image usb://img_backup.stk

Mode..................................... unknown
Data Type................................ Code

Management access will be blocked for the duration
of the transfer
Are you sure you want to start? (y/n) y
```

3 Copy the running-config to the USB flash drive.

```
console#copy running-config usb://rc_backup.scr

Mode............................ unknown
Data Type....................... Config Script
Source Filename................. temp-config.scr

Management access will be blocked for the duration
of the transfer
Are you sure you want to start? (y/n) y
```

4 Download the new image from the USB flash drive to the switch. The image overwrites the image that is not currently active.

```
console#copy usb://new_image.stk image

Mode................................ unknown
Data Type........................... Code

Management access will be blocked for the duration
of the transfer
Are you sure you want to start? (y/n) y
```

5 To activate the new image after it has been successfully downloaded to the switch, follow the procedures described in Upgrading the Firmware, starting with step 5.

15

Automatically Updating the Image and Configuration

The topics covered in this chapter include:

- Auto Configuration Overview
- What Are the Dependencies for DHCP Auto Configuration?
- Default Auto Configuration Values
- Managing Auto Configuration (Web)
- Managing Auto Configuration (CLI)
- Auto Configuration Example

Auto Configuration Overview

The Auto Configuration feature can automatically update the firmware image and obtain configuration information when the switch boots. Auto Configuration begins the automatic download and installation process when the switch or stack master is initialized and no configuration file (startup-config) is found, or when the switch boots and loads a saved configuration that has Auto Configuration enabled. Auto Configuration is enabled by default. Allow downgrade is also enabled by default.

The Auto Configuration feature includes two components:

- USB Auto Configuration
- DHCP Auto Install

If no configuration file is found and the Auto Configuration feature is enabled, the Auto Configuration process begins. If a USB device is connected to the PowerConnect switch USB port and contains the appropriate file, the switch uses the USB Auto Configuration feature to update the configuration or image. If the USB Auto Configuration fails - either because it is disabled,

no USB storage device is present, or no configuration or images files are present on the USB storage device, the switch uses the DHCP Auto Install process.

NOTE: Neither USB Configuration nor Auto Install is invoked if a valid configuration file is on the switch.Auto Install is not invoked if a valid configuration file is on the switch.

What Is USB Auto Configuration?

You can use the USB Auto Configuration feature to configure or upgrade one or more switches that have not been previously configured, such as when you deploy new switches. Before you deploy the switch, you perform the following steps:

1. Create a text file that contains IP addresses (and/or MAC addresses) and file names that are parsed and used by this feature. The MAC address used to identify the switch is the MAC address of the out-of-band interface (see "Switch Addresses" on page 97 for further information).
2. Copied the file onto a USB device.
3. Insert the USB device into the front-panel USB port on the PowerConnect switch.

When the Auto Configuration process starts and a factory default or empty configuration file is present on the switch, the feature automatically searches a plugged-in USB device for information.

What Files Does USB Auto Configuration Use?

The USB Auto Configuration feature uses the following file types:

- *.setup file for initial switch configuration
- *.text file for configuration information
- *.stk file for software image installation

The Auto Configuration file searches the USB device for a file with a *.setup extension. If only one .setup file is present, the switch uses the file. When multiple *.setup files are present, the switch uses only the powerconnnect.setup file. If no powerconnect.setup file is available, the switch checks for a file with a *.text configuration file and a *.stk image file. If multiple .text files exist, the switch uses the powerconnect.text file. If

multiple *.stk files are present, the switch uses the image with the highest (most recent) version. Finally, if no *.setup, *.text, or *.stk files are found, the switch proceeds to the DHCP Auto Configuration process.

What Is the DHCP Auto Configuration Process?

The switch can use a DHCP server to obtain configuration information from a TFTP server.

DHCP Auto Configuration is accomplished in three phases:

1. Assignment or configuration of an IP address for the switch
2. Assignment of a TFTP server
3. Obtaining a configuration file for the switch from the TFTP server

Auto Configuration is successful when an image or configuration file is downloaded to the switch or stack master from a TFTP server.

NOTE: The downloaded configuration file is not automatically saved to startup-config. You must explicitly issue a save request (**copy running-config startup-config**) in order to save the configuration.

Obtaining IP Address Information

DHCP is enabled by default on the Out-of-Band (OOB) interface. If an IP address has not been assigned, the switch issues requests for an IP address assignment.

A network DHCP server returns the following information:

- IP address and subnet mask to be assigned to the interface
- IP address of a default gateway, if needed for IP communication

After an IP address is assigned to the switch, if a hostname is not already assigned, Auto Configuration issues a DNS request for the corresponding hostname. This hostname is also displayed as the CLI prompt (as in response to the **hostname** command).

Obtaining Other Dynamic Information

The following information is also processed and may be returned by a BOOTP or DHCP server:

- Name of configuration file (the *file* field in the DHCP header or option 67) to be downloaded from the TFTP server.
- Identification of the TFTP server providing the file. The TFTP server can be identified by name or by IP address as follows:
  - hostname: DHCP option 66 or the *sname* field in the DHCP header)
  - IP address: DHCP option 150 or the *siaddr* field in the DHCP header

When a DHCP OFFER identifies the TFTP server more than once, the DHCP client selects one of the options in the following order: *sname*, option 66, option 150, *siaddr*. If the TFTP server is identified by hostname, a DNS server is required to translate the name to an IP address.

The DHCP client on the switch also processes the name of the text file (option 125, the V-I vendor-specific Information option) which contains the path to the image file.

Obtaining the Image

Auto Configuration attempts to download an image file from a TFTP server only if no configuration file was found in the internal flash, or even with a saved configuration file that has Auto Configuration enabled.

The network DHCP server returns a DHCP OFFER message with option 125. When configuring the network DHCP server for image downloads, you must include Option 125 and specify the Dell Enterprise Number, 674. Within the Dell section of option 125, sub option 5 must specify the path and name of a file on the TFTP server. This file is not the image file itself, but rather a text file that contains the path and name of the image file. Upon receipt of option 125, the switch downloads the text file from the TFTP server, reads the name of the image file, and downloads the image file from the TFTP server.

After the switch successfully downloads and installs the new image, it automatically reboots. The download or installation might fail for one of the following reasons:

- The path or filename of the image on the TFTP server does not match the information specified in DHCP option 125.
- The downloaded image is the same as the current image.
- The validation checks, such as valid CRC Checksum, fails.

If the download or installation was unsuccessful, a message is logged.

Obtaining the Configuration File

If the DHCP OFFER identifies a configuration file, either as option 67 or in the *file* field of the DHCP header, the switch attempts to download the configuration file.

NOTE: The configuration file is required to have a file type of *.cfg.

The TFTP client makes three unicast requests. If the unicast attempts fail, or if the DHCP OFFER did not specify a TFTP server address, the TFTP client makes three broadcast requests.

If the DHCP server does not specify a configuration file or download of the configuration file fails, the Auto Configuration process attempts to download a configuration file with the name `dell-net.cfg`. The switch unicasts or broadcasts TFTP requests for a network configuration file in the same manner as it attempts to download a host-specific configuration file.

The default network configuration file consists of a set of IP address-to-hostname mappings, using the command **ip host hostname** *address*. The switch finds its own IP address, as learned from the DHCP server, in the configuration file and extracts its hostname from the matching command. If the default network configuration file does not contain the switch's IP address, the switch attempts a reverse DNS lookup to resolve its hostname.

A sample `dell-net.cfg` file follows:

```
config
  ...
  ip host switch1 192.168.1.10
  ip host switch2 192.168.1.11
  ... <other hostname definitions>
exit
```

Once a hostname has been determined, the switch issues a TFTP request for a file named *`hostname`*`.cfg` , where *`hostname`* is the first thirty-two characters of the switch's hostname.

If the switch is unable to map its IP address to a hostname, Auto Configuration sends TFTP requests for the default configuration file `host.cfg`.

Table 15-1 summarizes the config files that may be downloaded and the order in which they are sought.

Table 15-1. Configuration File Possibilities

Order Sought	File Name	Description	Final File Sought
1	`bootfile.cfg`	Host-specific config file, ending in a *.cfg file extension	Yes
2	`dell-net.cfg`	Default network config file	No
3	`hostname.cfg`	Host-specific config file, associated with hostname.	Yes
4	`host.cfg`	Default config file	Yes

Table 15-2 displays the determining factors for issuing unicast or broadcast TFTP requests.

Table 15-2. TFTP Request Types

TFTP Server Address Available	Host-specific Switch Config Filename Available	TFTP Request Method
Yes	Yes	Issue a unicast request for the host-specific router config file to the TFTP server
Yes	No	Issue a unicast request for a default network or router config file to the TFTP server
No	Yes	Issue a broadcast request for the host-specific router config file to any available TFTP server
No	No	Issue a broadcast request for the default network or router config file to any available TFTP server

Monitoring and Completing the DHCP Auto Configuration Process

When the switch boots and triggers an Auto Configuration, a message displays on the console screen to indicate that the process is starting. After the process completes, the Auto Configuration process writes a log message. When Auto Configuration has successfully completed, you can execute a **show running-config** command to validate the contents of configuration.

Saving a Configuration

The Auto Configuration feature includes an AutoSave capability that allows the downloaded configuration to be automatically saved; however, AutoSave is disabled by default. If AutoSave has not been enabled, you must explicitly save the downloaded configuration in nonvolatile memory on the stack master. This makes the configuration available for the next reboot. In the CLI, this is performed by issuing a **write** command or **copy running-config startup-config** command and should be done after validating the contents of saved configuration.

Stopping and Restarting the Auto Configuration Process

You can terminate the Auto Configuration process at any time before the image or configuration file is downloaded. This is useful when the switch is disconnected from the network. Termination of the Auto Configuration process ends further periodic requests for a host-specific file.

The Auto Configuration process automatically starts after a reboot if the configuration file is not found on the switch. The configuration file will not be found if it has never been saved on the switch, or if you issue a command to erase the configuration file (**clear config** or **erase startup-config**).

Managing Downloaded Config Files

The configuration files downloaded by Auto Configuration are stored in the nonvolatile memory as .scr files. The files may be managed (viewed or deleted) along with files downloaded by the configuration scripting utility.

A file is not automatically deleted after it is downloaded. The file does not take effect upon a reboot unless you explicitly save the configuration (the saved configuration takes effect upon reboot). If you do not save the configuration downloaded by the Auto Configuration feature, the Auto Configuration process occurs again on a subsequent reboot. This may result in one of the previously downloaded files being overwritten.

What Are the Dependencies for DHCP Auto Configuration?

The Auto Configuration process from TFTP servers depends upon the following network services:

- A DHCP server must be configured on the network with appropriate services.
- An image file and a text file containing the image file name for the switch must be available from a TFTP server if DHCP image download is desired.
- A configuration file (either from bootfile (or) option 67 option) for the switch must be available from a TFTP server.
- The switch must be connected to the network and have a Layer 3 interface that is in an UP state.
- A DNS server must contain an IP address to hostname mapping for the TFTP server if the DHCP server response identifies the TFTP server by name.
- A DNS server must contain an IP address to hostname mapping for the switch if a `<hostname>.cfg` file is to be downloaded.
- If a default gateway is needed to forward TFTP requests, an IP helper address for TFTP needs to be configured on the default gateway.

Default Auto Configuration Values

Table 15-3 describes the Auto Configuration defaults.

Table 15-3. Auto Configuration Defaults

Feature	Default	Description
Auto Install Mode	Enabled	When the switch boots and no saved configuration is found, the Auto Configuration automatically begins.
Retry Count	3	When the DHCP or BootP server returns information about the TFTP server and bootfile, the switch makes three unicast TFTP requests for the specified bootfile. If the unicast attempts fail or if a TFTP server address was not provided, the switch makes three broadcast requests to any available TFTP server for the specified bootfile.
AutoSave	Disabled	If the switch is successfully auto-configured, the running configuration is not saved to the startup configuration.
AutoReboot	Enabled	After an image is successfully downloaded during the Auto Configuration process, the switch automatically reboots and makes the downloaded image the active image.

Managing Auto Configuration (Web)

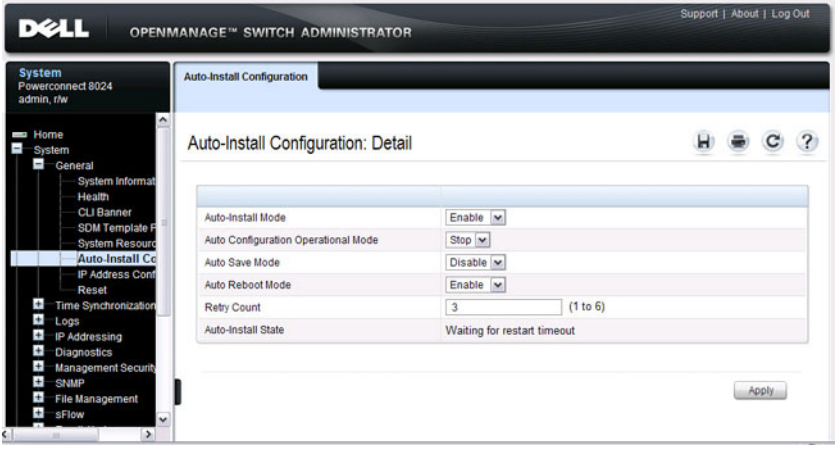
This section provides information about the OpenManage Switch Administrator pages to use to manage images and files on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

Auto-Install Configuration

Use the **Auto-Install Configuration** page to allow the switch to obtain network information (such as the IP address and subnet mask) and automatically download a host-specific or network configuration file during the boot process if no startup-config file is found.

To display the **Auto Configuration** page, click **System** → **General** → **Auto-Install Configuration** in the navigation panel.

Figure 15-1. Auto-Install Configuration



Managing Auto Configuration (CLI)

This section provides information about the commands you manage the Auto-Install Configuration feature on the switch. For more information about these commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Managing Auto Configuration

Beginning in Privileged EXEC mode, use the following commands to manually activate the Auto Configuration process and download a configuration script from a remote system to the switch, validate the script, and activate it.

NOTE: The Auto Configuration feature begins automatically when the switch is booted and no startup-config file is found or if the system boots and finds the **boot host dhcp** command in the startup-config file.

Command	Purpose
configure	Enter Global Configuration mode.
boot autoinstall start	Enable the Auto Configuration feature on the switch.
boot host dhcp	Enable Auto Configuration for the next reboot cycle. The command does not change the current behavior of Auto Configuration, but it does save the command to NVRAM.
boot host autosave	Allow the switch to automatically save the configuration file downloaded to the switch by the Auto Configuration feature.
boot host retrycount *retries*	Specify the number of attempts to download the file (by sending unicast TFTP requests, and if unsuccessful, broadcast TFTP requests) specified in the response from the DHCP server. The range for *retries* is 1–3.
boot host autoreboot	Allow the switch to automatically reboot when the image is successfully downloaded through the Auto Configuration feature.
exit	Exit to Privileged Exec mode.
show boot	Displays the current status of the Auto Configuration process.

Auto Configuration Example

A network administrator is deploying three PowerConnect switches and wants to quickly and automatically install the latest image and a common configuration file that configures basic settings such as VLAN creation and membership, RADIUS server settings, and 802.1X information. The configuration file also contains the command **boot host autosave** so that the downloaded configuration is automatically saved to the startup config.

Enabling DHCP Auto Configuration and Auto Image Download

If no configuration file is found during the boot process, the Auto Configuration feature uses the DHCP Auto Configuration process to download the configuration file to the switch. This example describes the procedures to complete the configuration.

To use DHCP auto configuration:

1. Create a default config file for the switches named `host.cfg`. For information about creating configuration files, see Managing Images and Files.
2. Upload the `host.cfg` file to the TFTP server.
3. Upload the image file to the TFTP server.
4. Configure an address pool on the DHCP server that contains the following information:
   a. The IP address (*yiaddr*) and subnet mask (option 1) to be assigned to the interface
   b. The IP address of a default gateway (option 3)
   c. DNS server address (option 6)
   d. Name of config file for each host
   e. Identification of the TFTP server by hostname (DHCP option 66 or the *sname* field in the DHCP header) or IP address (DHCP option 150 or the *siaddr* field in the DHCP header)
   f. Name of the text file (option 125, the V-I vendor-specific Information option) that contains the path to the image file.

5 Connect a port (OOB port for out-of-band management or any switch port for in-band management) on each switch to the network.

6 Boot the switches.

16

Monitoring Switch Traffic

This chapter describes sFlow features, Remote Monitoring (RMON), and Port Mirroring features.

The topics covered in this chapter include:

- Traffic Monitoring Overview
- Default Traffic Monitoring Values
- Monitoring Switch Traffic (Web)
- Monitoring Switch Traffic (CLI)
- Traffic Monitoring Configuration Examples

Traffic Monitoring Overview

The switch maintains statistics about network traffic that it handles. It also has embedded technology that collects and sends information about traffic to other devices. PowerConnect 8000-series and 8100-series switches include support for flow-based monitoring through sFlow and Remote Network Monitoring (RMON) agents.

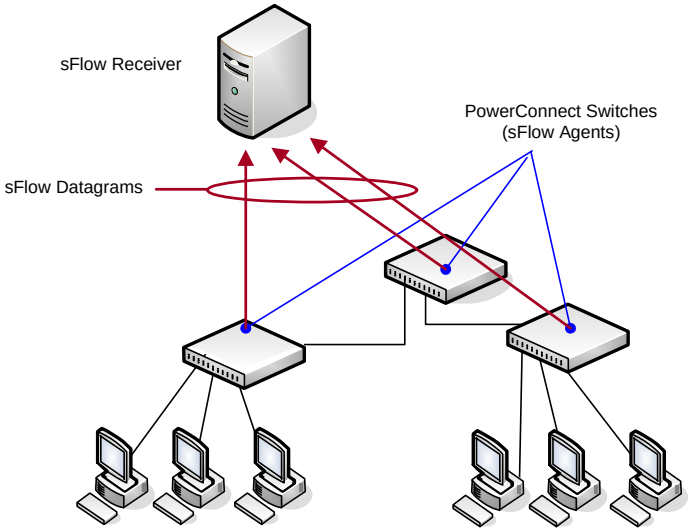
What is sFlow Technology?

sFlow is an industry standard technology for monitoring high-speed switched and routed networks. PowerConnect 8000-series and 8100-series switches software has a built-in sFlow agent that can monitor network traffic on each port and generate sFlow data to an sFlow receiver (also known as a collector). sFlow helps to provide visibility into network activity, which enables effective management and control of network resources. sFlow is an alternative to the NetFlow network protocol, which was developed by Cisco Systems. The switch supports sFlow version 5.

As illustrated in Figure 16-1, the sFlow monitoring system consists of sFlow Agents (such as PowerConnect 8000-series and 8100-series switches) and a central sFlow receiver. sFlow Agents use sampling technology to capture

traffic statistics from monitored devices. sFlow datagrams forward sampled traffic statistics to the sFlow Collector for analysis. You can specify up to eight different sFlow receivers to which the switch sends sFlow datagrams.

Figure 16-1. sFlow Architecture



The advantages of using sFlow are:

- It is possible to monitor all ports of the switch continuously, with no impact on the distributed switching performance.
- Minimal memory/CPU is required. Samples are not aggregated into a flow-table on the switch; they are forwarded immediately over the network to the sFlow receiver.
- The sFlow system is tolerant to packet loss in the network because statistical modeling means the loss is equivalent to a slight change in the sampling rate.
- sFlow receiver can receive data from multiple switches, providing a real-time synchronized view of the whole network.
- The receiver can analyze traffic patterns based on protocols found in the headers (e.g., TCP/IP, IPX, Ethernet, AppleTalk…). This alleviates the need for a layer 2 switch to decode and understand all protocols.

sFlow Sampling

The sFlow Agent in the PowerConnect software uses two forms of sampling:

- Statistical packet-based sampling of switched or routed Packet Flows
- Time-based sampling of counters

Packet Flow Sampling and Counter Sampling are performed by sFlow Instances associated with individual Data Sources within an sFlow Agent. Both types of samples are combined in sFlow datagrams. Packet Flow Sampling creates a steady, but random, stream of sFlow datagrams that are sent to the sFlow Collector. Counter samples may be taken opportunistically to fill these datagrams.

To perform Packet Flow Sampling, an sFlow Sampler Instance is configured with a Sampling Rate. Packet Flow sampling results in the generation of Packet Flow Records. To perform Counter Sampling, an sFlow Poller Instance is configured with a Polling Interval. Counter Sampling results in the generation of Counter Records. sFlow Agents collect Counter Records and Packet Flow Records and send them as sFlow datagrams to sFlow Collectors.

Packet Flow Sampling

Packet Flow Sampling, carried out by each sFlow instance, ensures that any packet observed at a Data Source has an equal chance of being sampled, irrespective of the Packet Flow(s) to which it belongs.

Packet Flow Sampling is accomplished as follows:

- A packet arrives on an interface.
- The Network Device makes a filtering decision to determine whether the packet should be dropped.
- If the packet is not filtered (dropped) a destination interface is assigned by the switching/routing function.
- A decision is made on whether or not to sample the packet.

  The mechanism involves a counter that is decremented with each packet. When the counter reaches zero a sample is taken.

- When a sample is taken, the counter indicating how many packets to skip before taking the next sample is reset. The value of the counter is set to a random integer where the sequence of random integers used over time is the Sampling Rate.

Counter Sampling

The primary objective of Counter Sampling is to efficiently, periodically export counters associated with Data Sources. A maximum Sampling Interval is assigned to each sFlow instance associated with a Data Source.

Counter Sampling is accomplished as follows:

- sFlow Agents keep a list of counter sources being sampled.
- When a Packet Flow Sample is generated the sFlow Agent examines the list and adds counters to the sample datagram, least recently sampled first. Counters are only added to the datagram if the sources are within a short period, 5 seconds say, of failing to meet the required Sampling Interval.
- Periodically, say every second, the sFlow Agent examines the list of counter sources and sends any counters that must be sent to meet the sampling interval requirement.

The set of counters is a fixed set.

What is RMON?

Like sFlow, RMON is a technology that enables the collection and analysis of a variety of data about network traffic. PowerConnect 8000-series and 8100-series switches software includes an RMON probe (also known as an RMON agent) that collect information and analyze packets. The data that is collected is defined in the RMON MIB, RFC 2819.

RMON is defined in an Internet Engineering Task Force (IETF) specification and is an extension of the SNMP MIB. You can view the RMON information locally on the switch or by using a generic RMON console on a network management station (NMS). SNMP does not need to be configured on the switch to view the RMON data locally. However, if you use a management station to view the RMON data that the switch collects and analyzes, you must configure the following SNMP settings:

- Set up the SNMP community string to be used by the SNMP manager at a given IP address.
- Specify the network management system IP address or permit management access from all IP addresses.

For more information about configuring SNMP, see "Configuring SNMP" on page 271.

The RMON agent in the switch supports the following groups:

- Group 1—Statistics. Contains cumulative traffic and error statistics.
- Group 2—History. Generates reports from periodic traffic sampling that are useful for analyzing trends.
- Group 3 —Alarm. Enables the definition and setting of thresholds for various counters. Thresholds can be passed in either a rising or falling direction on existing MIB objects, primarily those in the Statistics group. An alarm is triggered when a threshold is crossed and the alarm is passed to the Event group. The Alarm requires the Event Group.
- Group 9 —Event. Controls the actions that are taken when an event occurs. RMON events occur when:
  - A threshold (alarm) is exceeded
  - There is a match on certain filters.

NOTE: The switch supports RMON1.

What is Port Mirroring?

Port mirroring is used to monitor the network traffic that a port sends and receives. The Port Mirroring feature creates a copy of the traffic that the source port handles and sends it to a destination port. The source port is the port that is being monitored. The destination port is monitoring the source port. The destination port is where you would connect a network protocol analyzer to learn more about the traffic that is handled by the source port.

A port monitoring session includes one or more source ports that mirror traffic to a single destination port. The PowerConnect 8000-series and 8100-series switches support a single port monitoring session. LAGs (port channels) cannot be used as the source or destination ports.

For each source port, you can specify whether to mirror ingress traffic (traffic the port receives, or RX), egress traffic (traffic the port sends, or TX), or both ingress and egress traffic.

NOTE: You can create a DiffServ policy class definition that mirrors specific types of traffic to a destination port. For more information, see "Configuring Differentiated Services" on page 1103.

The packet that is copied to the destination port is in the same format as the original packet on the wire. This means that if the mirror is copying a received packet, the copied packet is VLAN tagged or untagged as it was received on the source port. If the mirror is copying a transmitted packet, the copied packet is VLAN tagged or untagged as it is being transmitted on the source port.

After you configure the port mirroring session, you can enable or disable the administrative mode of the session to start or stop the probe port from receiving mirrored traffic.

Why is Traffic Monitoring Needed?

Monitoring the traffic that the switch handles, as well as monitoring all traffic in the network, can help provide information about network performance and utilization. This information can be useful in network planning and resource allocation. Information about traffic flows can also help troubleshoot problems in the network.

Default Traffic Monitoring Values

The sFlow agent is enabled by default, but sampling and polling are disabled on all ports. Additionally, no sFlow receivers (collectors) are configured. Table 16-1 contains additional default values for the sFlow feature.

Table 16-1. sFlow Defaults

Parameter	Default Value
Receiver timeout for sampling	0
Receiver port	6343
Receiver Maximum Datagram Size	1400 bytes
Maximum header size	128 bytes

RMON is enabled by default, but no RMON alarms, events, or history statistic groups are configured.

Port mirroring is disabled, and no ports are configured as source or destination ports. After you configure a port mirroring session, the administrative mode is disabled until you explicitly enable it.

Monitoring Switch Traffic (Web)

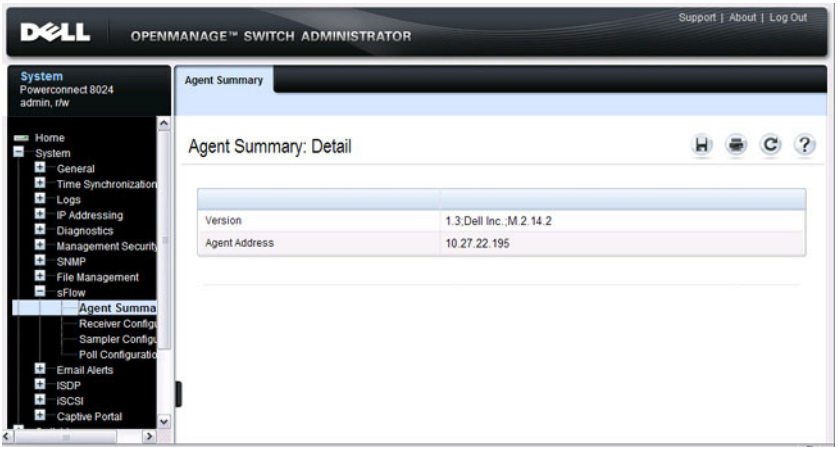
This section provides information about the OpenManage Switch Administrator pages to use to monitor network traffic on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

sFlow Agent Summary

Use the sFlow **Agent Summary** page to view information about sFlow MIB and the sFlow Agent IP address.

To display the **Agent Summary** page, click **System** → **sFlow** → **Agent Summary** in the navigation panel.

Figure 16-2. sFlow Agent Summary

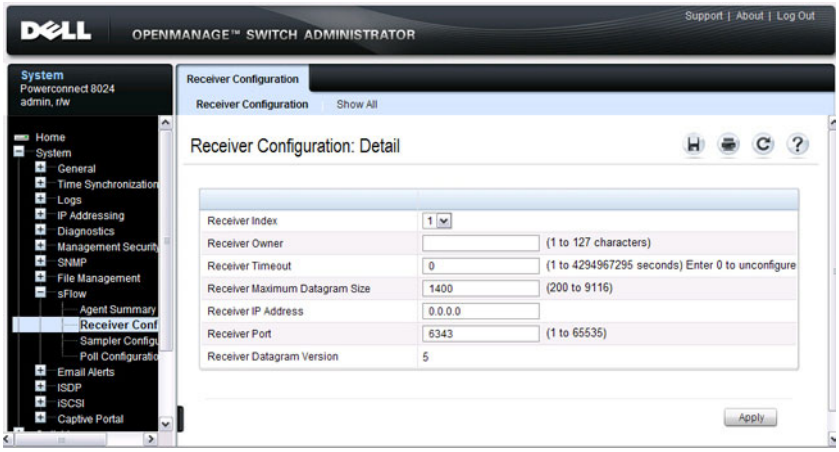


sFlow Receiver Configuration

Use the sFlow **Receiver Configuration** page to configure settings for the sFlow receiver to which the switch sends sFlow datagrams. You can configure up to eight sFlow receivers that will receive datagrams.

To display the **Receiver Configuration** page, click **System** → **sFlow** → **Receiver Configuration** in the navigation panel.

Figure 16-3. sFlow Receiver Configuration



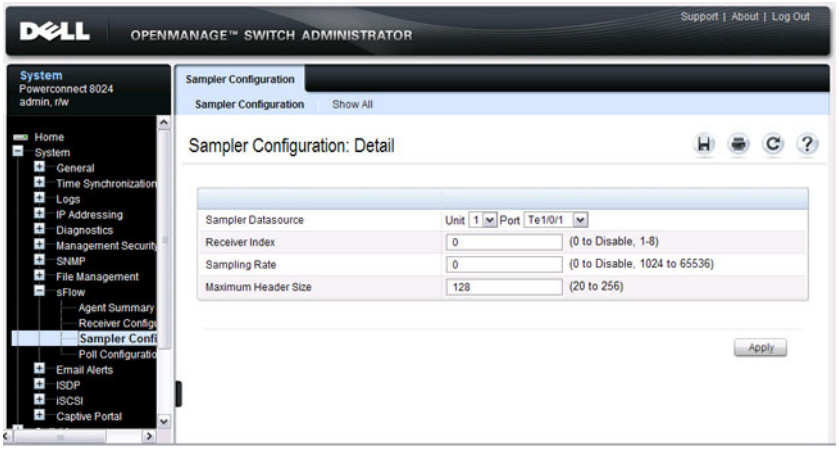
Click **Show All** to view information about configured sFlow receivers.

sFlow Sampler Configuration

Use the sFLow **Sampler Configuration** page to configure the sFlow sampling settings for switch ports.

To display the **Sampler Configuration** page, click **System** → **sFlow** → **Sampler Configuration** in the navigation panel.

Figure 16-4. sFlow Sampler Configuration



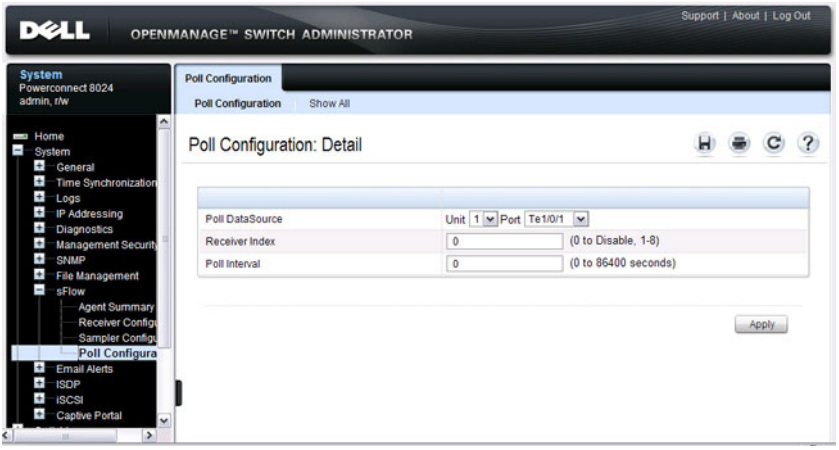
Click **Show All** to view information about configured sampler data sources.

sFlow Poll Configuration

Use the sFLow **Poll Configuration** page to configure how often a port should collect counter samples.

To display the **Sampler Configuration** page, click **System** → **sFlow** → **Sampler Configuration** in the navigation panel.

Figure 16-5. sFlow Poll Configuration



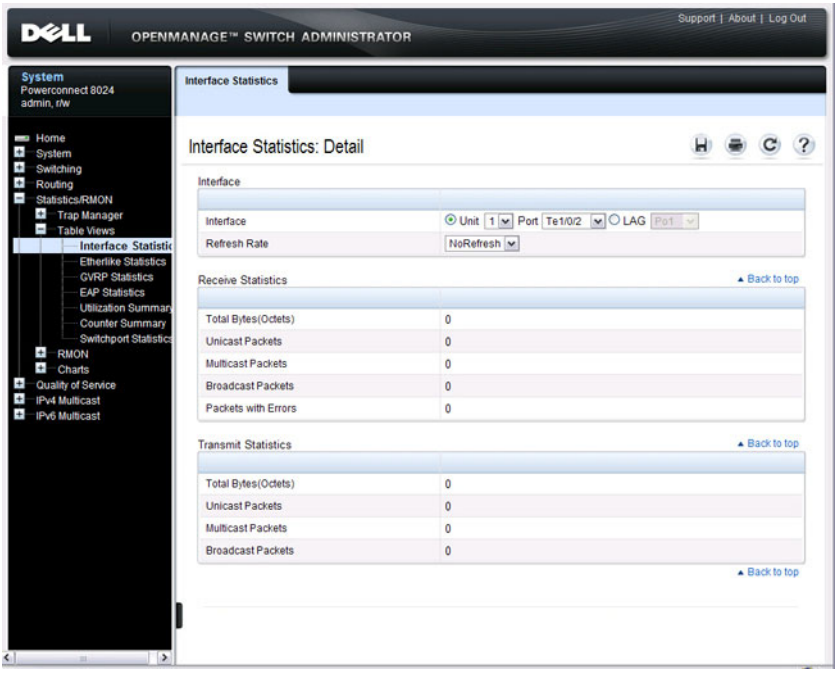
Click **Show All** to view information about the ports configured to collect counter samples.

Interface Statistics

Use the **Interface Statistics** page to display statistics for both received and transmitted packets. The fields for both received and transmitted packets are identical.

To display the page, click **Statistics/RMON → Table Views → Interface Statistics** in the navigation panel.

Figure 16-6. Interface Statistics

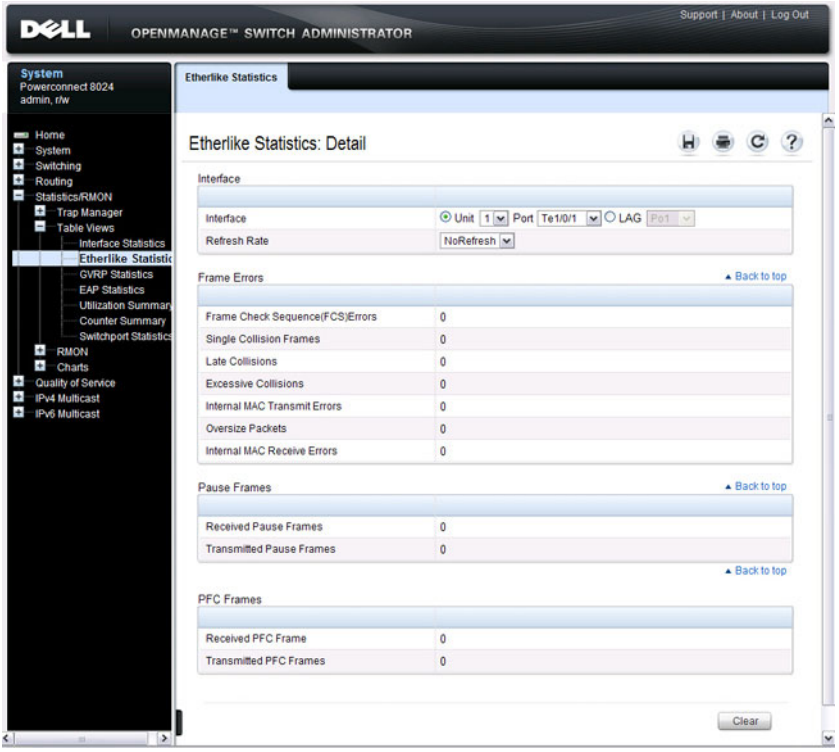


Etherlike Statistics

Use the **Etherlike Statistics** page to display interface statistics.

To display the page, click **Statistics/RMON → Table Views → Etherlike Statistics** in the navigation panel.

Figure 16-7. Etherlike Statistics

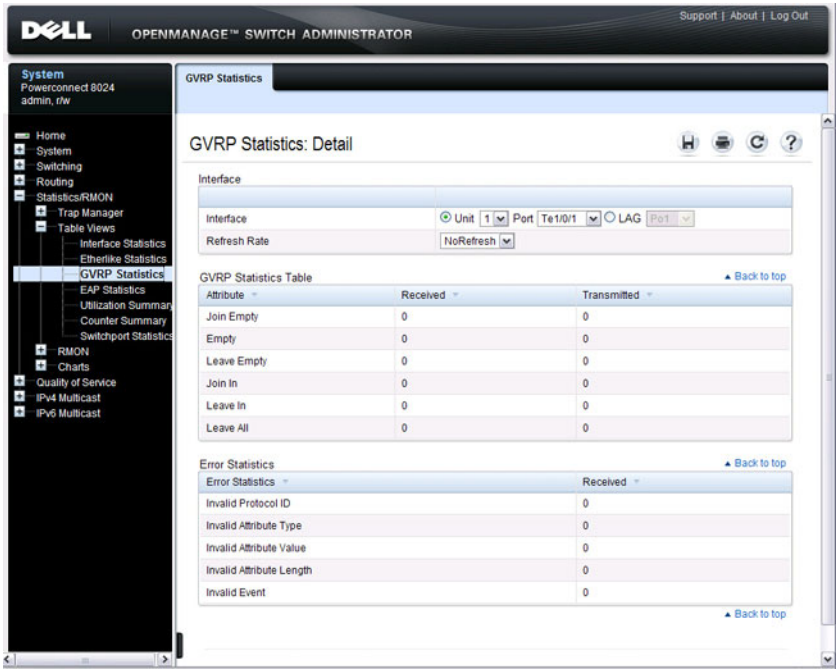


GVRP Statistics

Use the **GVRP Statistics** page to display switch statistics for GVRP.

To display the page, click **Statistics/RMON → Table Views → GVRP Statistics** in the navigation panel.

Figure 16-8. GVRP Statistics

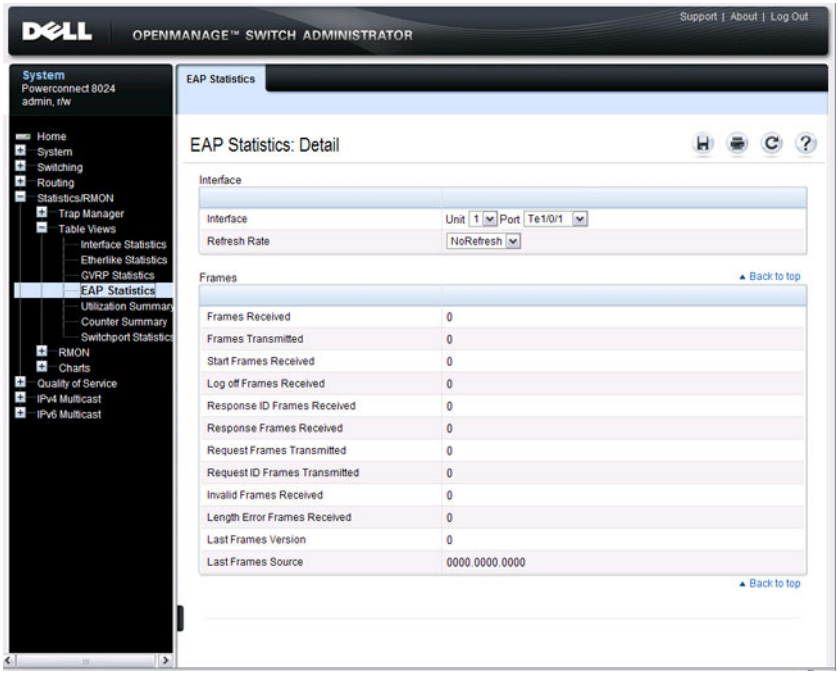


EAP Statistics

Use the **EAP Statistics** page to display information about EAP packets received on a specific port. For more information about EAP, see "Configuring Port and System Security" on page 457.

To display the **EAP Statistics** page, click **Statistics/RMON → Table Views → EAP Statistics** in the navigation panel.

Figure 16-9. EAP Statistics

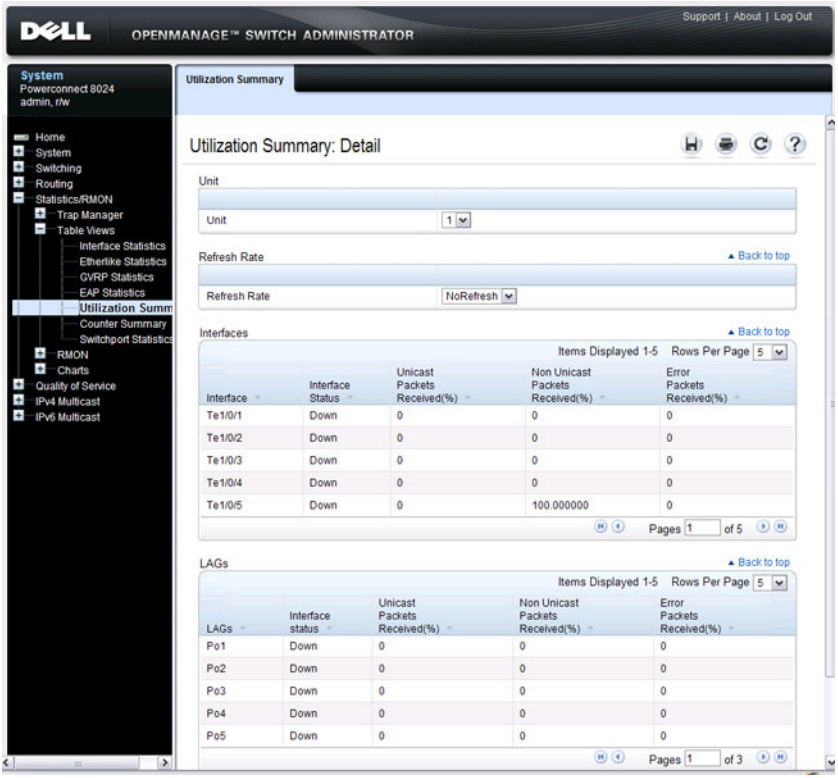


Utilization Summary

Use the **Utilization Summary** page to display interface utilization statistics.

To display the page, click **Statistics/RMON → Table Views → Utilization Summary** in the navigation panel.

Figure 16-10. Utilization Summary

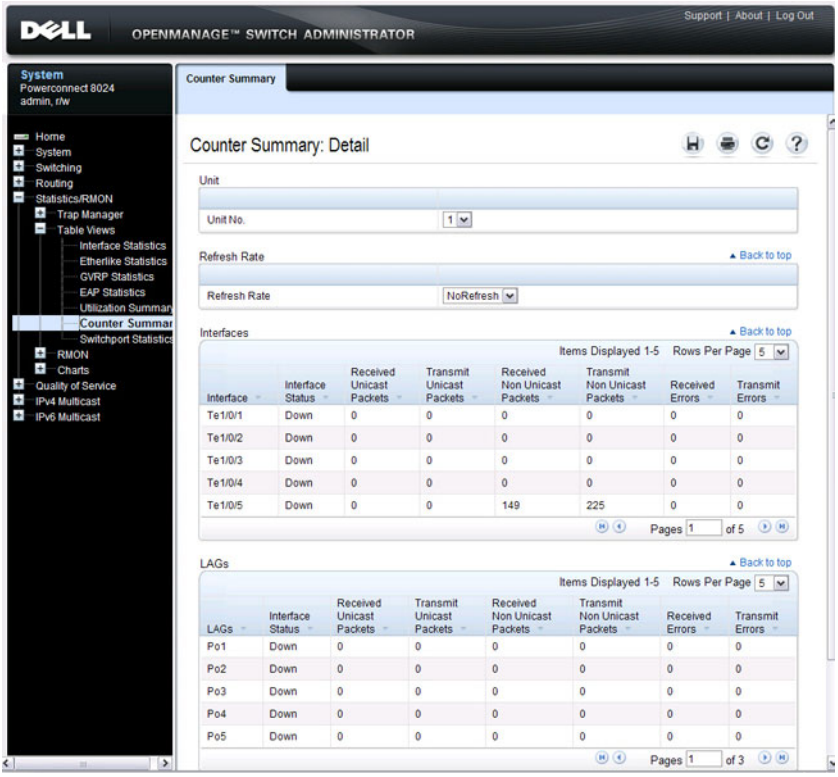


Counter Summary

Use the **Counter Summary** page to display interface utilization statistics in numeric sums as opposed to percentages.

To display the page, click **Statistics/RMON → Table Views → Counter Summary** in the navigation panel.

Figure 16-11. Counter Summary

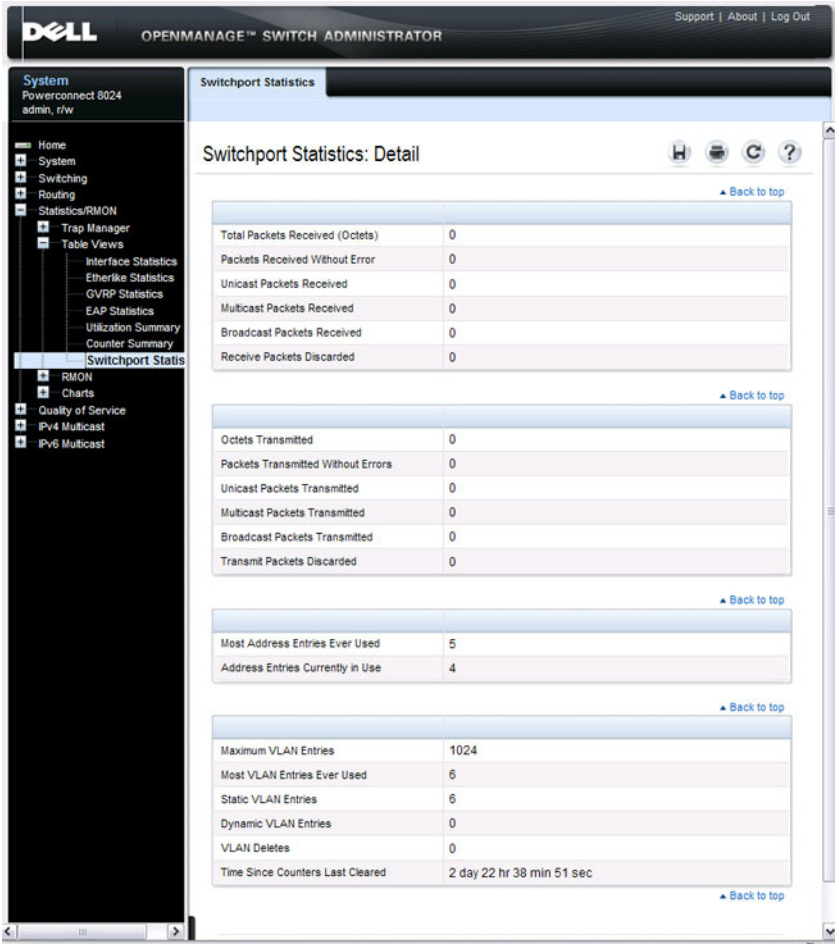


Switchport Statistics

Use the **Switchport Statistics** page to display statistical summary information about switch traffic, address tables, and VLANs.

To display the page, click **Statistics/RMON → Table Views → Switchport Statistics** in the navigation panel.

Figure 16-12. Switchport Statistics

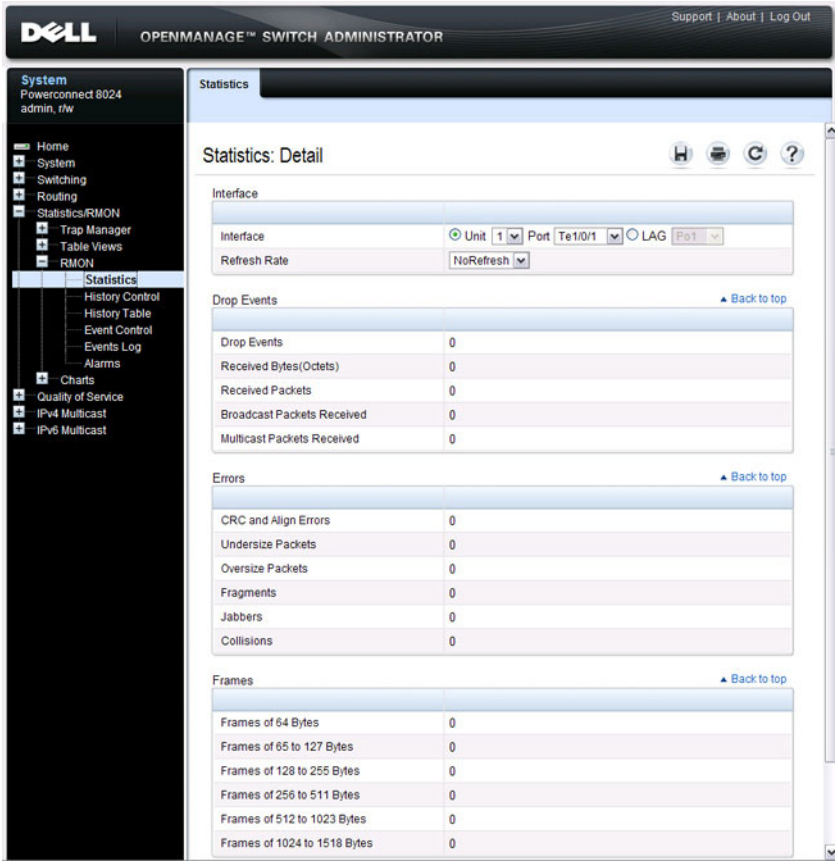


RMON Statistics

Use the **RMON Statistics** page to display details about switch use such as packet processing statistics and errors that have occurred on the switch.

To display the page, click **Statistics/RMON → RMON → Statistics** in the navigation panel.

Figure 16-13. RMON Statistics

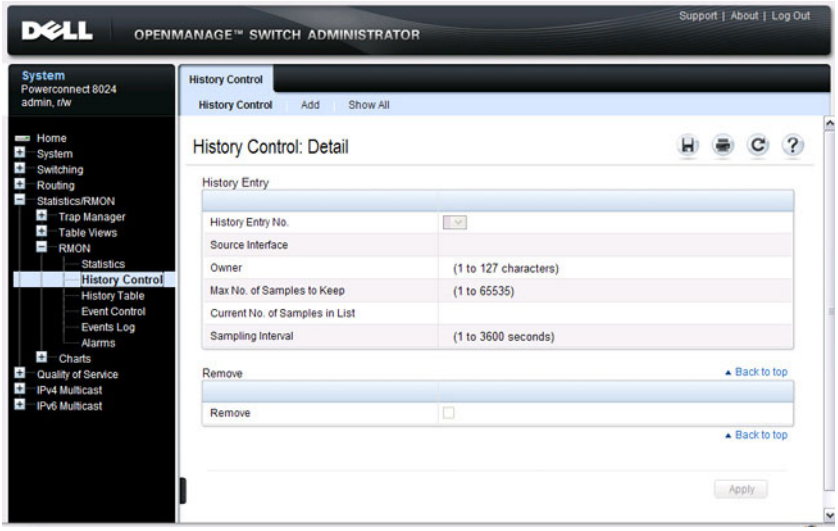


RMON History Control Statistics

Use the **RMON History Control** page to maintain a history of statistics on each port. For each interface (either a physical port or a port-channel), you can define how many buckets exist, and the time interval between each bucket snapshot.

To display the page, click **Statistics/RMON** → **RMON** → **History Control** in the navigation panel.

Figure 16-14. RMON History Control



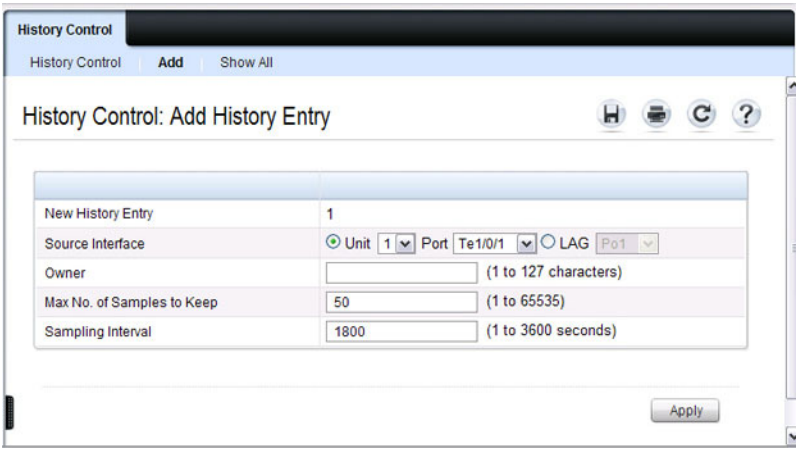
Adding a History Control Entry

To add an entry:

1. Open the **RMON History Control** page.
2. Click **Add**.

   The **Add History Entry** page displays.

Figure 16-15. Add History Entry



3 Select the port or LAG on which you want to maintain a history of statistics.

4 Specify an owner, the number of historical buckets to keep, and the sampling interval.

5 Click **Apply** to add the entry to the **RMON History Control Table**.

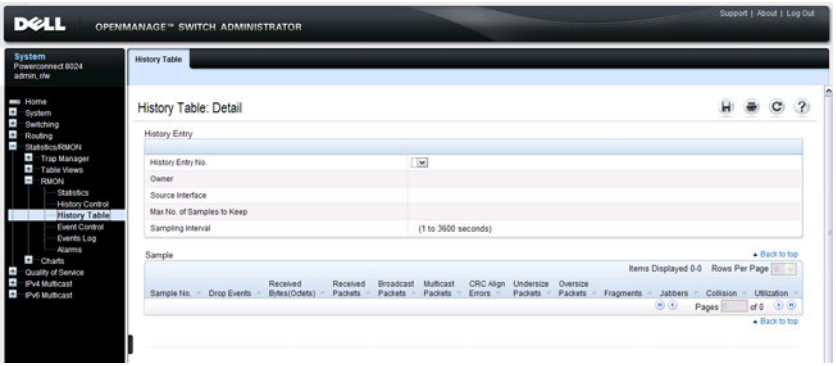
To view configured history entries, click the **Show All** tab. The **RMON History Control Table** displays. From this page, you can remove configured history entries.

RMON History Table

Use the RMON History Table page to display interface-specific statistical network samplings. Each table entry represents all counter values compiled during a single sample.

To display the **RMON History Table** page, click **Statistics/RMON → RMON → History Table** in the navigation panel.

Figure 16-16. RMON History Table

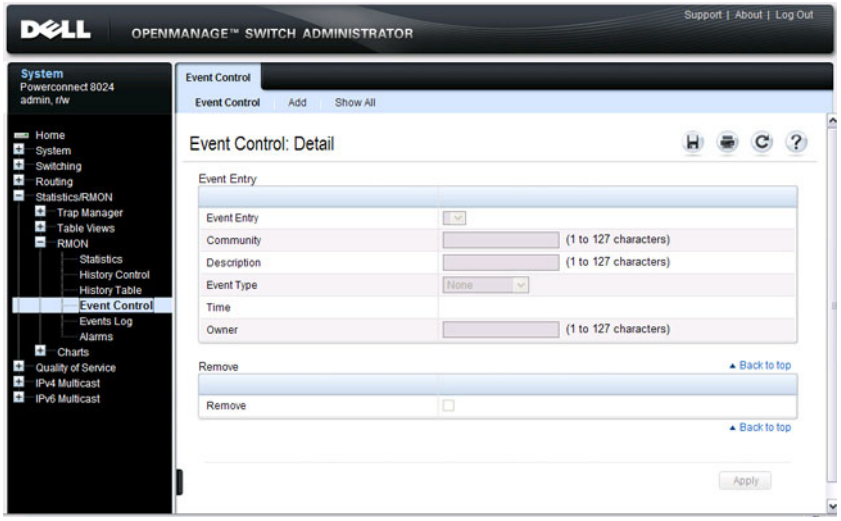


RMON Event Control

Use the **RMON Events Control** page to define RMON events. Events are used by RMON alarms to force some action when a threshold is crossed for a particular RMON counter. The event information can be stored in a log and/or sent as a trap to a trap receiver.

To display the page, click **Statistics/RMON → RMON → Event Control** in the navigation panel.

Figure 16-17. RMON Event Control



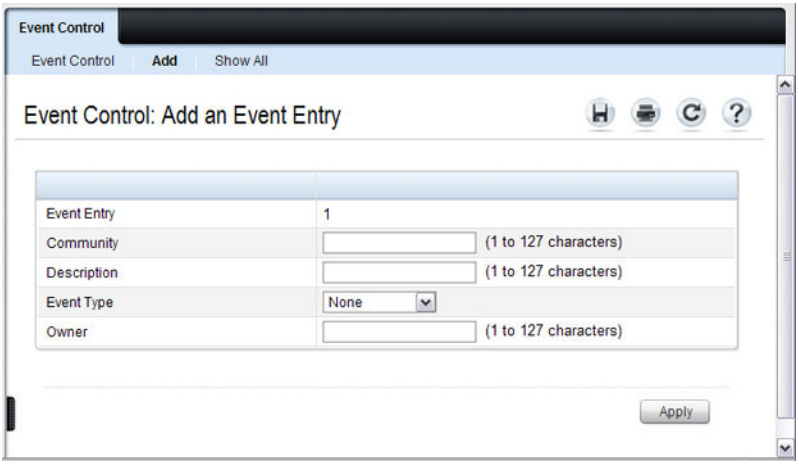
Adding an RMON Event

To add an event:

1. Open the **RMON Event Control** page.
2. Click **Add**.

   The **Add an Event Entry** page displays.

Figure 16-18. Add an Event Entry



3 If the event sends an SNMP trap, specify the SNMP community to receive the trap.

4 Optionally, provide a description of the event and the name of the event owner.

5 Select an event type.

6 Click **Apply**.

The event is added to the **RMON Event Table**, and the device is updated.

Viewing, Modifying, or Removing an RMON Event

To manage an event:

1 Open the **RMON Event Control** page.

2 Click **Show All** to display the **Event Control Table** page.

3 To edit an entry:

- **a** Select the **Edit** check box in for the event entry to change.
- **b** Modify the fields on the page as needed.

4 To remove an entry, select the **Remove** check box in for the event entry to remove.

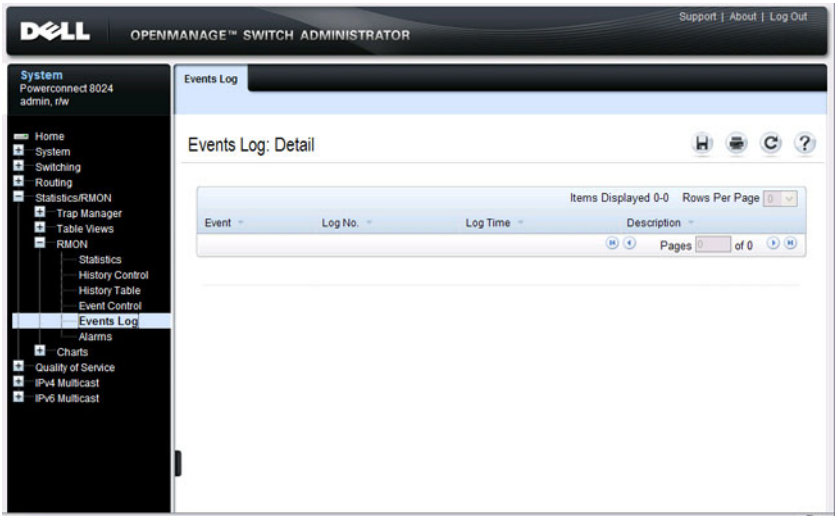
5 Click **Apply**.

RMON Event Log

Use the **RMON Event Log** page to display a list of RMON events.

To display the page, click **Statistics/RMON → RMON → Events Log** in the navigation panel.

Figure 16-19. RMON Event Log

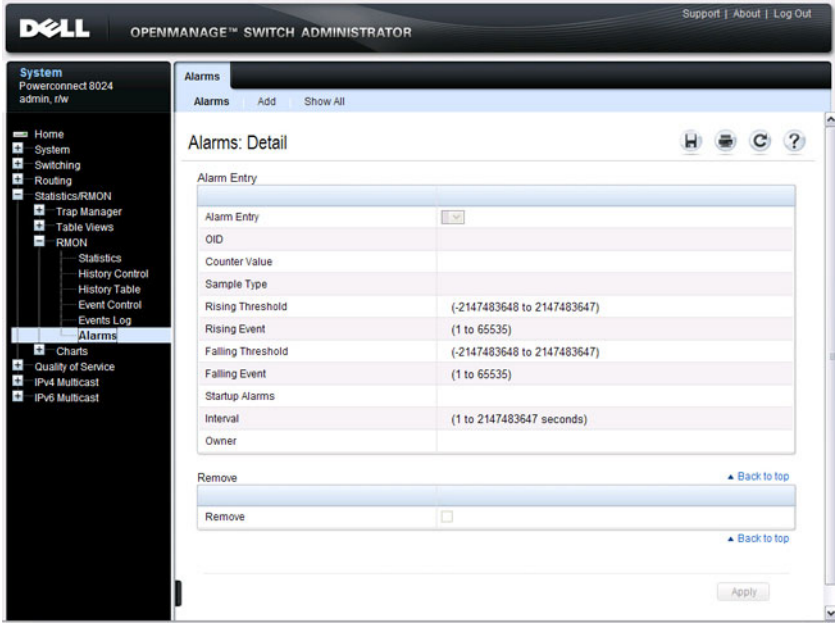


RMON Alarms

Use the **RMON Alarms** page to set network alarms. Alarms occur when certain thresholds are crossed for the configured RMON counters. The alarm triggers an event to occur. The events can be configured as part of the RMON Events group. For more information about events, see "RMON Event Log" on page 370.

To display the page, click **Statistics/RMON → RMON → Alarms** in the navigation panel.

Figure 16-20. RMON Alarms



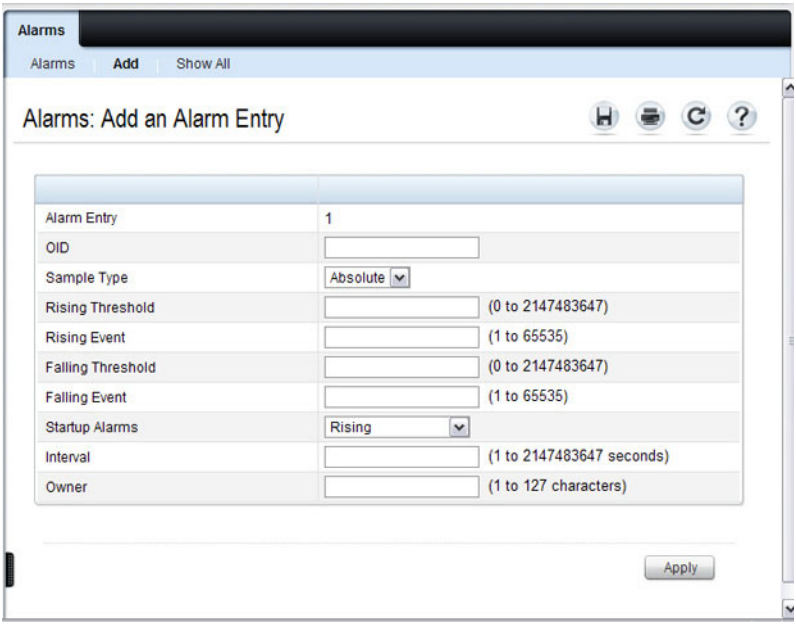
Adding an Alarm Table Entry

To add an alarm:

1. Open the **RMON Alarms** page.
2. Click **Add.**

   The **Add an Alarm Entry** page displays.

Figure 16-21. Add an Alarm Entry



3. Complete the fields on this page as needed. Use the help menu to learn more information about the data required for each field.
4. Click **Apply**.

   The RMON alarm is added, and the device is updated.

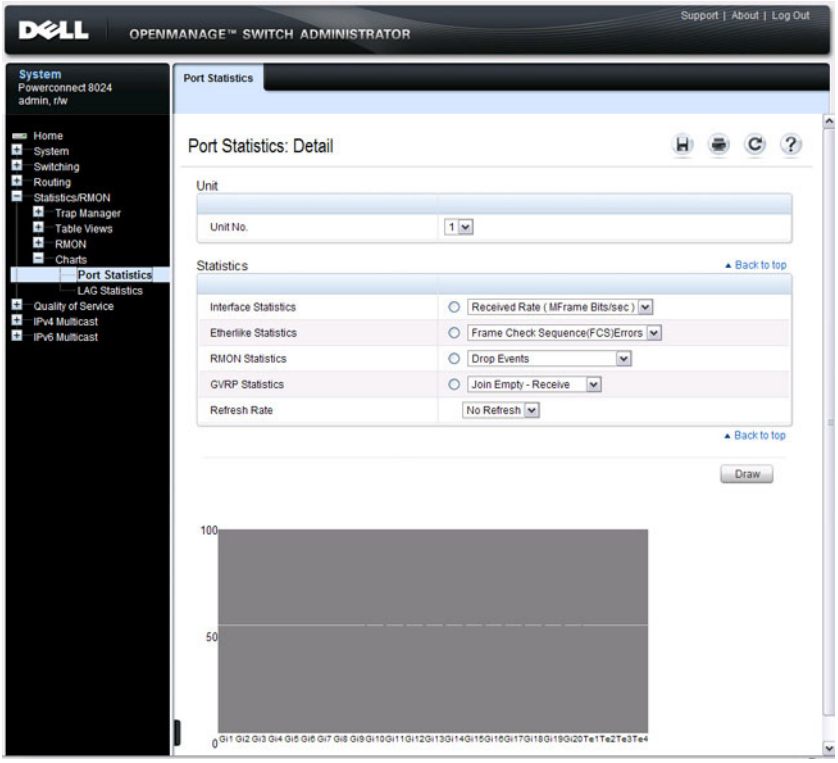
To view configured alarm entries, click the **Show All** tab. The **Alarms Table** displays. From this page, you can remove configured alarms.

Port Statistics

Use the **Port Statistics** page to chart port-related statistics on a graph.

To display the page, click **Statistics/RMON** → **Charts** → **Port Statistics** in the navigation panel.

Figure 16-22. Ports Statistics



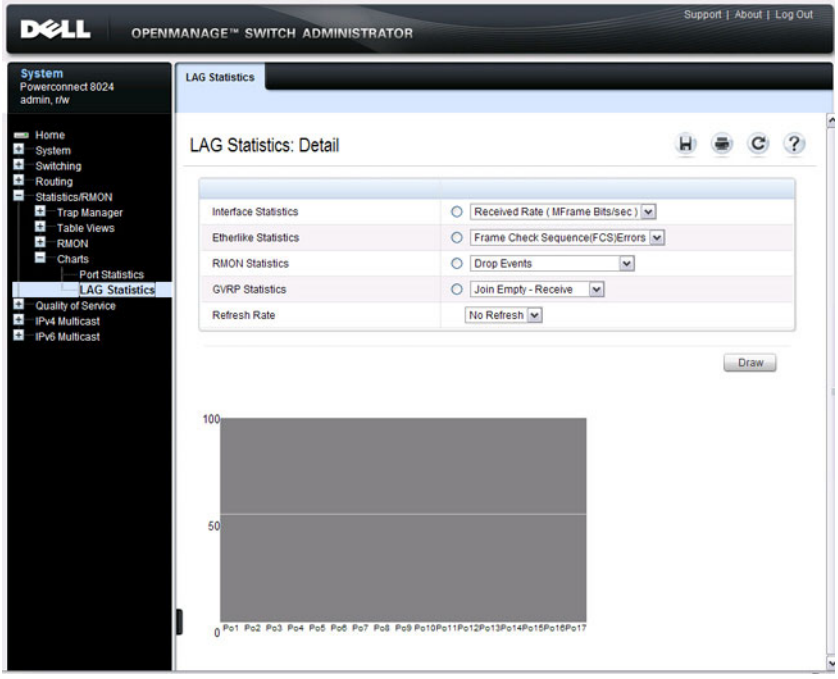
To chart port statistics, select the type of statistics to chart and (if desired) the refresh rate, then click **Draw**.

LAG Statistics

Use the **LAG Statistics** page to chart LAG-related statistics on a graph.

To display the page, click **Statistics/RMON → Charts → LAG Statistics** in the navigation panel.

Figure 16-23. LAG Statistics



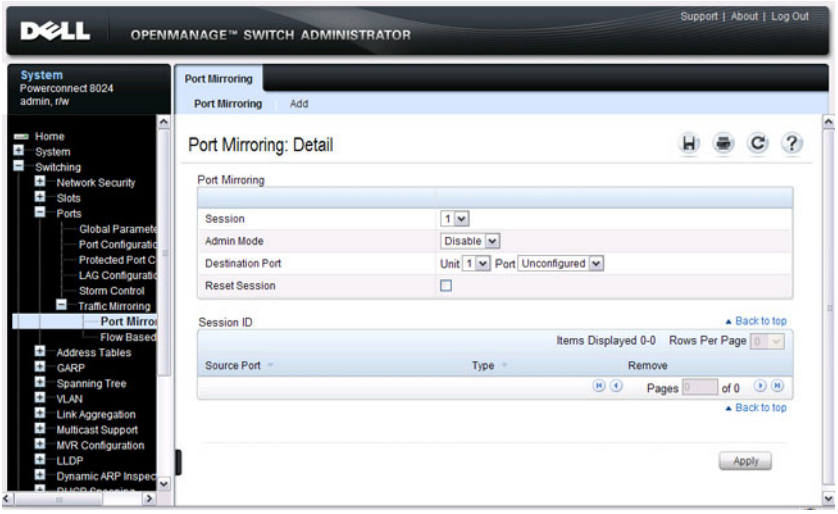
To chart LAG statistics, select the type of statistics to chart and (if desired) the refresh rate, then click **Draw**.

Port Mirroring

Use the **Port Mirroring** page to create a mirroring session in which all traffic that is sent or received (or both) on one or more source ports is mirrored to a destination port.

To display the **Port Mirroring** page, click **Switching** → **Ports** → **Traffic Mirroring** → **Port Mirroring** in the navigation panel.

Figure 16-24. Port Mirroring



Configuring a Port Mirror Session

To configure port mirroring:

1. Open the **Port Mirroring** page.
2. Click **Add**.

   The **Add Source Port** page displays.
3. Select the port to be mirrored.
4. Select the traffic to be mirrored.

Figure 16-25. Add Source Port



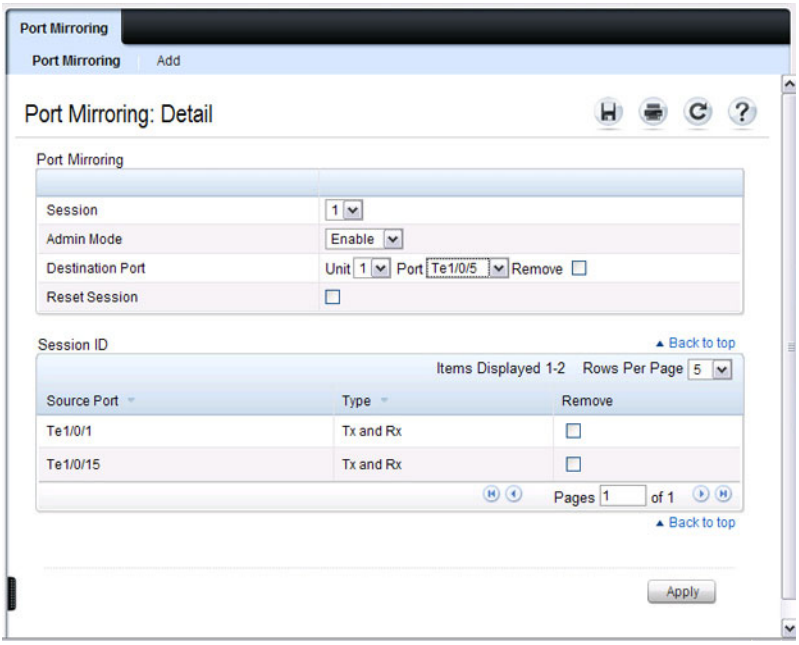
5 Click **Apply**.

6 Repeat the previous steps to add additional source ports.

7 Click **Port Mirroring** to return to the **Port Mirroring** page.

8 Enable the administrative mode and specify the destination port.

Figure 16-26. Configure Additional Port Mirroring Settings



9 Click **Apply**.

Monitoring Switch Traffic (CLI)

This section provides information about the commands you use to manage traffic monitoring features on the switch and to view information about switch traffic. For more information about these commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Configuring sFlow

Beginning in Privileged EXEC mode, use the following commands to configure the sFlow receiver and to configure the sampling and polling on switch interfaces.

Command	Purpose
configure	Enter Global Configuration mode
sflow *rcvr_index* **destination** *ip-address* [*port*]	Configure the address of the sFlow receiver and (optionally) the destination UDP port for sFlow datagrams. • *rcvr_index*—The index of this sFlow receiver (Range: 1–8). • *ip-address*—The sFlow receiver IP address. • *port* —The destination Layer 4 UDP port for sFlow datagrams. (Range: 1–65535).
sflow *rcvr_index* **destination owner** *owner_string* **timeout** *timeout*	Specify the identity string of the receiver and set the receiver timeout value. *timeout*—The number of seconds the configuration will be valid before it is automatically cleared. A value of 0 essentiality means the receiver is not configured.
sflow *rcvr_index* **maxdatagram** *size*	Specify the maximum number of data bytes that can be sent in a single sample datagram. The receiver should also be set this value to avoid fragmentation of the sFlow datagrams. (Range: 200–9116 bytes).

Command	Purpose
sflow *rcvr-index* **polling** *if_type if_number poll-interval*	Enable a new sFlow poller instance on an interface range. • *rcvr-index* — The sFlow Receiver associated with the poller (Range: 1–8). • *if_type if_number* — The list of interfaces to poll. The interface type can be Tengigabitethernet (te), for example te1/0/3-5 enables polling on ports 3, 4, and 5. • *poll-interval* — The sFlow instance polling interval. A value of *n* means once in *n* seconds a counter sample is generated. (Range: 0–86400).
sflow *rcvr-index* **sampling** *if_type if_number sampling-rate* [*size*]	Enable a new sflow sampler instance for the specified interface range. • *rcvr-index* — The sFlow Receiver for this sFlow sampler to which flow samples are to be sent. • *if_type if_number* — The list of interfaces to sample. The interface type can be Tengigabitethernet (te), for example te1/0/3-5 enables polling on ports 3, 4, and 5. • *sampling-rate* — The statistical sampling rate for packet sampling from this source. A sampling rate of 1 counts all packets. A value of *n* means that out of *n* incoming packets, 1 packet will be sampled. (Range: 1024 - 65536). • *size* — The maximum number of bytes that should be copied from the sampler packet (Range: 20 - 256 bytes).
interface *interface*	Enter interface configuration mode for the specified interface. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3** or **te1/0/3**. You can also specify a range of interfaces with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12.
sflow *rcvr-index* **polling** *poll-interval*	Enable a new sFlow poller instance for the interface.
sflow *rcvr-index* **sampling** *sampling-rate* [*size*]	Enable a new sflow sampler instance for the interface.
CTRL + Z	Exit to Privileged Exec mode.

Command	Purpose
show sflow agent	View information about the switch sFlow agent.
show sflow *index* **destination**	View information about a configured sFlow receivers.
show sflow *index* **polling**	View information about the configured sFlow poller instances for the specified receiver.
show sflow *index* **sampling**	View information about the configured sFlow sampler instances for the specified receiver.

Configuring RMON

Beginning in Privileged EXEC mode, use the following commands to configure RMON alarms, collection history, and events. The table also lists the commands you use to view information collected by the RMON probe.

Command	Purpose
configure	Enter Global Configuration mode
rmon event *number* [log] [**trap** *community*] [**description** *string*] [**owner** *string*]	Configure an RMON event. • *number* — The event index. (Range: 1–65535) • **log** — Specify that an entry is made in the log table for each event. • **trap** *community* — If the event is an SNMP trap to be sent, it is sent to the SNMP community specified by this octet string. (Range: 0-127 characters) • **description** *string* — A comment describing this event. (Range 0-127 characters) • **owner** *string* — Enter a name that specifies who configured this event. If unspecified, the name is an empty string.

Command	Purpose
rmon alarm *number variable interval* {**absolute** |**delta**} **rising-threshold** *value* [*event-number*] **rising-threshold** *value* [*event-number*] [**startup** *direction*] [**owner** *string*]	Add an alarm entry • *number* — The alarm index. (Range: 1–65535) • *variable* — A fully qualified SNMP object identifier that resolves to a particular instance of an MIB object. • *interval* — The interval in seconds over which the data is sampled and compared with the rising and falling thresholds. (Range: 1–4294967295) • **rising-threshold** *value* — Rising threshold value. (Range: 0–4294967295) • **rising-threshold** *value* — Falling threshold value. (Range: 0–4294967295) • event-number — The index of the event that is used when a rising or falling threshold is crossed. (Range: 1-65535) • **delta** — The sampling method for the selected variable and calculating the value to be compared against the thresholds. If the method is delta, the selected variable value at the last sample is subtracted from the current value, and the difference compared with the thresholds. • **absolute** — The sampling method for the selected variable and calculating the value to be compared against the thresholds. If the method is absolute, the value of the selected variable is compared directly with the thresholds at the end of the sampling interval. • **startup** *direction* — The type of startup alarm, which can be **rising**, **falling**, or **rising-falling**. • **owner** *string* — Enter a name that specifies who configured this alarm.
interface *interface*	Enter Interface Configuration mode for the specified port or LAG.

Command	Purpose
rmon collection history *index* [**owner** *ownername*] [**buckets** *bucket-number*] [**interval** *seconds*]	Enable an RMON MIB history statistics group on the interface. **NOTE:** You must configure RMON alarms and events before RMON collection history is able to display. • *index* — The requested statistics index group. (Range: 1–65535) • *ownername* — Records the RMON statistics group owner name. If unspecified, the name is an empty string. • *bucket-number* — A value associated with the number of buckets specified for the RMON collection history group of statistics. If unspecified, defaults to 50. (Range: 1 - 65535) • *seconds* — The number of seconds in each polling cycle. If unspecified, defaults to 1800. (Range: 1–3600)
CTRL + Z	Exit to Privileged EXEC mode.
show rmon {alarms |collection history | events | history | log | statistics}	View information collected by the RMON probe.

Viewing Statistics

Use the following commands in Privileged EXEC mode to view statistics about the traffic handled by the switch.

Command	Purpose
show interfaces counters [*if_type if_number* | **port-channel** *interface*]	Display the number of octets and packets handled by all interfaces or the specified interface.
show statistics {switchport | *interface*}	Display detailed statistics for a specific port or LAG, or for the entire switch. The *interface* variable includes the interface type and number.
show gvrp statistics *interface*	Displays GVRP statistics for the specified port or LAG.

Configuring Port Mirroring

Use the following commands in Privileged EXEC mode to configure a port mirroring session.

Command	Purpose
configure	Enter Global Configuration mode
monitor session *session_number* **source interface** {**cpu** | *interface*} [**rx** | **tx**]	Configure a source (monitored) port or CPU interface for a monitor session. • *session_number*—The monitoring session ID, which is always 1. • *interface*—The Ethernet interface to be monitored. • rx | tx — Monitor ingress (rx) or egress (tx) traffic. If you not specify, both ingress and egress traffic is monitored.
monitor session *session_number* **destination interface** *interface*	Configure a destination (probe) port for a monitor session. • *session_number*—The monitoring session ID, which is always 1. • *interface*—The Ethernet interface to which the monitored source traffic is copied.
monitor session *session_number* **mode**	Enable the administrative mode for the configured port mirroring session to start sending the traffic from the source port to the destination (probe) port.
exit	Exit to Privileged EXEC mode.
show monitor session 1	View information about the configured port mirroring session.

Traffic Monitoring Configuration Examples

This section contains the following examples:

- Configuring sFlow
- Configuring RMON

Configuring sFlow

This example shows how to configure the switch so that ports 10-15 and port 23 send sFlow datagrams to an sFlow receiver at the IP address 192.168.20.34. The receiver owner is receiver1, and the timeout is 100000 seconds. A counter sample is generated on the ports every 60 seconds (polling interval), and 1 out of every 8192 packets is sampled. Note that sFlow monitoring is not enabled until a receiver owner string is configured.

To configure the switch:

1. Configure information about the sFlow receiver.

```
console#configure
console(config)#sflow 1 destination 192.168.30.34
console(config)#sflow 1 destination owner
receiver1 timeout 100000
```

2. Configure the polling and sampling information for tengigabit Ethernet ports 10-20.

```
console(config)#sflow 1 polling te1/0/10-15 60
console(config)#sflow 1 sampling te1/0/10-15 8192
```

3. Configure the polling and sampling information for tengigabit Ethernet port 23.

```
console(config)#interface te1/0/23
console(config-if-Te1/0/23)#sflow 1 polling 60
console(config-if-Te1/0/23)#sflow 1 sampling 8192
```

4. Verify the configured information.

```
console#show sflow 1 destination

Receiver Index..................... 1
Owner String....................... receiver1
Time out........................... 99994
IP Address:........................ 192.168.30.34
```

```
Address Type..................... 1
Port............................. 6343
Datagram Version................. 5
Maximum Datagram Size............ 1400

console#show sflow 1 polling

 Poller          Receiver  Poller
 Data Source     Index     Interval
 -----------     -------   -------
te1/0/10          1         60
te1/0/11          1         60
te1/0/12          1         60
te1/0/13          1         60
te1/0/14          1         60
te1/0/15          1         60
te1/0/23          1         60

console#show sflow 1 sampling

 Sampler       Receiver  Packet           Max Header
 Data Source   Index     Sampling Rate    Size
 -----------   -------   -------------    ----------
te1/0/10        1         8192            128
te1/0/11        1         8192            128
te1/0/12        1         8192            128
te1/0/13        1         8192            128
te1/0/14        1         8192            128
te1/0/15        1         8192            128
te1/0/23        1         8192            128
```

Configuring RMON

This example generates a trap and creates a log entry when the number of inbound packets are undeliverable due to errors increases by 20 or more.

First, an RMON event is created. Then, the alarm is created. The event (event 1) generates a trap and creates a log entry. The alarm is configured for the MIB object ifInErrors (OID: 1.3.6.1.2.1.2.2.1.14.1). The OID is the variable. The alarm checks the variable every 30 seconds to compare the MIB counter to the configured rising and falling thresholds. If the rise is equal to or greater than 20, event 1 goes into effect.

To configure the switch:

1. Create the event. The trap is sent to the private SNMP community.

```
console#configure
console(config)#rmon event 1 description
"emergency event" log trap private
```

2. Create the alarm.

```
console(config)#rmon alarm 1
1.3.6.1.2.1.2.2.1.14.1 30 delta rising-threshold
20 1 falling-threshold 1
```

3. Verify the configuration.

```
console#show rmon events

Index  Description   Type     Community  Owner Last time sent
-------------------------------------------------------------
1      emergency    log-trap  private          0 days 0h:0m:0s
       event

console#show rmon alarms

Index    OID                          Owner
-----------------------------------------------
1       1.3.6.1.2.1.2.2.1.14.1
```

17

Configuring iSCSI Optimization

This chapter describes how to configure Internet Small Computer System Interface (iSCSI) optimization, which enables special quality of service (QoS) treatment for iSCSI traffic.

The topics covered in this chapter include:

- iSCSI Optimization Overview
- Default iSCSI Optimization Values
- Configuring iSCSI Optimization (Web)
- Configuring iSCSI Optimization (CLI)
- iSCSI Optimization Configuration Examples

iSCSI Optimization Overview

iSCSI optimization provides a means of monitoring iSCSI sessions and iSCSI traffic on the switch. This is accomplished by monitoring, or "snooping," traffic to detect packets used by iSCSI stations to establish iSCSI sessions and connections. Data from these exchanges may optionally be used to create classification rules to assign traffic between the stations to a configured traffic class. The traffic classification affects how the packets in the flow are queued and scheduled for egress on the destination port.

What Does iSCSI Optimization Do?

In networks containing iSCSI initiators and targets, iSCSI Optimization helps to monitor iSCSI sessions or give iSCSI traffic preferential QoS treatment. Dynamically-generated classifier rules are used to direct the iSCSI data traffic to queues that can be given the desired preference characteristics over other data traveling through the switch. This may help to avoid session interruptions during times of congestion that would otherwise cause iSCSI packets to be dropped. However, in systems where a large proportion of traffic is iSCSI, it may also interfere with other network control-plane traffic, such as ARP or LACP.

The preferential treatment of iSCSI traffic needs to be balanced against the needs of other critical data in the network.

How Does the Switch Detect iSCSI Traffic Flows?

The switch detects iSCSI session establishment and termination packets by installing classifier rules that trap iSCSI protocol packets to the CPU for examination. Devices that initiate iSCSI sessions generally use well-known TCP ports 3260 or 860 to contact targets. When iSCSI optimization is enabled, by default the switch identifies IP packets to or from these ports as iSCSI session traffic. You can configure the switch to monitor traffic for additional port numbers or port number-target IP address combinations, and you can remove the well-known port numbers from monitoring. You can also associate a target name with a configured target TCP port entry.

How Is Quality of Service Applied to iSCSI Traffic Flows?

The iSCSI CoS mode is configurable and controls whether CoS queue assignment and/or packet marking is performed on iSCSI traffic. When the iSCSI CoS mode is enabled, the CoS policy is applied to packets in detected iSCSI sessions. When the iSCSI CoS mode is disabled, iSCSI sessions and connections are detected and shown in the status tables, but no CoS policy is applied to packets.

You can configure whether the iSCSI optimization feature uses the VLAN priority or IP DSCP mapping to determine the traffic class queue. By default, iSCSI flows are assigned to the highest VLAN priority tag or DSCP value mapped to the highest queue not used for stack management or voice VLAN. Use the **classofservice dot1p-mapping** command or the **Quality of Service →**

Class of Service → Mapping Table Configuration page to configure the relevant Class of Service parameters for the queue in order to complete the setting.

You can configure whether iSCSI frames are remarked to contain the configured VLAN priority tag or IP DSCP when forwarded through the switch.

How Does iSCSI Optimization Use ACLs?

iSCSI Optimization borrows ACL lists from the global system pool. ACL lists allocated by iSCSI Optimization reduce the total number of ACLs available for use by the network operator. Enabling iSCSI Optimization uses one ACL list to monitor for iSCSI sessions. Each monitored iSCSI session utilizes two rules from additional ACL lists up to a maximum of two ACL lists. This means that the maximum number of ACL lists allocated by iSCSI is three.

What Information Does the Switch Track in iSCSI Traffic Flows?

Packets are examined to find the following data, which is used in tracking the session and creating the classifier entries that enable QoS treatment:

- Initiator's IP Address
- Target's IP Address
- ISID (Initiator defined session identifier)
- Initiator's IQN (iSCSI Qualified Name)
- Target's IQN
- Initiator's TCP Port
- Target's TCP Port

If no iSCSI traffic is detected for a session for a configurable aging period, the session data is cleared.

How Does iSCSI Optimization Interact With Dell EqualLogic Arrays?

The iSCSI feature includes auto-provisioning support with the ability to detect directly connected Dell EqualLogic SAN storage arrays and automatically reconfigure the switch to enhance storage traffic flows.

The PowerConnect 8000-series and 8100-series switches uses LLDP, a vendor-neutral protocol, to discover Dell EqualLogic devices on the network. LLDP is enabled by default. For more information about LLDP, see "Discovering Network Devices" on page 637.

When the switch detects a Dell EqualLogic array, the following actions occur:

- An MTU of 9216 is enabled on all ports and port-channels, if it is not already enabled.
- Spanning-Tree portfast is enabled on the interface identified by LLDP.
- Unicast storm control is disabled on the interface identified by LLDP.

If the iSCSI CoS policy feature is enabled on the switch and an EqualLogic array is detected, the switch applies additional iSCSI CoS policies to the EqualLogic inter-array traffic on TCP ports 9876 and 25555. If the iSCSI CoS policy is disabled and EqualLogic arrays are present, the additional CoS policy is removed globally.

What Occurs When iSCSI Optimization Is Enabled or Disabled?

When iSCSI is enabled on the switch, the following actions occur:

- Flow control is globally enabled, if it is not already enabled.
- iSCSI session snooping is enabled
- iSCSI LLDP monitoring starts to automatically detect Dell EqualLogic arrays.

If the iSCSI feature is disabled on the switch, iSCSI resources are released and the detection of Dell EqualLogic arrays by using LLDP is disabled. Disabling iSCSI does not remove the MTU, flow control, portfast or storm control configuration applied as a result of enabling iSCSI. iSCSI Optimization is enabled by default.

How Does iSCSI Optimization Interact with DCBx?

The Data Center Bridging Exchange (DCBx) component supports the reception, decoding, and transmission of the Application Priority TLV. In general, if the Application Priority TLV has been received from the configuration source, it will be transmitted to the other auto configuration ports. The DCBx component contains a control to generate the Application Priority TLV for iSCSI if it is not already present in the DCBX information. DCBx generates an Application Priority TLV whenever the following conditions are met:

- An EqualLogic array has been detected on the port
- iSCSI CoS is enabled using a VPT value

The generated Application Priority TLV will contain the following values (in addition to any other information contained in the TLV):

- AE Selector=14
- AE Protocol=3260
- AE Priority=priority configured for iSCSI PFC by the **iscsi cos vpt** command (default priority is 4)

The existing application priority entries being transmitted, if any, will not be disturbed.

NOTE: If it is desired to utilize DCBX to configure lossless transport of iSCSI using PFC, the operator MUST configure a VLAN end-to-end in order to transport the VLAN priority tag and ensure proper COS treatment on every network enabled device, including CNAs and the EQL arrays.

How Does iSCSI Optimization Interact with Dell Compellent Arrays?

Dell PowerConnect switches support a macro that may be used to configure a port connected to a Dell Compellent storage array. The name of the macro is `profile-compellent-nas`. The macro takes a single argument: the interface identifier to which the Dell Compellent array is connected. The macro disables unicast storm control and sets the spanning tree configuration on the port to portfast. For an example of how to execute the macro, see "Configuring iSCSI Optimization Between Servers and a Disk Array" on page 399.

Default iSCSI Optimization Values

Table 17-1 shows the default values for the iSCSI optimization feature.

Table 17-1. iSCSI Optimization Defaults

Parameter	Default Value
iSCSI Optimization Global Status	Enabled
iSCSI CoS mode	Disabled
Classification	iSCSI packets are classified by VLAN instead of by DSCP values.
VLAN Priority tag	iSCSI flows are assigned by default the highest 802.1p VLAN priority tag mapped to the highest queue not used for the voice VLAN.
DSCP	When DSCP is selected as the classification, iSCSI flows are assigned by default the highest DSCP tag mapped to the highest queue not used for the voice VLAN.
Remark	Non configured.
iSCSI Session Aging Time	10 minutes
iSCSI Optimization Target Ports	iSCSI well-known ports 3260 and 860 are configured as default (with no IP address or name) but can be removed as any other configured target.

Configuring iSCSI Optimization (Web)

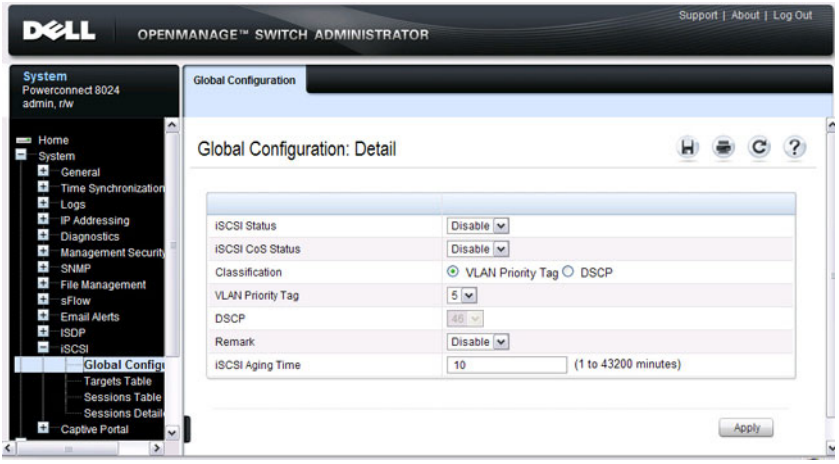
This section provides information about the OpenManage Switch Administrator pages to use to the iSCSI features on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click at the top of the page.

iSCSI Global Configuration

Use the **Global Configuration** page to allow the switch to snoop for iSCSI sessions/connections and to configure QoS treatment for packets where the iSCSI protocol is detected.

To access the iSCSI **Global Configuration** page, click **System** → **iSCSI** → **Global Configuration** in the navigation panel.

Figure 17-1. iSCSI Global Configuration

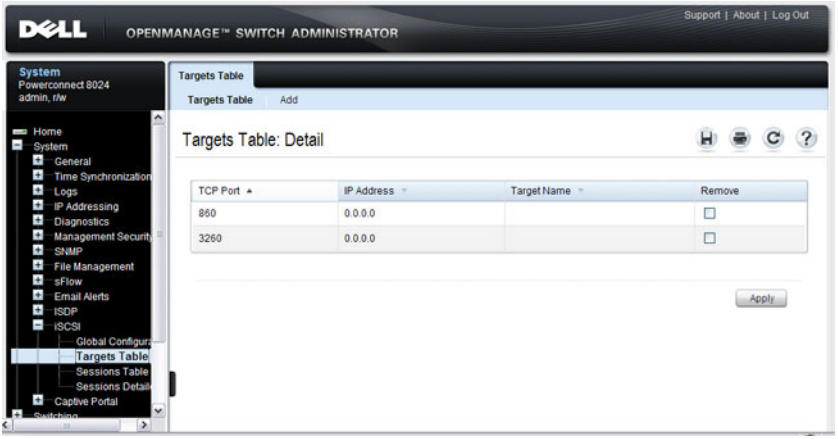


iSCSI Targets Table

Use the **Targets Table** page to view and configure iSCSI targets on the switch.

To access the **Targets Table** page, click **System** → **iSCSI** → **Targets** in the navigation panel.

Figure 17-2. iSCSI Targets Table



To add an iSCSI Target, click **Add** at the top of the page and configure the relevant information about the iSCSI target.

Figure 17-3. Add iSCSI Targets

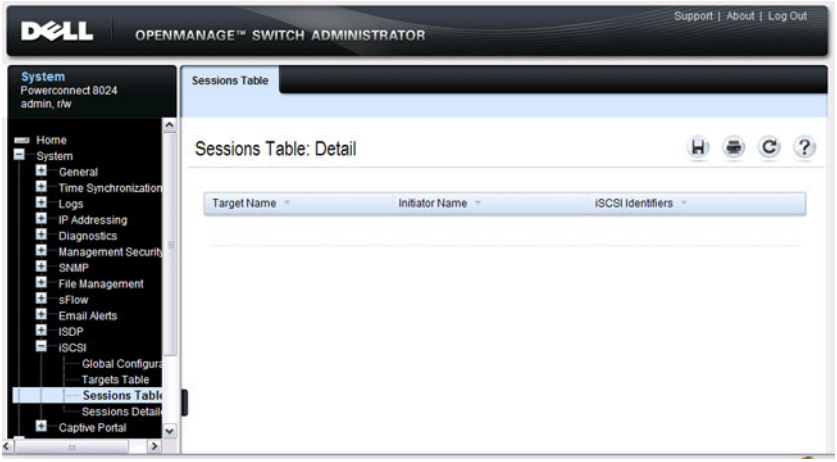


iSCSI Sessions Table

Use the **Sessions Table** page to view summary information about the iSCSI sessions that the switch has discovered. An iSCSI session occurs when an iSCSI initiator and iSCSI target communicate over one or more TCP connections. The maximum number of iSCSI sessions is 192.

To access the **Sessions Table** page, click **System** → **iSCSI** → **Sessions Table** in the navigation panel.

Figure 17-4. iSCSI Sessions Table

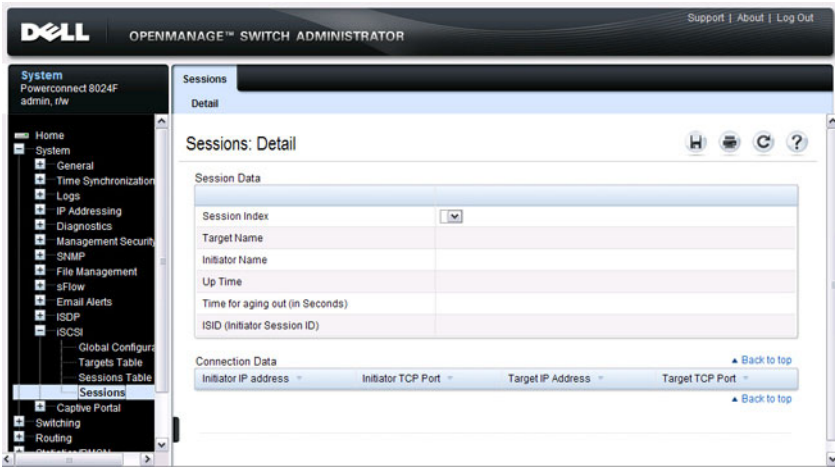


iSCSI Sessions Detailed

Use the **Sessions Detailed** page to view detailed information about an iSCSI sessions that the switch has discovered.

To access the **Sessions Detailed** page, click **System → iSCSI → Sessions Detailed** in the navigation panel.

Figure 17-5. iSCSI Sessions Detail



Configuring iSCSI Optimization (CLI)

This section provides information about the commands you use to configure iSCSI settings on the switch. For more information about the commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Command	Purpose
configure	Enter Global Configuration mode.
iscsi enable	Globally enable iSCSI optimization.
iscsi target port *tcp-port-1* [*tcp-port-2...tcp-port-16*] [**address** *ip-address*] [**name** *targetname*]	Configure an iSCSI target port and, optionally, address and name. • *tcp-port-n*—TCP port number or list of TCP port numbers on which the iSCSI target listens to requests. Up to 16 TCP ports can be defined in the system in one command or by using multiple commands. • *ip-address*—IP address of the iSCSI target. When the no form of this command is used, and the tcp port to be deleted is one bound to a specific IP address, the address field must be present. • *targetname*—iSCSI name of the iSCSI target. The name can be statically configured; however, it can be obtained from iSNS or from sendTargets response. The initiator must present both its iSCSI Initiator Name and the iSCSI Target Name to which it wishes to connect in the first login request of a new session or connection.

Command	Purpose
iscsi cos {**enable** | **disable** | **vtp** *vtp* | **dscp** *dscp* [**remark**]	Set the quality of service profile that will be applied to iSCSI flows. • **enable**—Enables application of preferential QoS treatment to iSCSI frames • **disable**—Disables application of preferential QoS treatment to iSCSI frames. • **vpt/dscp**—The VLAN Priority Tag or DSCP value to assign received iSCSI session packets. • **remark**—Mark the iSCSI frames with the configured DSCP value when egressing the switch.
iscsi aging time *time*	Set aging time (range: 1–43,200 seconds) for iSCSI sessions.
exit	Exit to Privilege Exec mode.
show iscsi	Display iSCSI settings.
show iscsi sessions	Display iSCSI session information.

iSCSI Optimization Configuration Examples

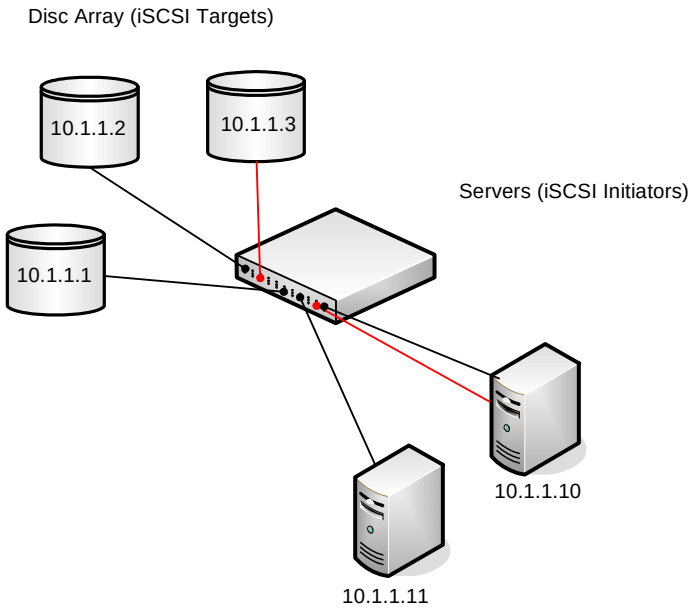
iSCSI optimization is enabled by default with the appropriate settings to operate properly is almost all configurations. However, you find it necessary to alter those settings, the following procedure illustrates the configuration steps required.

Configuring iSCSI Optimization Between Servers and a Disk Array

Figure 17-6 illustrates a PowerConnect 8000-series and 8100-series switches connecting two servers (iSCSI initiators) to a disk array (iSCSI targets).

An iSCSI application running on the switch has installed priority filters to ensure that iSCSI traffic that is part of these two sessions receives priority treatment when forwarded in hardware.

Figure 17-6. iSCSI Optimization



The following commands show how to configure the iSCSI example depicted in Figure 17-6.

1. Enable iSCSI optimization on the switch if it has been previously disabled (iSCSI optimization is enabled by default).

```
console#config
console(config)#iscsi enable
```

2. Configure the switch to associate the DSCP priority 45 (and the queue that is mapped to it) with detected iSCSI session traffic. The **remark** keyword indicates that the switch should add this priority marking on packets as it forwards them.

```
console(config)#iscsi cos dscp 45 remark
console(config)#exit
```

   The default target port and IP address criteria is used to determine which packets are snooped for iSCSI session data (ports 860 and 3260; any IP address).

3. If the array is a Compellent storage array, execute the Compellent macro on the ports attached to the array:

```
console#config
console(config)#macro global apply profile-
compellent-nas $interface_name te1/0/21
console(config)#macro global apply profile-
compellent-nas $interface_name te1/0/22
console(config)#macro global apply profile-
compellent-nas $interface_name te1/0/23
```

18

Configuring a Captive Portal

This chapter describes how to configure the Captive Portal feature.

The topics covered in this chapter include:

- Captive Portal Overview
- Default Captive Portal Behavior and Settings
- Configuring the Captive Portal (Web)
- Configuring a Captive Portal (CLI)
- Captive Portal Configuration Example

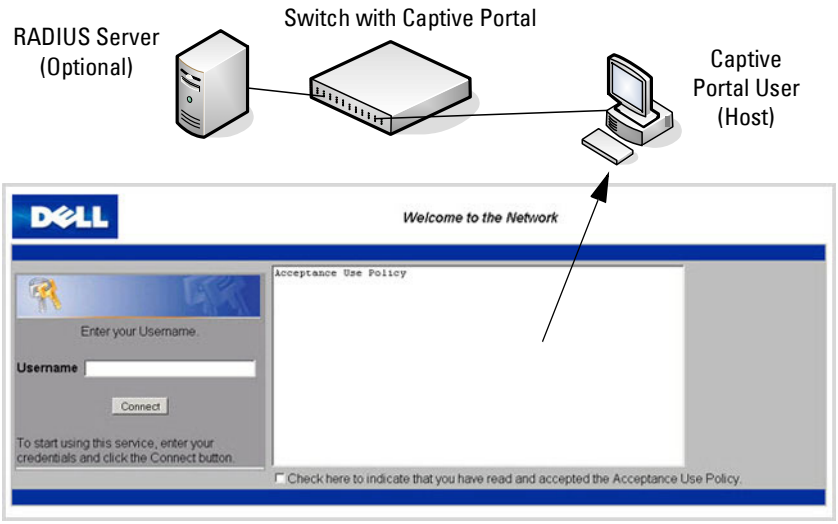
Captive Portal Overview

A Captive Portal helps manage or restrict network access. Captive Portals are often used in locations that provide wired Internet access to customers, such as business centers and hotels. For example, a hotel might provide an Ethernet port in each room so that guests can connect to the Internet during their stay. The hotel might charge for Internet use, or the hotel might allow guests to connect only after they indicate that they have read and agree to the acceptable use policy.

What Does a Captive Portal Do?

The Captive Portal feature allows you to require a user to enter login information on a custom Web page before gaining access to the network. When the user connects to the port and opens a browser, the user is presented with a welcome screen. To gain network access, the user must enter a username (for guest access) or a username and password (for authenticated access) and accept the terms of use. You can also configure the Captive Portal feature to redirect the user to another web page after successful authentication, for example your company home page.

Figure 18-1. Connecting to the Captive Portal



Default Captive Portal Welcome Screen (Displays in Captive Portal User's Browser)

The Captive Portal feature blocks hosts connected to the switch from accessing the network until user verification has been established. You can configure Captive Portal verification to allow access for both guest and authenticated users. Authenticated users must be validated against a database of authorized Captive Portal users before access is granted. The database can be stored locally on the switch or on a RADIUS server.

Is the Captive Portal Feature Dependent on Any Other Feature?

If you require RADIUS authentication, you must configure the RADIUS server information on the switch (see "Using RADIUS Servers to Control Management Access" on page 192). You must also configure the RADIUS attributes for Captive Portal users on the RADIUS server. For information about the RADIUS attributes to configure, see Table 18-2.

For a list of RADIUS attributes that the switch supports, see "Which RADIUS Attributes Does the Switch Support?" on page 194.

You can configure the switch to send SNMP trap messages to any enabled SNMP Trap Receivers for several Captive Portal events, such as when a Captive Portal user has an authentication failure or when a Captive Portal user successfully connects to the network. If you enable the traps, the switch also writes a message to the trap log when the event occurs. To enable the Captive Portal traps, see "Configuring SNMP Notifications (Traps and Informs)" on page 299.

What Factors Should Be Considered When Designing and Configuring a Captive Portal?

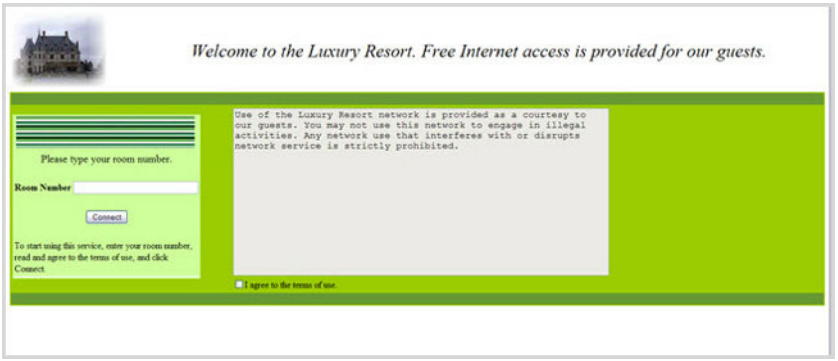
Before enabling the Captive Portal feature, decide what type (or types) of authentication to require. Since the PowerConnect 8000-series and 8100-series switches support up to 10 different Captive Portal instances, you can configure one Captive Portal that requires a username and password and another that only requires the username. For each Captive Portal, you can customize the welcome screen, including the colors and logo.

If you require authentication, consider the number of users that must exist in the user database. The local user database supports up to 128 users. If you need to support more than 128 authenticated users, you must use a remote RADIUS server for authentication.

You can specify whether the captive portal uses HTTP or HTTPS as the protocol during the user verification process. HTTP does not use encryption during verification, and HTTPS uses the Secure Sockets Layer (SSL), which requires a certificate to provide encryption. The certificate is presented to the user at connection time.

The initial Web page that a user sees when he or she connects to the Captive Portal can be customized. You can change the logo, color schemes, welcome messages, and all text on the page, including the field and button labels. The welcome page the user sees after a successful verification or authentication can also be customized.

Figure 18-2. Customized Captive Portal Welcome Screen



How Does Captive Portal Work?

When a port is enabled for Captive Portal, all the traffic coming onto the port from the unverified clients are dropped except for the ARP, DHCP, DNS and NETBIOS packets. These packets are allowed to be forwarded by the switch so that the unverified clients can get an IP address and are able to resolve the hostname or domain names. Data traffic from verified clients goes through as expected. If an unverified client opens a web browser and tries to connect to the network, the Captive Portal redirects all the HTTP/HTTPS traffic from the unverified clients to the authenticating server on the switch. A Captive Portal web page is sent back to the unverified client. If the verification mode for the Captive Portal associated with the port is Guest, the client can be verified without providing authentication information. If the verification mode is Local or RADIUS, the client must provide credentials that are compared against the information in the Local or RADIUS client database. After the user successfully provides the required information, the Captive Portal feature grants access to the network.

What Captive Portal Pages Can Be Customized?

You can customize the following three Captive Portal pages:

- Authentication Page —This page displays when a client attempts to connect to the network. You can customize the images, text, and colors that display on this page.
- Logout Page — If the user logout mode is enabled, this page displays in a pop-up window after the user successfully authenticates. This window contains the logout button.
- Logout Success Page — If the user logout mode is enabled, this page displays after a user clicks the logout button and successfully deauthenticates.

Understanding User Logout Mode

The User Logout Mode feature allows a user who successfully authenticates to the network through the captive portal to explicitly deauthenticate from the network. When User Logout Mode is disabled or the user does not specifically request logout, the connection status will remain authenticated until the Captive Portal deauthenticates the user based on the configured session timeout value. In order for the user logout feature to function properly, the client browser must have JavaScript enabled an must allow popup windows.

Localizing Captive Portal Pages

The Captive Portal localization feature allows you to create up to five language-specific web pages for each captive portal as long as all pages use the same verification type; either guest or authorized user web pages. This allows you to create pages in a variety of languages to accommodate a diverse group of users.

To customize the pages that the user sees, click the language tab. By default, the English tab is available. The settings for the **Authentication Page** display.

Default Captive Portal Behavior and Settings

Captive Portal is disabled by default. If you enable Captive Portal, no interfaces are associated with the default Captive Portal. After you associate an interface with the Captive Portal and globally enable the Captive Portal feature, a user who connects to the switch through that interface is presented with the Captive Portal Welcome screen shown in Figure 18-3.

Figure 18-3. Default Captive Portal Welcome Screen



The user types a name in the Username field, selects the Acceptance Use Policy check box, and clicks **Connect** to gain network access. By default, the user does not need to be defined in a database or enter a password to access the network because the default verification mode is Guest. Note that duplicate Username entries can exist in this mode because the client IP and MAC addresses are obtained for identification.

Table 18-1 shows the default values for the Captive Portal feature.

Table 18-1. Default Captive Portal Values

Feature	Value
Global Captive Portal Operational Status	Disabled
Additional HTTP or HTTPS Ports	Disabled Captive Portal can be configured to use an additional HTTP and/or HTTPS port (in support of Proxy networks).

Table 18-1. Default Captive Portal Values

Feature	Value
Authentication Timeout	300 seconds
Configured Captive Portals	1
Captive Portal Name	Default
Protocol Mode	HTTP
Verification Mode	Guest
URL Redirect Mode	Off
User Group	1-Default
Session Timeout	86400 seconds
Local Users	None configured
Interface associations	None
Interface status	Not blocked If the Captive Portal is blocked, users cannot gain access to the network through the Captive Portal. Use this function to temporarily protect the network during unexpected events, such as denial of service attacks.
Supported Captive Portal users	1024
Supported local users	128
Supported Captive Portals	10

Configuring the Captive Portal (Web)

This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring Captive Portal settings on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

Captive Portal Global Configuration

Use the **Captive Portal Global Configuration** page to control the administrative state of the Captive Portal feature and configure global settings that affect all captive portals configured on the switch.

To display the **Captive Portal Global Configuration** page, click **System → Captive Portal → Global Configuration.**

Figure 18-4. Captive Portal Global Configuration



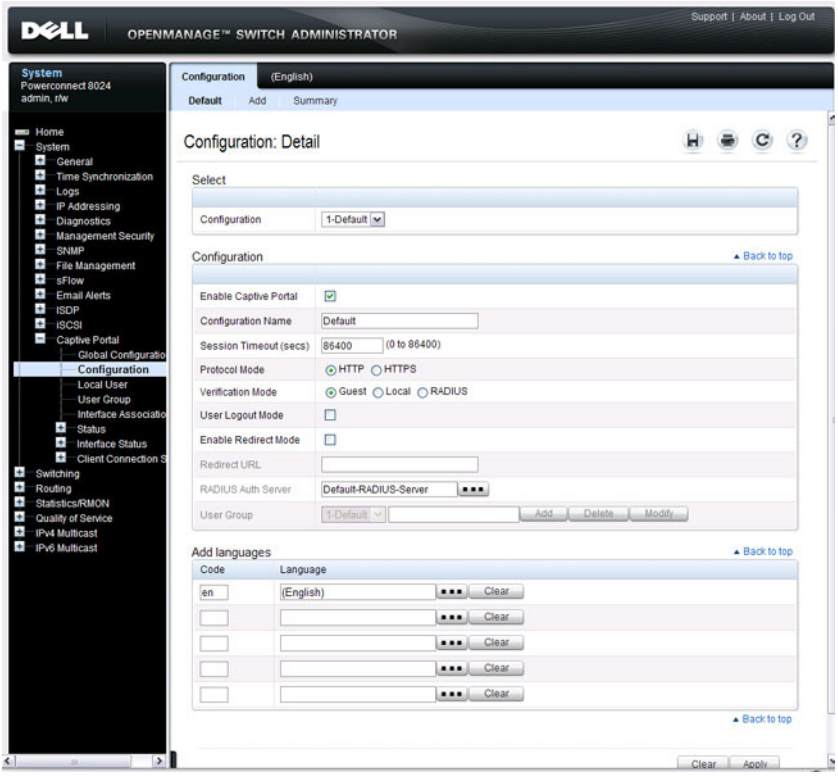
Captive Portal Configuration

Use the **Captive Portal Configuration** page to view summary information about captive portals on the system, add a captive portal, and configure existing captive portals.

The switch supports 10 Captive Portal configurations. Captive Portal configuration 1 is created by default and cannot be deleted. Each captive portal configuration can have unique guest or group access modes and a customized acceptance use policy that displays when the client connects.

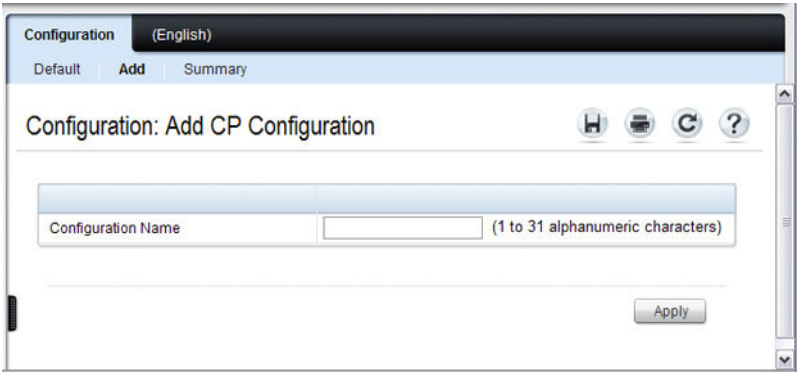
To display the **Captive Portal Configuration** page, click **System → Captive Portal → Configuration.**

Figure 18-5. Captive Portal Configuration



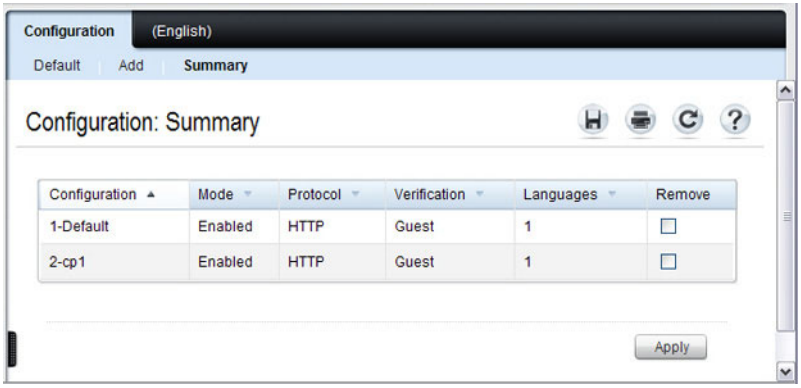
From the **Captive Portal Configuration** page, click **Add** to create a new Captive Portal instance.

Figure 18-6. Add Captive Portal Configuration



From the **Captive Portal Configuration** page, click **Summary** to view summary information about the Captive Portal instances configured on the switch.

Figure 18-7. Captive Portal Summary



Customizing a Captive Portal

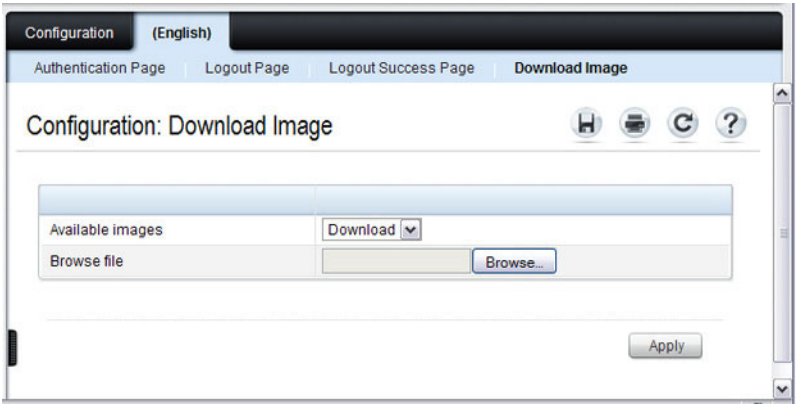
The procedures in this section customize the pages that the user sees when he or she attempts to connect to (and log off of) a network through the captive portal. These procedures configure the English version of the Default captive portal.

To configure the switch:

1. From the **Captive Portal Configuration** page click the **(English)** tab.The settings for the **Authentication Page** display, and the links to the Captive Portal customization appear.
2. Click **Download Image** to download one or more custom images to the switch. You can use a downloaded custom image for the branding logo (default: Dell logo) on the Authentication Page and Logout Success page, the account image (default: blue banner with keys) on the Authentication Page, and the background image (default: blank) on the Logout Success Page.

   NOTE: The image to download must be accessible from your local system. The image should be 5 KB max, 200x200 pixels, GIF or JPG format.

Figure 18-8. Captive Portal Authentication Page



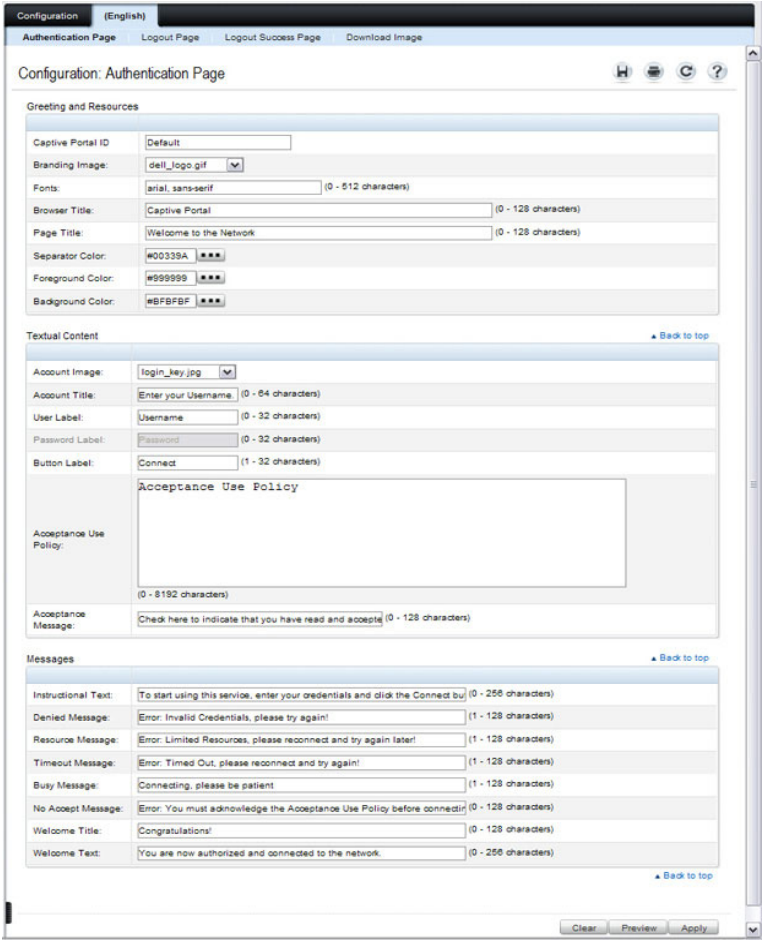
3. Make sure Download is selected in the **Available Images** menu, and click **Browse**.

4 Browse to the directory where the image to be downloaded is located and select the image.

5 Click **Apply** to download the selected file to the switch.

6 To customize the Authentication Page, which is the page that a user sees upon attempting to connect to the network, click the **Authentication Page** link.

Figure 18-9. Captive Portal Authentication Page



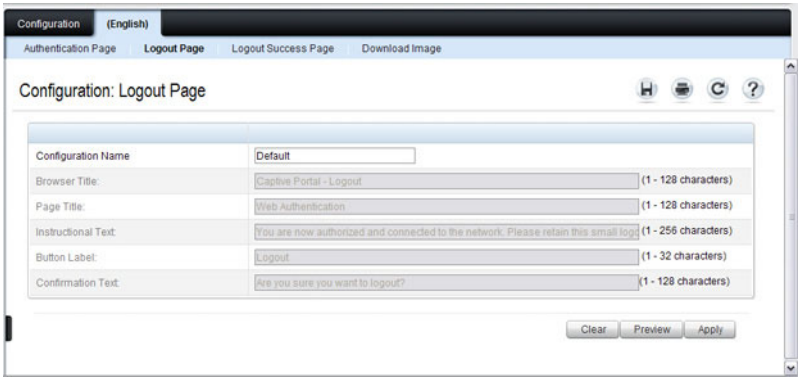
7 Select the branding image to use and customize other page components such as the font for all text the page displays, the page title, and the acceptance use policy.

8 Click **Apply** to save the settings to the running configuration or click **Preview** to view what the user will see. To return to the default views, click **Clear**.

9 Click the **Logout Page** link to configure the page that contains the logout window.

NOTE: You can configure the Logout Page settings only if the User Logout Mode is selected on the **Configuration** page. The User Logout Mode allows an authenticated client to deauthenticate from the network.

Figure 18-10. Captive Portal Logout Page

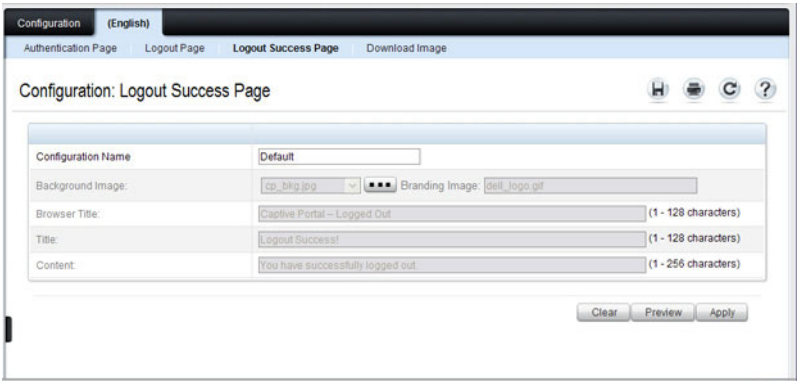


10 Customize the look and feel of the Logout Page, such as the page title and logout instructions.

11 Click **Apply** to save the settings to the running configuration or click **Preview** to view what the user will see. To return to the default views, click **Clear**.

12 Click the **Logout Success Page** link to configure the page that contains the logout window. A user is required to logout only if the User Logout Mode is selected on the **Configuration** page.

Figure 18-11. Captive Portal Logout Success Page



13 Customize the look and feel of the Logout Page, such as the background image and successful logout message.

14 Click **Apply** to save the settings to the running configuration or click **Preview** to view what the user will see. To return to the default views, click **Clear**.

Local User

You can configure a portal to accommodate guest users and authorized users. Guest users do not have assigned user names and passwords. Authorized users provide a valid user name and password that must first be validated against a local database or RADIUS server. Authorized users can gain network access once the switch confirms the user's credentials.

By default, each Captive Portal instance contains the default group. The default group can be renamed, or a different group can be created and assigned to each Captive Portal instance. A Captive Portal instance can be associated to one user group only. A user, however, can be assigned to multiple groups.

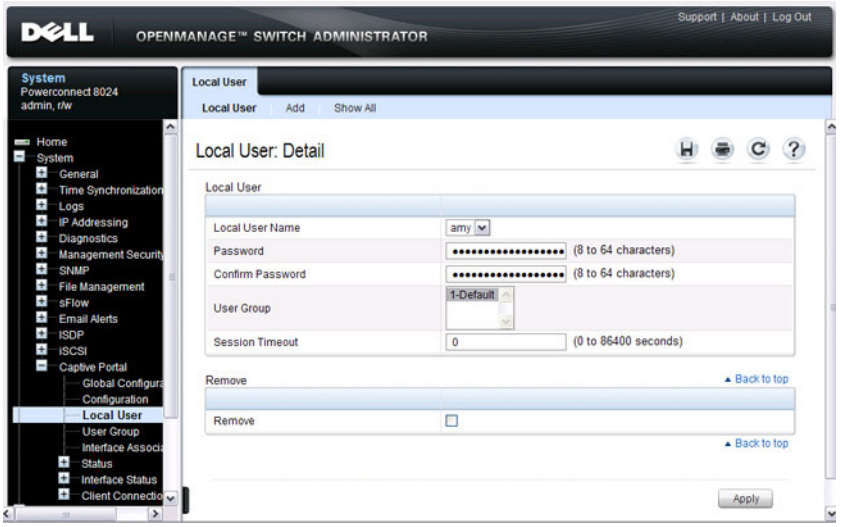
The **Local User** page allows you to add authorized users to the local database, which can contain up to 128 user entries. You can also add and delete users from the local database from the **Local User** page.

To display the **Local User** page, click **System → Captive Portal → Local User.**

Figure 18-12 shows the **Local User** page after a user has been added. If no users have been added to the switch, many of the fields do not display on the screen.

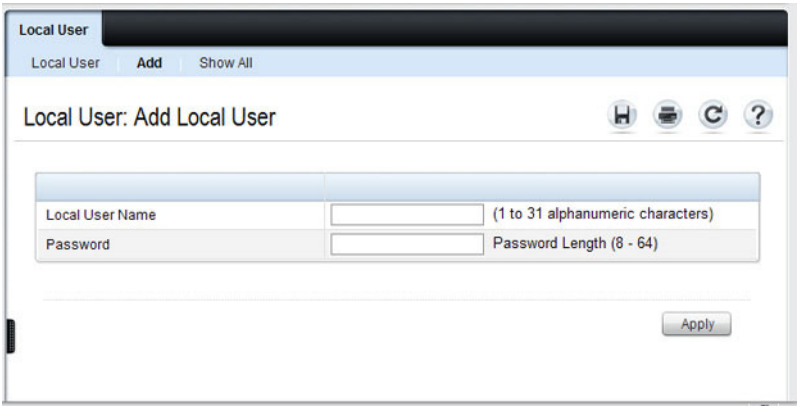
NOTE: Multiple user groups can be selected by holding the CTRL key down while clicking the desired groups.

Figure 18-12. Local User Configuration



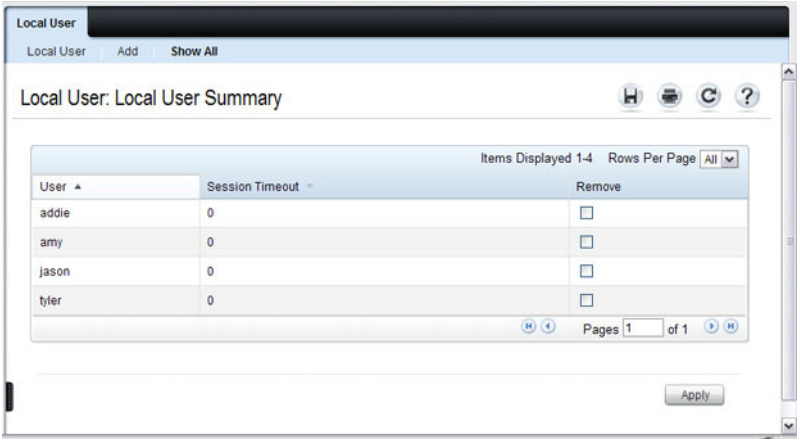
From the **Local User** page, click **Add** to add a new user to the local database.

Figure 18-13. Add Local User



From the **Local User** page, click **Show All** to view summary information about the local users configured in the local database.

Figure 18-14. Captive Portal Local User Summary



To delete a configured user from the database, select the Remove check box associated with the user and click **Apply**.

Configuring Users in a Remote RADIUS Server

You can use a remote RADIUS server client authorization. You must add all users to the RADIUS server. The local database does not share any information with the remote RADIUS database.

Table 18-2 indicates the RADIUS attributes you use to configure authorized captive portal clients. The table indicates both RADIUS attributes and vendor-specific attributes (VSA). VSAs are denoted in the Attribute column and are comma delimited (vendor ID, attribute ID).

Table 18-2. Captive Portal User RADIUS Attributes

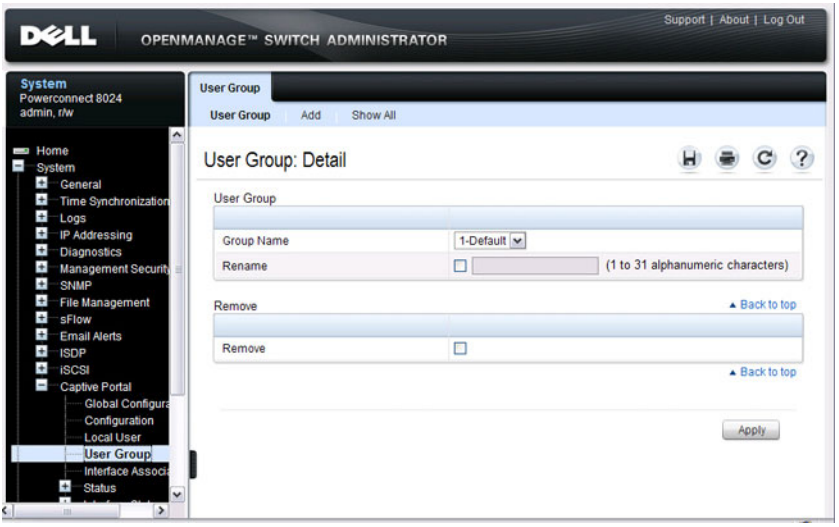
Attribute	Number	Description	Range	Usage	Default
User-Name	1	User name to be authorized	1-32 characters	Required	None
User-Password	2	User password	8-64 characters	Required	None
Session-Timeout	27	Logout once session timeout is reached (seconds). If the attribute is 0 or not present then use the value configured for the captive portal.	Integer (seconds)	Optional	0
Dell-Captive-Portal-Groups	6231, 127	A comma-delimited list of group names that correspond to the configured CP instance configurations.	String	Optional	None. The default group is used if not defined here

User Group

You can assign Local Users to User Groups that you create. If the Verification Mode is Local or RADIUS, you assign a User Group to a Captive Portal Configuration. All users who belong to the group are permitted to access the network through this portal. The User Group list is the same for all Captive Portal configurations on the switch.

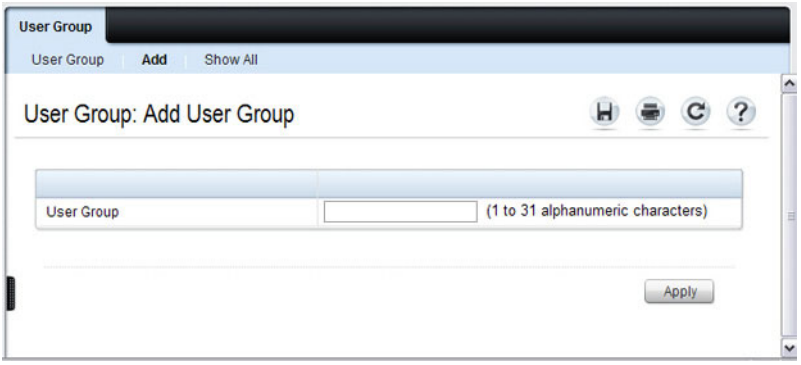
To display the **User Group** page, click **System → Captive Portal → User Group.**

Figure 18-15. User Group



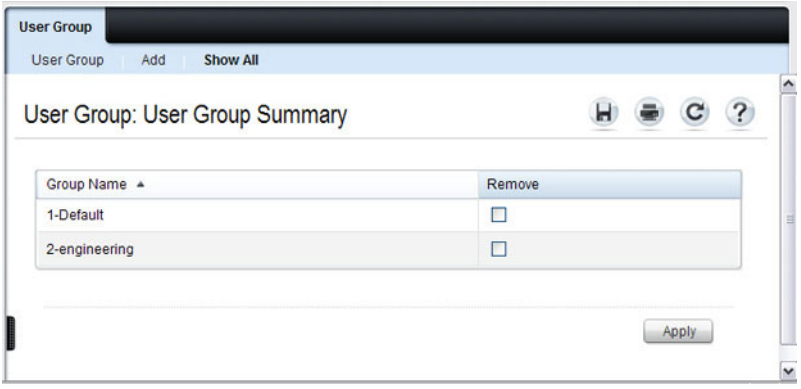
From the **User Group** page, click **Add** to configure a new user group.

Figure 18-16. Add User Group



From the **User Group** page, click **Show All** to view summary information about the user groups configured on the switch.

Figure 18-17. Captive Portal User Group Summary



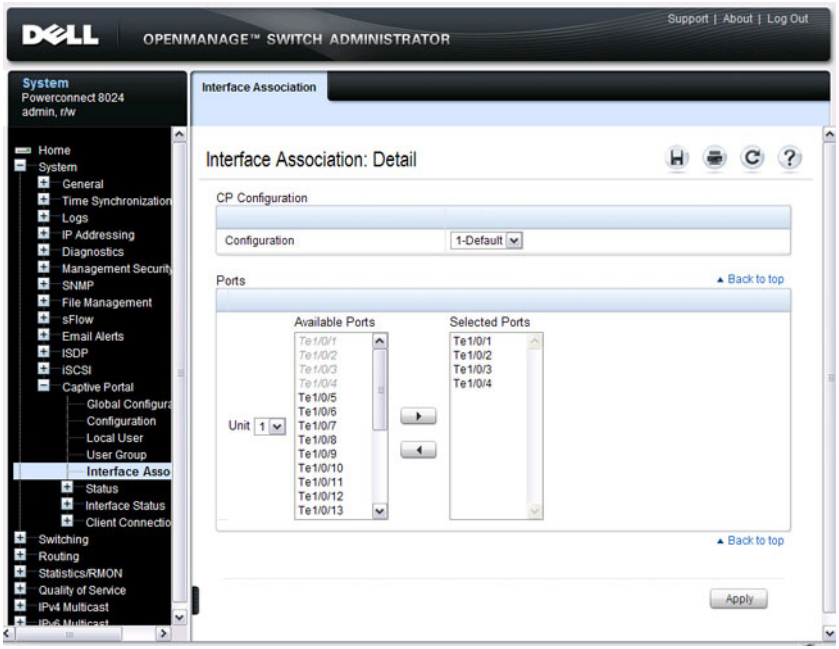
To delete a configured group, select the Remove check box associated with the group and click **Apply**.

Interface Association

From the **Interface Association** page, you can associate a configured captive portal with specific interfaces. The captive portal feature only runs on the interfaces that you specify. A captive portal can have multiple interfaces associated with it, but an interface can be associated to only one Captive Portal at a time.

To display the **Interface Association** page, click **System → Captive Portal → Interface Association.**

Figure 18-18. Captive Portal Interface Association



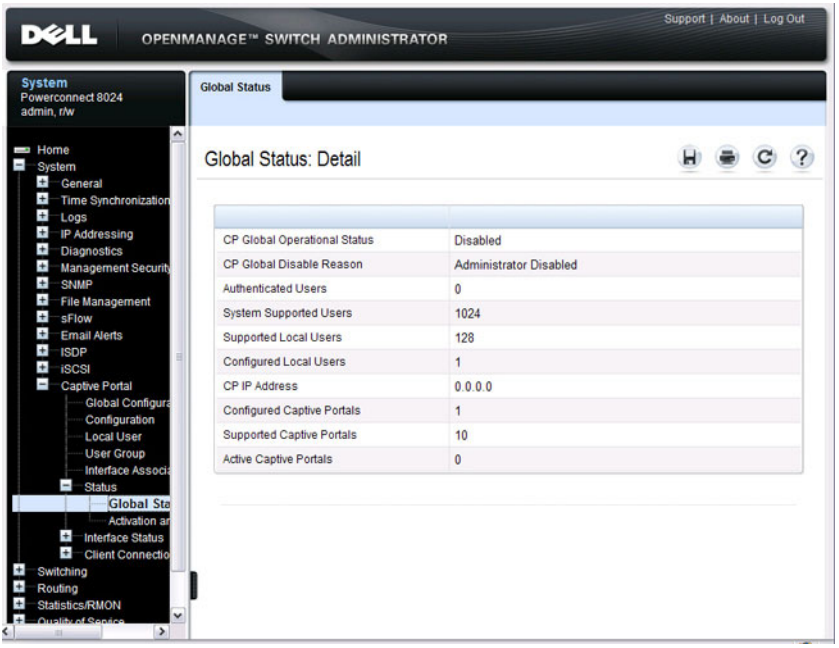
NOTE: When you associate an interface with a Captive Portal, the interface is disabled in the Interface List. Each interface can be associated with only one Captive Portal at a time.

Captive Portal Global Status

The **Captive Portal Global Status** page contains a variety of information about the Captive Portal feature. From the **Captive Portal Global Status** page, you can access information about the Captive Portal activity and interfaces.

To display the **Global Status** page, click **System → Captive Portal → Status → Global Status.**

Figure 18-19. Captive Portal Global Status



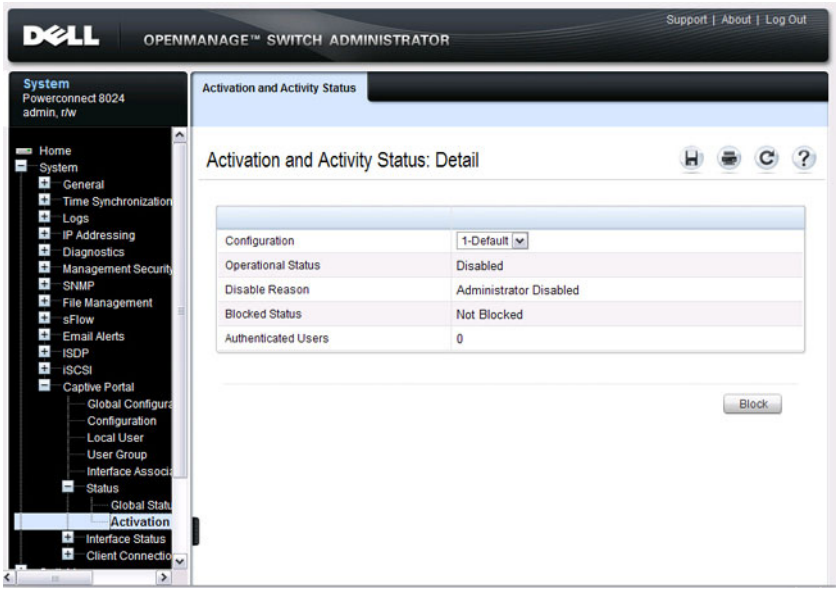
Captive Portal Activation and Activity Status

The **Captive Portal Activation and Activity Status** page provides information about each Captive Portal configured on the switch.

The **Captive Portal Activation and Activity Status** page has a drop-down menu that contains all captive portals configured on the switch. When you select a captive portal, the activation and activity status for that portal displays.

To display the **Activation and Activity Status** page, click **System** → **Captive Portal** → **Status** → **Activation and Activity Status.**

Figure 18-20. Captive Portal Activation and Activity Status



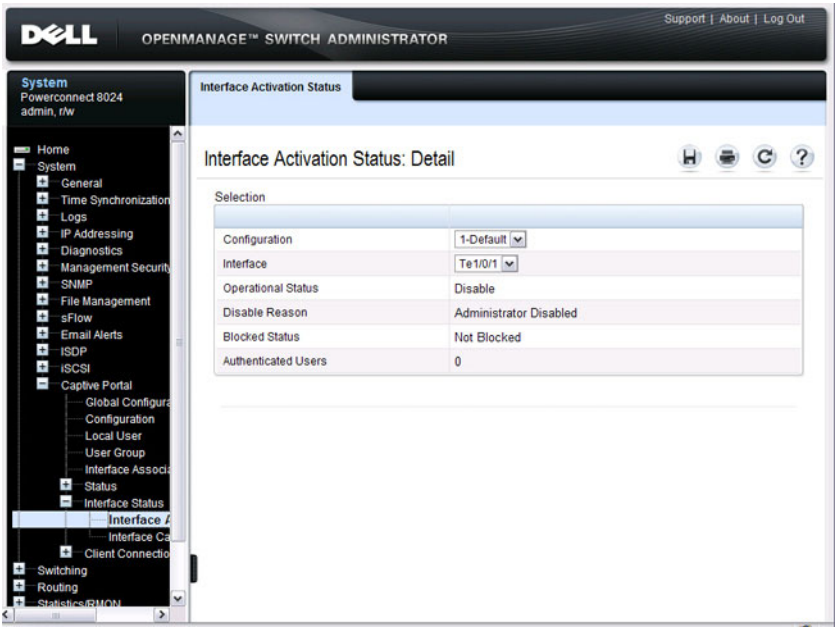
NOTE: Use the Block and Unblock buttons to control the blocked status. If the Captive Portal is blocked, users cannot gain access to the network through the Captive Portal. Use this function to temporarily protect the network during unexpected events, such as denial of service attacks.

Interface Activation Status

The **Interface Activation Status** page shows information for every interface assigned to a captive portal instance.

To display the **Interface Activation Status** page, click **System → Captive Portal → Interface Status → Interface Activation Status.**

Figure 18-21. Interface Activation Status

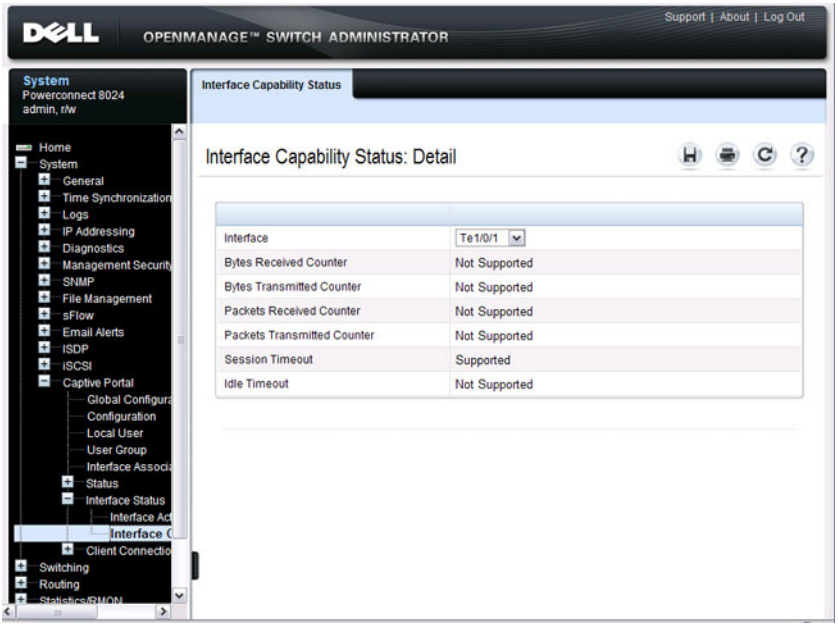


Interface Capability Status

The **Interface Capability Status** page contains information about interfaces that can have CPs associated with them. The page also contains status information for various capabilities. Specifically, this page indicates what services are provided through the Captive Portal to clients connected on this interface. The list of services is determined by the interface capabilities.

To display the **Interface Capability Status** page, click **System → Captive Portal → Interface Status → Interface Capability Status.**

Figure 18-22. Interface Capability Status

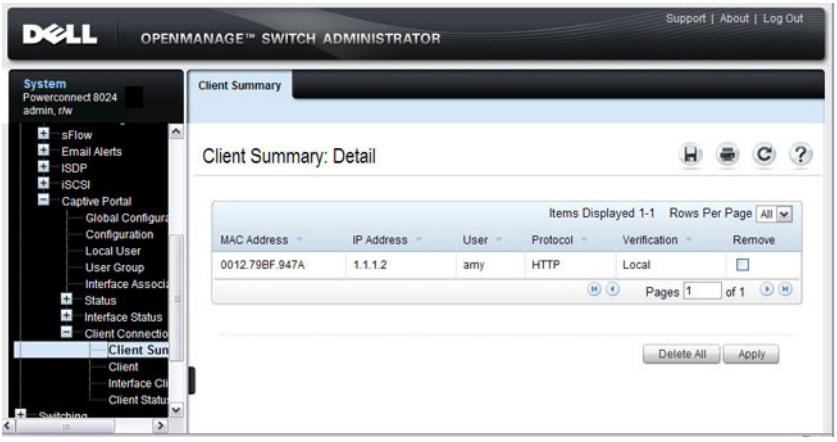


Client Summary

Use the **Client Summary** page to view summary information about all authenticated clients that are connected through the captive portal. From this page, you can manually force the captive portal to disconnect one or more authenticated clients. The list of clients is sorted by client MAC address.

To display the **Client Summary** page, click **System → Captive Portal → Client Connection Status → Client Summary.**

Figure 18-23. Client Summary



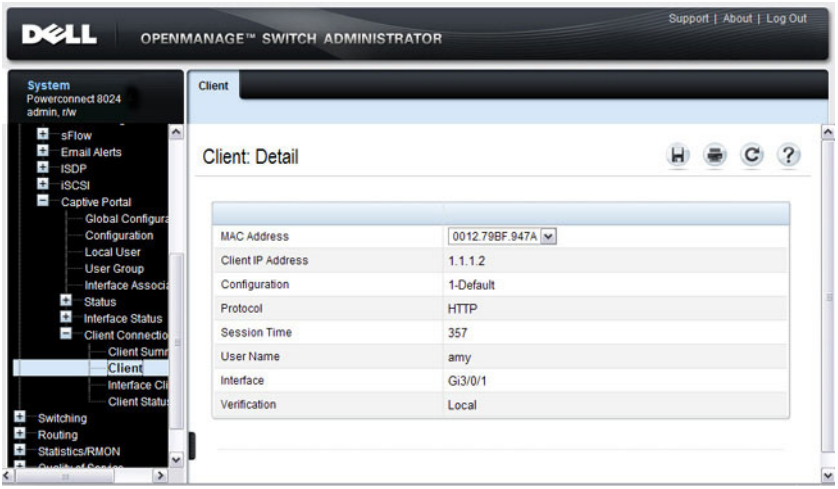
To force the captive portal to disconnect an authenticated client, select the **Remove** check box next to the client MAC address and click **Apply**. To disconnect all clients from all captive portals, click **Delete All.**

Client Detail

The **Client Detail** page shows detailed information about each client connected to the network through a captive portal.

To display the **Client Detail** page, click **System → Captive Portal → Client Connection Status → Client Detail.**

Figure 18-24. Client Detail

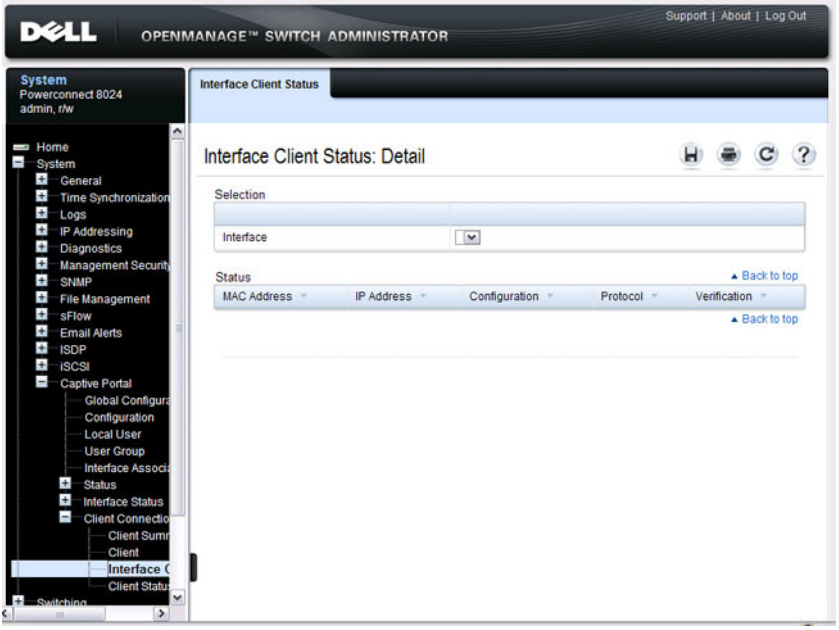


Captive Portal Interface Client Status

Use the **Interface Client Status** page to view clients that are authenticated to a specific interface.

To display the **Interface Client Status** page, click **System → Captive Portal → Client Connection Status → Interface Client Status.**

Figure 18-25. Interface - Client Status

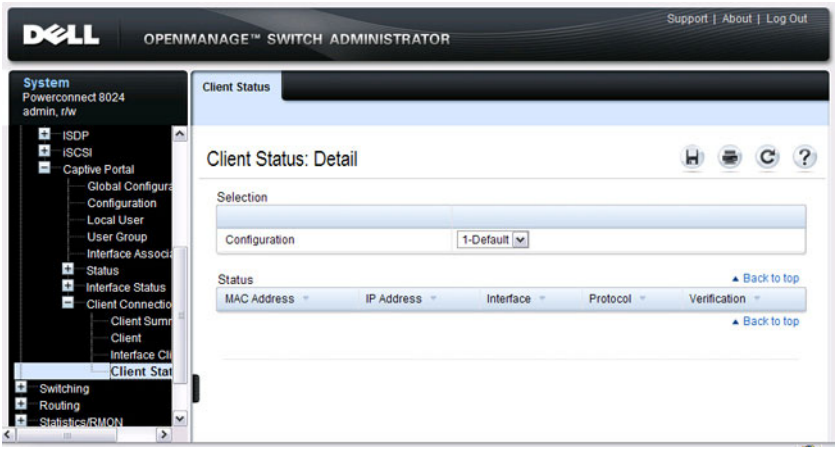


Captive Portal Client Status

Use the **Client Status** page to view clients that are authenticated to a specific Captive Portal configuration.

To display the **Client Status** page, click **System → Captive Portal → Client Connection Status → Client Status.**

Figure 18-26. Captive Portal - Client Status



Configuring a Captive Portal (CLI)

This section provides information about the commands you use to create and configure Captive Portal settings. For more information about the commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Configuring Global Captive Portal Settings

Beginning in Privileged EXEC mode, use the following commands to configure global Captive Portal settings.

Command	Purpose
configure	Enter global configuration mode.
captive-portal	Enter Captive Portal mode.
http port *port-num*	(Optional) Configure an additional HTTP port for Captive Portal to monitor. Use this command on networks that use an HTTP proxy server. *port-num* — The port number to monitor (Range: 1–65535, excluding ports 80, 443, and the configured switch management port).
https port *port-num*	(Optional) Configure an additional HTTPS port for Captive Portal to monitor. Use this command on networks that use an HTTPS proxy server. *port-num* — The port number to monitor Range: 1–65535, excluding ports 80, 443, and the configured switch management port).
authentication timeout *timeout*	(Optional) Configure the number of seconds the user has to enter valid credentials into the verification page. If the user exceeds the configured timeout, the verification page needs to be served again in order for the client to gain access to the network. *timeout* — The authentication timeout (Range: 60–600 seconds).
enable	Globally enable the Captive Portal feature.

Command	Purpose
CTRL + Z	Exit to Privileged EXEC mode.
show captive-portal [**status**]	View the Captive Portal administrative and operational status. Use the **status** keyword to view additional global Captive Portal information and summary information about all configured Captive Portal instances.

Creating and Configuring a Captive Portal

Beginning in Privileged EXEC mode, use the following commands to create a Captive Portal instance and configure its settings.

Command	Purpose
configure	Enter global configuration mode.
captive-portal	Enter Captive Portal mode.
configuration *cp-id*	Enter the captive portal instance mode *cp-id* — The Captive Portal instance (Range: 1–10). The Captive Portal configuration identified by CP ID 1 is the default CP configuration.
name *string*	Add a name to the Captive Portal instance. *string* — CP configuration name (Range: 1–32 characters).
protocol {http | https}	Specify whether to use HTTP or HTTPs during the Captive Portal user verification process.
verification {guest | local | radius}	Specify how to process user credentials the user enters on the verification page. • guest — Allows access for unauthenticated users (users that do not have assigned user names and passwords). • local — Authenticates users against a local user database. • radius — Authenticates users against a remote RADIUS database.
radius-auth-server *name*	Specify the name of the RADIUS server to use for RADIUS verification. Use the commands described in "Using RADIUS Servers to Control Management Access" on page 192 to configure RADIUS server settings for the switch.

Command	Purpose
user-logout	(Optional) Enable user logout mode to allow an authenticated client to deauthenticate from the network. If this option is clear or the user does not specifically request logout, the client connection status remains authenticated until the CP deauthenticates the user, for example by reaching the idle timeout or session timeout values.
redirect	(Optional) Enable the redirect mode for a Captive Portal configuration so that the user is redirected to a specific Web page after the verification or authentication process. When the redirect mode is not enabled, the user sees the Captive Portal welcome page after the verification or authentication process.
redirect-url *url*	(Optional) Specify the web page that the users sees after successful verification or authentication through the Captive Portal. *url* — The URL for redirection (Range: 1–512 characters).
group *group-number*	(For Local and RADIUS verification) Configure the group number associated with this Captive Portal configuration. By default, only the default group exists. To assign a different user group to the Captive Portal instance, you must first configure the group. *group-number* — The number of the group to associate with this configuration (Range: 1–10)
session-timeout *timeout*	(Optional) Enter the number of seconds to wait before terminating a session. A user is logged out once the session timeout is reached. You can also set the session timeout for each user if the Captive Portal requires authentication. *timeout* — Session timeout. 0 indicates timeout not enforced (Range: 0–86400 seconds)
interface *interface*	Associate an interface with this Captive Portal. (The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**.
enable	Enable the Captive Portal instance.

Command	Purpose
block	(Optional) Block all traffic for a Captive Portal configuration. If the Captive Portal is blocked, users cannot gain access to the network through the Captive Portal. Use this function to temporarily protect the network during unexpected events, such as denial of service attacks.
CTRL + Z	Exit to Privileged EXEC mode.
show captive-portal configuration *cp-id* [**status** | **interface**]	View summary information about a Captive Portal instance. • *cp-id* — The Captive Portal instance (Range: 1–10). • **status** — View additional information about the Captive Portal instance. • interface — View information about the interface(s) associated with the specified Captive Portal.
show captive-portal interface configuration *cp-id* **status**	View information about the interfaces associated with the specified Captive Portal instance. *cp-id* — The Captive Portal instance (Range: 1–10).

NOTE: To return the default Captive Portal instance to its default values, use the **clear** command in the Captive Portal Instance mode. You must also use the **no interface** *interface* command to remove any associated interfaces from the instance.

Configuring Captive Portal Groups and Users

Beginning in Privileged EXEC mode, use the following commands to create a Captive Portal group. You can use the default group, or you can create a new group.

Command	Purpose
configure	Enter global configuration mode.
captive-portal	Enter Captive Portal mode.
user group *group-id* [**name** *name*]	Configure a group. Each Captive Portal that requires authentication has a group associated with it. Only the users who are members of that group can be authenticated if they connect to the Captive Portal. • *group-id* — Group ID (Range: 1–10). • *name* — Group name (Range: 1–32 characters).
user *user-id* **name** *name*	Create a new user for the local user authentication database. • *user-id*—User ID (Range: 1–128). • *name*—user name (Range: 1–32 characters).
user *user-id* **password** *password*	Configure the password for the specified user. • *user-id*—User ID (Range: 1–128). • *password*—User password (Range: 8–64 characters).
user *user-id* **group** *group-id*	Associate a group with a Captive Portal user. A user can be associated with more than one group. • *user-id* — User ID (Range: 1–128). • *group-id* — Group ID (Range: 1–10).
user *user-id* **session-timeout** *timeout*	Enter the number of seconds to wait before terminating a session for the specified user. The user is logged out once the session timeout is reached. • *user-id* — User ID (Range: 1–128). • *timeout* — Session timeout. 0 indicates timeout not enforced (Range: 0–86400 seconds)

Command	Purpose
user group *group-id* **moveusers** *new-group-id*	(Optional) Move all of the users in a group to a different group. This command removes the users from the group specified by *group-id*. • *group-id* — Group ID (Range: 1–10). • *new-group-id* — Group ID (Range: 1–10).
CTRL + Z	Exit to Privileged EXEC mode.
show captive-portal user [*user-id*]	View summary information about all users configured in the local database. Specify the user ID to view additional information about a user. *user-id* — User ID (Range: 1–128).
clear captive portal users	(Optional) Delete all captive portal user entries from the local database.

Managing Captive Portal Clients

The commands in this section are all executed in Privileged EXEC mode. Use the following commands to view and manage clients that are connected to a Captive Portal.

Command	Purpose
show captive-portal configuration [*cp-id*] **client status**	Display information about the clients authenticated to all Captive Portal configurations or a to specific configuration. *cp-id* — The Captive Portal instance (Range: 1–10).
show captive-portal interface *interface* **client status**	Display information about clients authenticated on all interfaces or no a specific interface. *interface* — Specific Ethernet interface, such as te1/0/8.
show captive-portal client [*macaddr*] **status**	Display client connection details or a connection summary for connected Captive Portal users. *macaddr* — The MAC address of the client.
captive-portal client deauthenticate *macaddr*	Deauthenticate a specific captive portal client. *macaddr* — The MAC address of the client.

Captive Portal Configuration Example

The manager of a resort and conference center needs to provide wired Internet access to each guest room at the resort and in each conference room. Due to legal reasons, visitors and guests must agree to the resort's acceptable use policy to gain network access. Additionally, network access from the conference rooms must be authenticated. The person who rents the conference room space receives a list username and password combinations upon arrival. Hotel employees have their own Captive Portal.

The network administrator for the resort and conference center decides to configure the three Captive Portals Table 18-3 describes.

Table 18-3. Captive Portal Instances

Captive Portal Name	Description
Guest	Free Internet access is provided in each guest room, but guests must enter a name and agree to the acceptable use policy before they can gain access. The manager wants guests to be redirected to the resort's home web page upon successful verification. No logout is required.
Conference	Because physical access to the conference rooms is less secure than access to each guest room, the manager wants to ensure that people who connect to the network through a port in a conference room are authenticated. The Conference Captive Portal uses the local database for authentication.
Employee	To gain network access, resort employees must enter a username and password that is stored on a RADIUS server.

Configuration Overview

The following steps provide an overview of the process you use to configure the Captive Portal feature.

To configure the switch:

1. If you plan to use a RADIUS server for authentication, configure the RADIUS server settings on the switch.
2. If authentication is required, configure the user groups to associate with each Captive Portal.
3. Create (add) the Captive Portals.
4. Configure the Captive Portal settings for each Captive Portal, such as the verification mode.
5. Associate interfaces with the Captive Portal instances.
6. Download the branding images, such as the company logo, to the switch.

   The images you download must be accessible from the switch, either on the system you use to manage the switch or on a server that is on the same network as the switch.

   NOTE: You must use the web interface to download images.

7. Customize the authentication, logout, and logout success web pages that a Captive Portal user will see.

   Dell recommends that you use Use Dell OpenManage Administrator to customize the Captive Portal authentication, logout, and logout success pages. A **Preview** button is available to allow you to see the pages that a Captive Portal user will see.
8. If you use the local database for user authentication, configure the users on the switch.
9. If you use a RADIUS server for authentication, add the users to the database on the RADIUS server.
10. Associate interfaces with the Captive Portal instances.
11. Globally enable Captive Portal.

Detailed Configuration Procedures

Use the following steps to perform the Captive Portal configuration:

1. Configure the RADIUS server information on the switch.

   In this example, the RADIUS server IP address is 192.168.2.188, and the RADIUS server name is luxury-radius.

   ```
   console#configure
   console(config)#radius-server host 192.168.12.182
   console(Config-auth-radius)#name luxury-radius
   console(Config-auth-radius)#exit
   ```

2. Configure the Captive Portal groups.

   ```
   console(config)#captive-portal
   console(config-CP)#user group 2 name Conference
   console(config-CP)#user group 3 name Employee
   console(config-CP)#exit
   ```

3. Configure the Guest Captive Portal.

   ```
   console(config)#captive-portal
   console(config-CP)#configuration 2
   console(config-CP 2)#name Guest
   console(config-CP 2)#redirect
   console(config-CP 2)#redirect-url
   http://www.luxuryresorturl.com
   console(config-CP 2)#interface te1/0/1
   console(config-CP 2)#interface te1/0/2
   ...
   console(config-CP 2)#interface te1/0/4
   console(config-CP 2)#exit
   ```

4. Configure the Conference Captive Portal.

   ```
   console(config-CP)#configuration 3
   console(config-CP 3)#name Conference
   console(config-CP 3)#verification local
   console(config-CP 3)#group 2
   console(config-CP 4)#interface te1/0/8
   ...
   console(config-CP 4)#interface te1/0/15
   console(config-CP 3)#exit
   ```

5. Configure the Employee Captive Portal.

```
console(config-CP)#configuration 4
console(config-CP 4)#name Employee
console(config-CP 4)#verification radius
console(config-CP 4)#group 3
console(config-CP 4)#interface te1/0/18
...
console(config-CP 4)#interface te1/0/22
console(config-CP 4)#exit
```

6. Use the web interface to customize the Captive Portal pages that are presented to users when they attempt to connect to the network.

   NOTE: Captive Portal page customization is supported only through the Web interface. For information about customizing the Captive Portal pages, see "Customizing a Captive Portal" on page 411.

7. Add the Conference users to the local database.

```
console(config-CP)#user 1 name EaglesNest1
console(config-CP)#user 1 password
Enter password (8 to 64 characters): *********
Re-enter password: *********
console(config-CP)#user 1 group 2
```

   Continue entering username and password combinations to populate the local database.

8. Add the User-Name, User-Password, Session-Timeout, and Dell-Captive-Portal-Groups attributes for each employee to the database on the RADIUS server.

9. Globally enable the Captive Portal.

```
console(config-CP)#enable
```

19

Configuring Port Characteristics

This chapter describes how to configure physical switch port characteristics, including settings such as administrative status and maximum frame size. This chapter also describes the link dependency feature.

The topics covered in this chapter include:

- Port Overview
- Default Port Values
- Configuring Port Characteristics (Web)
- Configuring Port Characteristics (CLI)
- Port Configuration Examples

Port Overview

A port is a physical interface. Cables physically connect ports on devices such as PCs or servers to ports on the switch to provide access to the network. The type of physical ports available on your PowerConnect 8000-series and 8100-series switches depends on the model.

What Physical Port Characteristics Can Be Configured?

Table 19-1 provides a summary of the physical characteristics that can be configured on the switch ports.

Table 19-1. Port Characteristics

Feature	Description
Administrative status	Controls whether the port is administratively enabled or disabled.
Description	Provides a text-based description of the port.
Auto negotiation	Enables a port to advertise its transmission rate, duplex mode and flow control abilities to its partner.

Table 19-1. Port Characteristics (Continued)

Feature	Description
Speed	Specifies the transmission rate for frames.
Duplex mode	Specifies whether the interface supports transmission between the switch and the connected client in one direction at a time (half) or both directions simultaneously (both).
Maximum frame size	Indicates the maximum frame size that can be handled by the port.
Flow control	This is a global setting that affects all ports. For more information about this feature, see "Configuring Port-Based Traffic Control" on page 665.
Storm control	For more information about this feature, see "Configuring Port-Based Traffic Control" on page 665.
Port security	For more information about this feature, see "Configuring Port and System Security" on page 457.
Protected port	For more information about this feature, see "Configuring Port-Based Traffic Control" on page 665.

What is Link Dependency?

The link dependency feature provides the ability to enable or disable one or more ports based on the link state of one or more different ports. With link dependency enabled on a port, the link state of that port is dependent on the link state of another port. For example, if port A is dependent on port B and the switch detects a link loss on port B, the switch automatically brings down the link on port A. When the link is restored to port B, the switch automatically restores the link to port A.

You can create a maximum of 16 groups.

Link Action

The link action specifies the action that the group members will take when the dependent port is down. The group members can transition to the same state as the dependant port, or they can transition to the opposite state. In other words, if the link action is **down** and the dependent port goes down, the members ports will go down as well. Conversely, when the link action is **up** and the dependant link goes down, the group member ports are enabled (brought up).

Creating a link dependency group with the **up** link action essentially creates a backup link for the dependent link and alleviates the need to implement STP to handle the fail-over.

Link Dependency Scenarios

The Link Dependency feature supports the scenarios in the following list.

NOTE: Whether the member ports or LAGs are brought *up* or *down* depends on the link action.

- Port dependent on port — If a port loses the link, the switch brings up/down the link on another port.
- Port dependent on LAG — If all ports in a channel-group lose the link, the switch brings up/down the link on another port.
- LAG dependent on port — If a port loses the link, the switch brings up/down all links in a channel-group.
- Multiple port command — If a group of ports lose their link, the switch brings up/down the link on another group of ports.
- Overlapping ports — Overlapping ports on different groups will be brought up/down only if both dependent ports lose the link.

What Interface Types are Supported?

The physical ports on the switch include the out-of-band (OOB) interface and 10-Gigabit Ethernet switch ports. The OOB interface supports a limited set of features and is for switch management only. The Ethernet switch ports support many logical features that are often supported by logical interfaces. The switch supports the following types of logical interfaces:

- Port-based VLANs — For more information, see "Configuring VLANs" on page 539.
- VLAN routing interfaces — For more information, see "Configuring Routing Interfaces" on page 843.
- Link Aggregation Groups (LAGs), which are also called port channels) — For more information, see "Configuring Link Aggregation" on page 781.
- Tunnels — For more information, see "Configuring Routing Interfaces" on page 843.
- Loopback interfaces — For more information, see "Configuring Routing Interfaces" on page 843.

What is Interface Configuration Mode?

When you use the CLI to configure physical or logical characteristics for an interface, you must enter Interface Configuration Mode for that interface. To enter the mode, type the keyword **interface** followed by the interface type and additional information to identify the interface, such as the interface number.

To enter Interface Configuration mode for a physical switch port, the following information is required:

- Type — For physical switch ports, the type is 10-Gibabit Ethernet (tengigabitethernet or te) for 10,000 Mbps Ethernet ports.
- Stack member number— The PowerConnect 8000-series and 8100-series switches are standalone, non-stacking switches, so the member number is always 1.
- Module (slot) number—For the PowerConnect 8000-series and 8100-series switches, the slot number is always 0.

- Port number—The number assigned to the port. For front-panel ports the port number is written above or below each port. Odd-numbered ports are on the top row, and even-numbered ports are on the bottom row. The port numbers increase from left to right.

For example, to enter Interface Configuration mode for 10-Gigabit Ethernet port 10, use the following command:

```
console(config)#interface tengigabitEthernet 1/0/10
```

NOTE: When you enter Interface Configuration mode, the command prompt changes and identifies the interface. In the previous example, the command prompt becomes `console(config-if-Te1/0/10)#.`

For many features, you can configure a range of interfaces. When you enter Interface Configuration mode for multiple interfaces, the commands you execute apply to all interfaces specified in the range.

To enter Interface Configuration mode for a range of interfaces, include the keyword **range** and specify the interfaces to configure. For example, to apply the same configuration to ports 1-10 on a standalone switch, use the following command:

```
console(config)#interface range tengigabitEthernet
1/0/1-10
```

To enter Interface Configuration mode for ports 3, 4, 5, 12, and 14 on a standalone switch, use the following command:

```
console(config)#interface range tengigabitEthernet
1/0/3-5,1/0/12,1/0/14
```

NOTE: You can switch to another interface or range of interfaces by entering the interface command while in Interface Configuration mode. It is not necessary to exit Interface Configuration mode to select a different interface.

Default Port Values

Table 19-2 lists the default values for the port characteristics that this chapter describes.

Table 19-2. Default Port Values

Feature	Description
Administrative status	All ports are enabled
Description	None defined
Auto negotiation	Enabled
Speed	Auto negotiate
Duplex mode	Auto negotiate
Flow control	Enabled
Maximum frame size	1518
Link Dependency	None configured

Configuring Port Characteristics (Web)

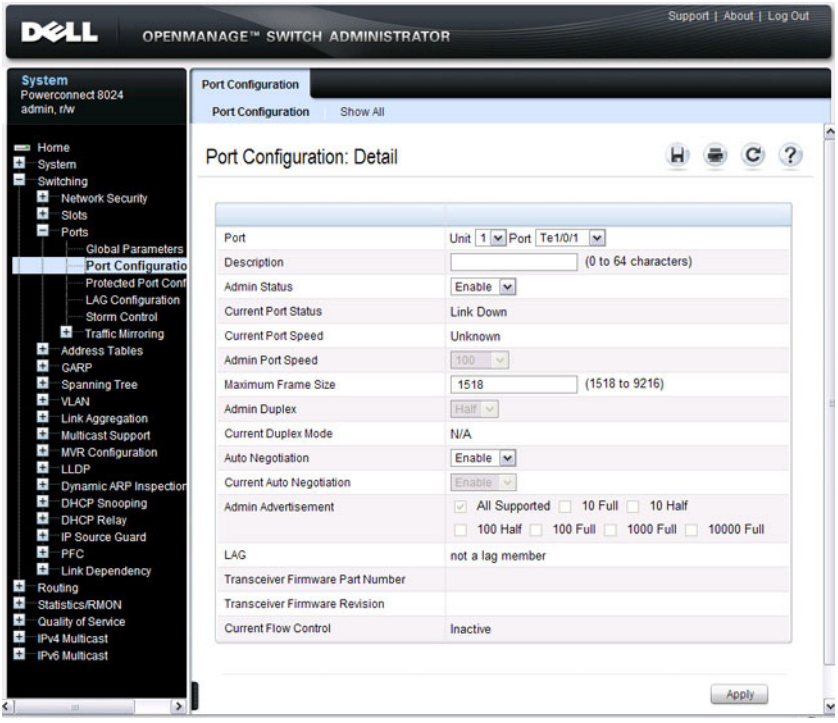
This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring port characteristics on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

Port Configuration

Use the **Port Configuration** page to define port parameters.

To display the **Port Configuration** page, click **Switching** → **Ports** → **Port Configuration** in the navigation panel.

Figure 19-1. Port Configuration

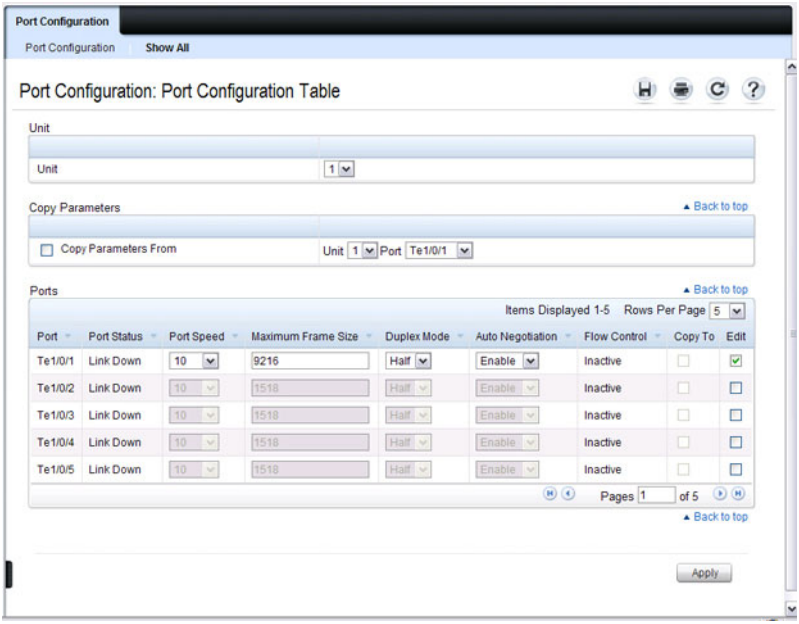


Configuring Multiple Ports

To configure port settings on multiple ports:

1. Open the **Port Configuration** page.
2. Click **Show All** to display the **Port Configuration Table** page.
3. In the **Ports** list, select the check box in the **Edit** column for the port to configure.
4. Select the desired settings.
5. Click **Apply**.

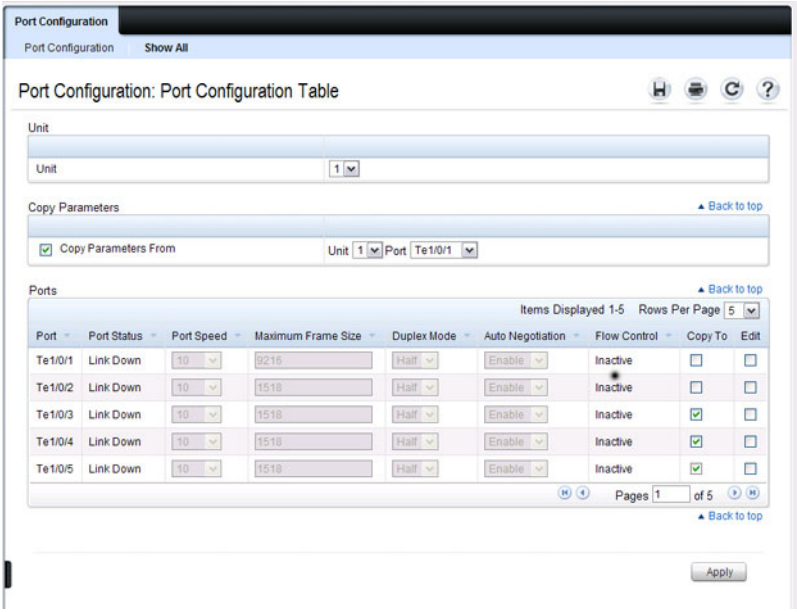
Figure 19-2. Configure Port Settings



6. Select the **Copy Parameters From** check box, and select the port with the settings to apply to other ports.
7. In the **Ports** list, select the check box(es) in the **Copy To** column that will have the same settings as the port selected in the **Copy Parameters From** field.

In the following example, Ports 3, 4, and 5 will be updated with the settings that are applied to Port 1.

Figure 19-3. Copy Port Settings



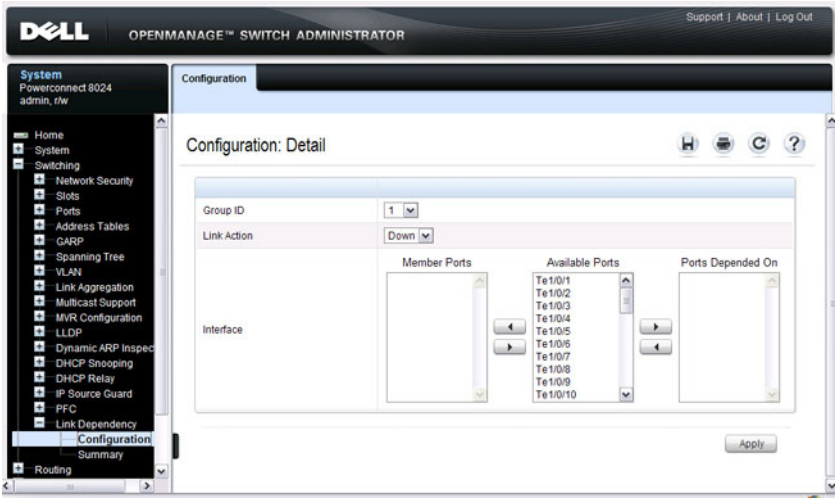
8 Click **Apply**.

Link Dependency Configuration

Use the **Link Dependency Configuration** page to create link dependency groups. You can create a maximum of 16 dependency groups. The page displays the groups whether they have been configured or not.

To display the **Link Dependency Configuration** page, click **Switching → Link Dependency → Configuration** in the navigation panel.

Figure 19-4. Link Dependency Configuration



Creating a Link Dependency Group

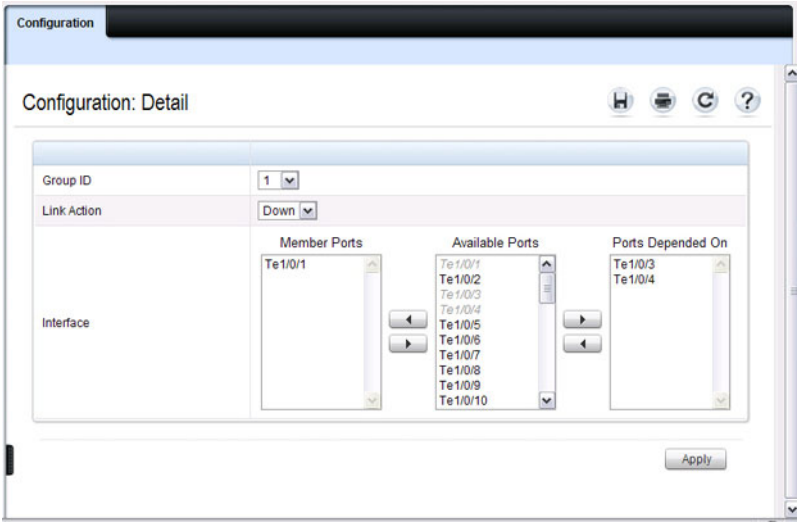
To create link dependencies:

1. Open the **Link Dependency Configuration** page.
2. In the **Group ID** field, select the ID of the group to configure.
3. Specify the link action.
4. To add a port to the **Member Ports** column, click the port in the **Available Ports** column, and then click the < button to the left of the **Available Ports** column. Ctrl + click to select multiple ports.

5 To add a port to the **Ports Depended On** column, click the port in the **Available Ports** column, and then click the > button to the right of the **Available Port**s column.

In the following example, Group 1 is configured so that Port 3 is dependent on Port 4.

Figure 19-5. Link Dependency Group Configuration



6 Click **Apply**.

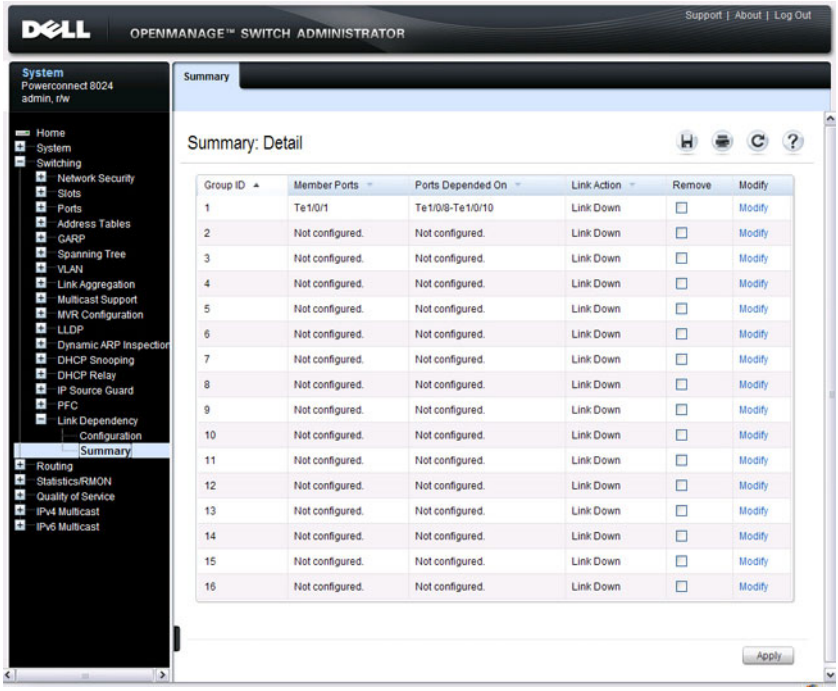
The Link Dependency settings for the group are modified, and the device is updated.

Link Dependency Summary

Use the **Link Dependency Summary** page to view all link dependencies on the system and to access the **Link Dependency Configuration** page. You can create a maximum of 16 dependency groups. The page displays the groups whether they have been configured or not.

To display the **Link Dependency Summary** page, click **Switching** → **Link Dependency** → **Link Dependency Summary** in the navigation panel.

Figure 19-6. Link Dependency Summary



To configure a group, click the **Modify** link associated with the ID of the group to configure. Clicking the **Modify** link takes you to the **Link Dependency Configuration** page. The Group ID is automatically selected based on the link that was clicked.

Configuring Port Characteristics (CLI)

This section provides information about the commands you use to configure port characteristics. For more information about the commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Configuring Port Settings

Beginning in Privileged EXEC mode, use the following commands to configure various port settings.

Command	Purpose
configure	Enter Global Configuration mode.
interface *interface*	Enter interface configuration mode for the specified interface. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**.
	You can also specify a range of interfaces with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12.
description *string*	Add a description to the port. The text string can be from1-64 characters.
shutdown	Administratively disable the interface.
speed {10 |100|1000|10000 | **auto** [100|1000|10000]}	Configure the speed of a given Ethernet interface or allow the interface to automatically detect the speed.
	If you use the 10, 100, or 1000 keywords with the auto keyword, the port auto negotiates only at the specified speeds.
	On combo ports, it is possible to configure auto negotiation even if only the fiber interface is active. The auto negotiation settings will be utilized when the copper port is active. Auto negotiation settings are ignored for the fiber ports.

Command	Purpose
duplex {half | full | auto}	Configure the full/half duplex operation of a given Ethernet interface, or enable duplex auto negotiation. Fiber ports must always be configured full-duplex. auto negotiation is never used on fiber ports.
mtu *size*	Enable jumbo frames on an interface by adjusting the maximum size of a packet.
CTRL + Z	Exit to Privileged EXEC mode.
show interfaces status	Show summary information about all interfaces.
show interfaces configuration	View a summary of the configuration for all ports.
show interfaces advertise	View a summary of the speeds that are advertised on each port.
show interfaces description	View configured descriptions for all ports.
show interfaces detail *interface*	View detailed information about the specified port.

Configuring Link Dependencies

Beginning in Privileged EXEC mode, use the following commands to configure ports that are dependent on the state of other ports.

Command	Purpose
configure	Enter Global Configuration mode.
link-dependency group *group_id*	Enter the link-dependency mode to configure a link-dependency group.
add *interface*	Add member ports to the group.
	The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**. You can also add port channels (LAGs) as members by using the keyword **port-channel** followed by an ID.
	You can also specify a range of interfaces. For example, **interface tengigabitethernet 1/0/8-10,1/0/20** configures interfaces 8, 9, 10 and 20.

Command	Purpose
depends-on *interface*	Specify the port(s) upon which the member ports are dependent. For information about the *interface* variable, see the previous command description.
action {**down**|**up**}	Specifies the action the member ports take when the dependent link goes down. • **down**—When the dependent link is down, the group members are down (the members are up otherwise). • **up**—When the dependent link goes down, the group members are brought up (the members are down otherwise)
CTRL + Z	Exit to Privileged EXEC mode.
show link-dependency [**group** *group_id*]	View link dependency settings for all groups or for the specified group, along with the group state.

Port Configuration Examples

This section contains the following examples:

- Configuring Port Settings
- Configuring a Link Dependency Groups

Configuring Port Settings

The commands in this example specify the speed and duplex mode for port 1 (tengigabitethernet 1/0/1) and change the MTU size for ports 10, 11, 12, 20, and 25.

To configure the switch:

1. Enter Interface Configuration mode for port 1.

   ```
   console#configure
   console(config)#interface tengigabitEthernet 1/0/1
   ```

2. Change the speed and duplex settings for the port.

   ```
   console(config-if-Te1/0/1)#speed 100
   console(config-if-Te1/0/1)#duplex full
   console(config-if-Te1/0/1)#exit
   ```

3. Enter Interface Configuration mode for ports 10, 11, 12, 20, and 24.

   ```
   console(config)#interface range tengigabitEthernet
   1/0/10-12,1/0/20,1/0/24
   ```

4. Enable jumbo frame support on the interfaces.

   ```
   console(config-if)#mtu 9216
   console(config-if)#CTRL + Z
   ```

5. View summary information about the ports

   ```
   console#show interfaces configuration

   Port       Type                Duplex  Speed    Neg   Admin St.
   ---------  -------------       ------  -------  ----  -----
   Te1/0/1    10G - Level         N/A     Unknown  Auto  Up
   Te1/0/2    10G - Level         N/A     Unknown  Auto  Up
   Te1/0/3    10G - Level         N/A     Unknown  Auto  Up
   Te1/0/4    10G - Level         N/A     Unknown  Auto  Up
   Te1/0/5    10G - Level         N/A     Unknown  Auto  Up

   --More-- or (q)uit
   ```

Configuring a Link Dependency Groups

The commands in this example create two link dependency groups. Group 1 has port 3 as a member port that is dependent on port 4. The group uses the default link action, which is down. This means that if port 4 goes down, port 3 goes down. When port 4 returns to the up state, port 3 is brought back up. In Group 2, port 6 dependent on port-channel (LAG) 1, and the link action is up. If port-channel 1 goes down, port 6 is brought up. This also means that when port-channel 1 is up, port 6 is down.

To configure the switch:

1. Enter the configuration mode for Group 1.

```
console#configure
console(config)#link-dependency group 1
```

2. Configure the member and dependency information for the group.

```
console(config-linkDep-group-1)#add
tengigabitethernet 1/0/3
console(config-linkDep-group-1)#depends-on
tengigabitethernet 1/0/4
console(config-linkDep-group-1)#exit
```

3. Enter the configuration mode for Group 2

```
console(config)#link-dependency group 2
console(config-linkDep-group-2)#add
tengigabitethernet 1/0/6
console(config-linkDep-group-2)#depends-on port-
channel 1
console(config-linkDep-group-2)#action up
console(config-linkDep-group-2)#CTRL + Z
```

4. View the configured link dependency groups.

```
console#show link-dependency

GroupId  Member Ports     Ports Depended On    Link Action
-------  -------------   -----------------    ----------
1        Te1/0/3          Te/0/4              Link Down
2        Te/0/6          ch1                  Link Up
```

20

Configuring Port and System Security

This chapter describes how to configure port-based security features, which control access to the network through the switch ports, and the denial of service (DoS) feature.

Port-based security includes IEEE 802.1X authentication and port MAC locking.

- IEEE 802.1X provides an authentication mechanism to devices connected to the switch. Network access is permitted only to authorized devices (clients).
- Port MAC locking is used to enable security on a per-port basis. When a port is locked, only packets with allowable source MAC addresses can be forwarded. All other packets are discarded. Port-MAC locking allows a configurable limit to the number of source MAC addresses that can be learned on a port.

NOTE: Port-based security can also be accomplished by using Access Control Lists (ACLs). For information about configuring ACLs, see "Configuring Access Control Lists" on page 501.

The topics covered in this chapter include:

- IEEE 802.1X
- Port Security (Port-MAC Locking)
- Denial of Service

IEEE 802.1X

What is IEEE 802.1X?

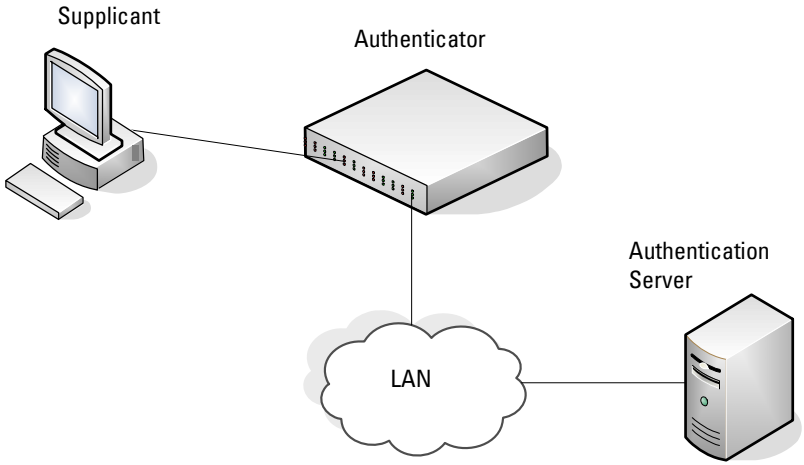
The IEEE 802.1X standard provides a means of preventing unauthorized access by supplicants (clients) to the services the switch offers, such as access to the LAN.

The 802.1X network has three components:

- **Supplicant** — The client connected to the authenticated port that requests access to the network.
- **Authenticator** — The network device that prevents network access prior to authentication.
- **Authentication Server** — The network server (such as a RADIUS server) that performs the authentication on behalf of the authenticator, and indicates whether the user is authorized to access system services.

Figure 20-1 shows the 802.1X network components.

Figure 20-1. IEEE 802.1X Network



As shown in Figure 20-1, the PowerConnect 8000-series and 8100-series switches is the authenticator and enforces the supplicant (a PC) that is attached to an 802.1X-controlled port to be authenticated by an authentication server (a RADIUS server). The result of the authentication process determines whether the supplicant is authorized to access services on that controlled port.

For a list of RADIUS attributes that the switch supports, see "Using RADIUS Servers to Control Management Access" on page 192.

What are the 802.1X Port States?

The 802.1X port state determines whether to allow or prevent network traffic on the port. A port can configured to be in one of the following 802.1X control modes:

- Auto (default)
- MAC-based
- Force-authorized
- Force-unauthorized.

These modes control the behavior of the port. The port state is either Authorized or Unauthorized.

If the port is in the authorized state, the port sends and receives normal traffic without client port-based authentication. When a port is in an unauthorized state, it ignores supplicant authentication attempts and does not provide authentication services to the client. By default, when 802.1X is globally enabled on the switch, all ports are in Auto, which means the port will be unauthorized until a successful authentication exchange has taken place.

In addition to authorized, unauthorized, and automode, the 802.1X mode of a port can be MAC based, as the following section describes.

NOTE: Only MAC-Based and Automode actually use 802.1X to authenticate. Authorized and Unauthorized modes are manual overrides.

What is MAC-Based 802.1X Authentication?

MAC-based authentication allows multiple supplicants connected to the same port to each authenticate individually. For example, a 5-port hub might be connected to a single port on the switch. Each host connected to the hub must authenticate separately in order to gain access to the network.

The hosts are distinguished by their MAC addresses.



NOTE: By default, all ports are in VLAN Access mode. A port that uses MAC-based authentication should be configured to be in General mode.

When multiple hosts (for example, a PC, a printer, and a phone in the same office) are connected to the switch on the same port, each of the connected hosts authenticates separately with the RADIUS server.

If a port uses MAC-based 802.1X authentication, the option to use MAC Authentication Bypass (MAB) is available. MAB is a supplemental authentication mechanism that allows 802.1X unaware clients – such as printers, fax machines, and some IP phones — to authenticate to the network using the client MAC address as an identifier.

The known and allowable MAC address and corresponding access rights of the client must be pre-populated in the authentication server.

When a port configured for MAB receives traffic from an unauthenticated client, the switch (Authenticator):

- Sends a EAP Request packet to the unauthenticated client
- Waits a pre-determined period of time for a response
- Retries – resends the EAP Request packet up to three times
- Considers the client to be 802.1X unaware client (if it does not receive an EAP response packet from that client)

The authenticator sends a request to the authentication server with the MAC address of the client in a hexadecimal format as the username and the MD5 hash of the MAC address as the password. The authentication server checks its database for the authorized MAC addresses and returns an Access-Accept or an Access-Reject response, depending on whether the MAC address is found in the database. MAB also allows 802.1X-unaware clients to be placed in a RADIUS-assigned VLAN or to apply a specific Filter ID to the client traffic.

NOTE: MAB initiates only after the dot1x guest VLAN period times out. If the client responds to any of the EAPOL identity requests, MAB does not initiate for that client.

What is the Role of 802.1X in VLAN Assignment?

PowerConnect 8000-series and 8100-series switches allow a port to be placed into a particular VLAN based on the result of the authentication or type of 802.1X authentication a client uses when it accesses the switch. The authentication server can provide information to the switch about which VLAN to assign the supplicant.

When a host connects to a switch that uses an authentication server to authenticate, the host authentication can typically have one of three outcomes:

- The host is authenticated.
- The host attempts to authenticate but fails because it lacks certain security credentials.
- The host is a guest and does not try to authenticate at all (802.1X unaware).

You can create three separate VLANs on the switch to handle a host depending on whether the host authenticates, fails the authentication, or is a guest. The RADIUS server informs the switch of the selected VLAN as part of the authentication.

Authenticated and Unauthenticated VLANs

Hosts that authenticate normally use a VLAN that includes access to network resources. Hosts that fail the authentication might be denied access to the network or placed on a *quarantine* VLAN with limited network access.

Much of the configuration to assign authenticated hosts to a particular VLAN takes place on the 802.1X authenticator server (for example, a RADIUS server). If you use an external RADIUS server to manage VLANs, you configure the server to use Tunnel attributes in Access-Accept messages in order to inform the switch about the selected VLAN. These attributes are defined in RFC 2868, and their use for dynamic VLAN is specified in RFC 3580.

The VLAN attributes defined in RFC3580 are as follows:

- Tunnel-Type=VLAN (13)
- Tunnel-Medium-Type=802
- Tunnel-Private-Group-ID=VLANID

VLANID is 12-bits and has a value between 1 and 4093.

Dynamic VLAN Creation

If RADIUS-assigned VLANs are enabled though the Authorization Network RADIUS configuration option, the RADIUS server is expected to include the VLAN ID in the 802.1X tunnel attributes of its response message to the switch. If dynamic VLAN creation is enabled on the switch and the RADIUS-assigned VLAN does not exist, then the assigned VLAN is dynamically created. This implies that the client can connect from any port and can get assigned to the appropriate VLAN. This gives flexibility for clients to move around the network without much additional configuration required.

Guest VLAN

The Guest VLAN feature allows a switch to provide a distinguished service to unauthenticated users. This feature provides a mechanism to allow users access to hosts on the guest VLAN. For example, a company might provide a guest VLAN to visitors and contractors to permit network access that allows visitors to connect to external network resources, such as the Internet, with no ability to browse information on the internal LAN.

In port-based 802.1X mode, when a client that does not support 802.1X is connected to an unauthorized port that is 802.1X-enabled, the client does not respond to the 802.1X requests from the switch. Therefore, the port remains in the unauthorized state, and the client is not granted access to the network. If a guest VLAN is configured for that port, then the port is placed in the configured guest VLAN and the port is moved to the authorized state, allowing access to the client. However, if the port is in MAC-based 802.1X authentication mode, it will not move to the authorized state. MAC-based mode makes it possible for both authenticated and guest clients to use the same port at the same time.

 **NOTE:** MAB and the guest VLAN feature are mutually exclusive on a port.

Client devices that are 802.1X-supplicant-enabled authenticate with the switch when they are plugged into the 802.1X-enabled switch port. The switch verifies the credentials of the client by communicating with an authentication server. If the credentials are verified, the authentication server informs the switch to *unblock* the switch port and allows the client unrestricted access to the network; i.e., the client is a member of an internal VLAN.

Guest VLAN mode can be configured on a per-port basis. If a client does not attempt authentication on a port, and the port is configured for the guest VLAN, the client is assigned to the guest VLAN configured on that port. The port is assigned a guest VLAN ID and is moved to the authorized status. When the guest VLAN is disabled, users authorized by the guest VLAN are removed.

What is Monitor Mode?

The monitor mode is a special mode that can be enabled in conjunction with 802.1X authentication. Monitor mode provides a way for network administrators to identify possible issues with the 802.1X configuration on the switch without affecting the network access to the users of the switch. It allows network access even in case where there is a failure to authenticate but logs the results of the authentication process for diagnostic purposes.

The monitor mode can be configured globally on a switch. If the switch fails to authenticate a user for any reason (for example, RADIUS access reject from RADIUS server, RADIUS timeout, or the client itself is Dot1x unaware), the client is authenticated and is undisturbed by the failure condition(s). The reasons for failure are logged for tracking purposes.

Table 20-1 provides a summary of the 802.1X Monitor Mode behavior.

Table 20-1. IEEE 802.1X Monitor Mode Behavior

Case	Sub-case	Regular Dot1x	Dot1x Monitor Mode
RADIUS/IAS Success	Success	Port State: Permit VLAN: Assigned Filter: Assigned	Port State: Permit VLAN: Assigned Filter: Assigned
	Incorrect NAS Port	Port State: Deny	Port State: Permit VLAN: Default PVID of the port

Table 20-1. IEEE 802.1X Monitor Mode Behavior (Continued)

Case	Sub-case	Regular Dot1x	Dot1x Monitor Mode
	Invalid VLAN Assignment	Port State: Deny	Port State: Permit VLAN: Default PVID of the port
	Invalid Filter-id	Port State: Deny	Port State: Permit VLAN: Default PVID of the port
	Bad RADIUS packet	Port State: Deny	Port State: Permit VLAN: Default PVID of the port
RADIUS/IAS Failure	Default behavior	Port State: Deny	Port State: Permit VLAN: Default PVID of the port
	Unauth VLAN enabled	Port State: Permit VLAN: Unauth	Port State: Permit VLAN: Unauth
RADIUS Timeout	Default behavior	Port State: Deny	Port State: Permit VLAN: Default PVID of the port
	Unauth VLAN enabled	Port State: Deny	Port State: Permit VLAN: Unauth
EAPOL Timeout	Default behavior	Port State: Deny	Port State: Permit
3 × EAPOL Timeout (Guest VLAN timer expiry or MAB timer expiry)	Guest VLAN enabled	Port State: Permit VLAN: Guest	Port State: Permit VLAN: Guest
	MAB Success Case	Port State: Permit VLAN: Assigned Filter: Assigned	Port State: Permit VLAN: Assigned Filter: Assigned
	MAB Fail Case	Port State: Deny	Port State: Permit VLAN: Default PVID of the port

Table 20-1. IEEE 802.1X Monitor Mode Behavior (Continued)

Case	Sub-case	Regular Dot1x	Dot1x Monitor Mode
Supplicant Timeout		Port State: Deny	Port State: Deny
Port/Client Authenticated on Guest VLAN	Delete Guest VLANID through Dot1Q	Port State: Deny	Port State: Permit VLAN: Default PVID of the port

How Does the Authentication Server Assign DiffServ Filters?

The PowerConnect 8000-series and 8100-series switches allow the external 802.1X Authenticator or RADIUS server to assign DiffServ policies to users that authenticate to the switch. When a host (supplicant) attempts to connect to the network through a port, the switch contacts the 802.1X authenticator or RADIUS server, which then provides information to the switch about which DiffServ policy to assign the host (supplicant). The application of the policy is applied to the host after the authentication process has completed.

For additional guidelines about using an authentication server to assign DiffServ policies, see "Configuring Authentication Server DiffServ Filter Assignments" on page 491.

What is the Internal Authentication Server?

The Internal Authentication Server (IAS) is a dedicated database for localized authentication of users for network access through 802.1X. In this database, the switch maintains a list of username and password combinations to use for 802.1X authentication. You can manually create entries in the database, or you can upload the IAS information to the switch.

If the authentication method for 802.1X is IAS, the switch uses the locally stored list of username and passwords to provide port-based authentication to users instead of using an external authentication server.

NOTE: The IAS database does not handle VLAN assignments or DiffServ policy assignments.

Default 802.1X Values

Table 20-2 lists the default values for the 802.1X features.

Table 20-2. Default Port-Based Security Values

Feature	Description
Global 802.1X status	Disabled
802.1X authentication method	none
Per-port 802.1X status	Disabled
Port state	automode
Periodic reauthentication	Disabled
Seconds between reauthentication attempts	3600
Authentication server timeout	30 seconds
Resending EAP identity Request	30 seconds
Quiet period	60 seconds
Supplicant timeout	30 seconds
Max EAP request	2 times
Maximum number of supplicants per port for MAC-based authentication mode	16
Guest VLAN	Disabled
Unauthenticated VLAN	Disabled
Dynamic VLAN creation	Disabled
RADIUS-assigned VLANs	Disabled
IAS users	none configured
Port security	Unlocked
Port security traps	Disabled
Maximum learned MAC addresses	100 (when locked)
Monitor mode	Disabled

Configuring IEEE 802.1X (Web)

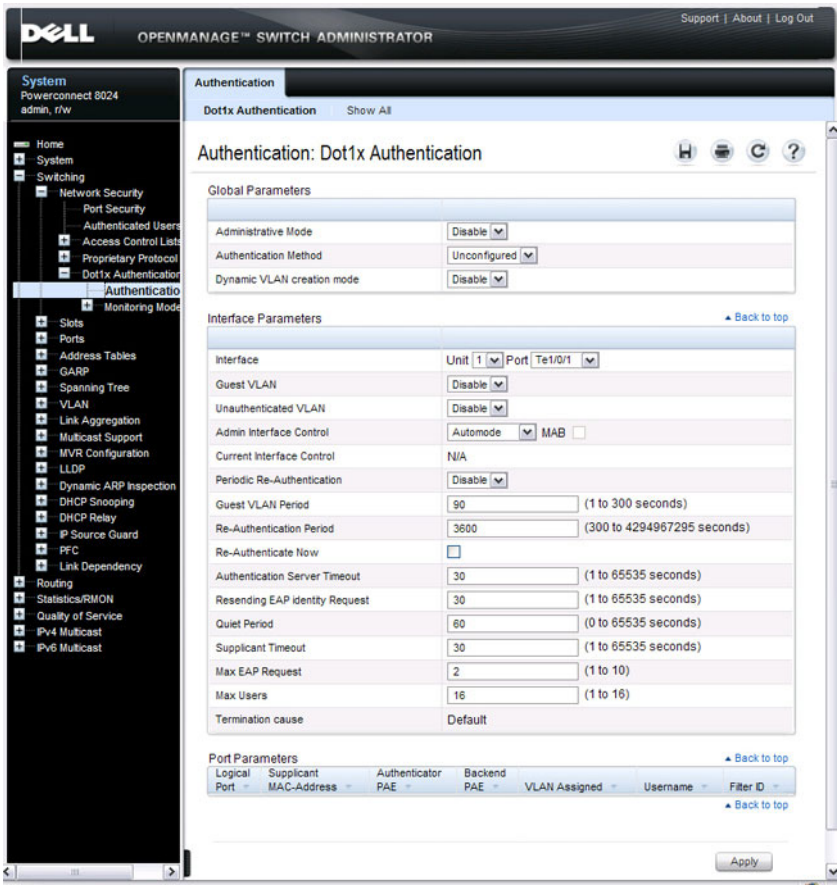
This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring the IEEE 802.1X features and Port Security on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click  at the top of the page.

Dot1x Authentication

Use the **Dot1x Authentication** page to configure the 802.1X administrative mode on the switch and to configure general 802.1X parameters for a port.

To display the **Dot1x Authentication** page, click **Switching** → **Network Security** → **Dot1x Authentication** → **Authentication** in the navigation panel.

Figure 20-2. Dot1x Authentication



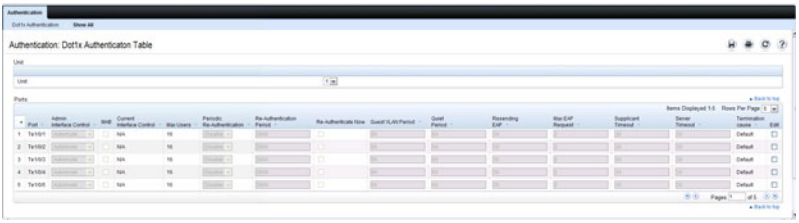
Configuring 802.1X Settings on Multiple Ports

To configure 802.1X authentication on multiple ports:

1. Open the **Dot1x Authentication** page.
2. Click **Show All** to display the **Dot1x Authentication Table** page.
3. In the Ports list, select the check box in the **Edit** column for the port to configure.

4 Select the desired settings to change for all ports that are selected for editing.

Figure 20-3. Configure Dot1x Settings



5 Click **Apply**.

Re-Authenticating One Port

To reauthenticate a port:

1 Open the **Dot1x Authentication** page.

2 Click **Show All**.

The **Dot1x Authentication Table** displays.

3 Check **Edit** to select the Unit/Port to re-authenticate.

4 Check **Reauthenticate Now**.

5 Click **Apply**.

The authentication process is restarted on the specified port.

Re-Authenticating Multiple Ports in the Dot1x Authentication Table

To reauthenticate multiple ports:

1 Open the **Dot1x Authentication** page.

2 Click **Show All**.

The **Dot1x Authentication Table** displays.

3 Check **Edit** to select the Units/Ports to re-authenticate.

4 To re-authenticate on a periodic basis, set **Periodic Re-Authentication** to Enable, and specify a **Re-Authentication Period** for all desired ports.

5 To re-authenticate immediately, check **Reauthenticate Now** for all ports to be re-authenticated.

6 Click **Apply**.

The authentication process is restarted on the specified ports (either immediately or periodically).

Changing Administrative Port Control

To change the administrative port control:

1 Open the **Dot1x Authentication** page.

2 Click **Show All**.

The **Dot1x Authentication Table** displays.

3 Scroll to the right side of the table and select the **Edit** check box for each port to configure. Change **Admin Port Control** to **Authorized**, **Unauthorized**, or **Automode** as needed for chosen ports. Only **MAC-Based** and **Automode** actually use 802.1X to authenticate. **Authorized** and **Unauthorized** are manual overrides.

4 Click **Apply**.

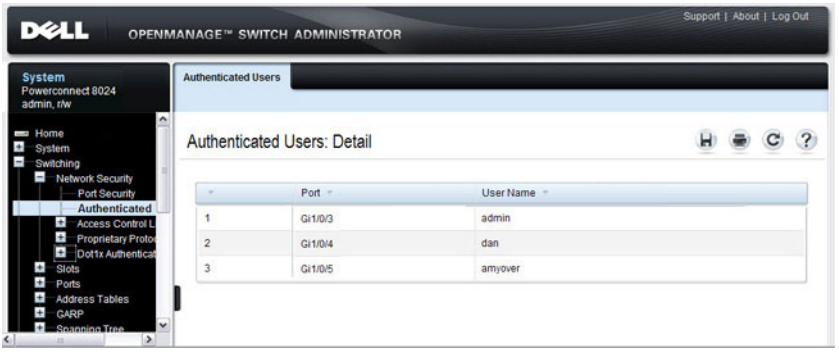
Admin Port Control is updated for the specified ports, and the device is updated.

Authenticated Users

The **Authenticated Users** page is used to display lists of ports that have authenticated users.

To display the **Authenticated Users** page, click **Switching** → **Network Security** → **Authenticated Users** in the navigation panel.

Figure 20-4. Network Security Authenticated Users



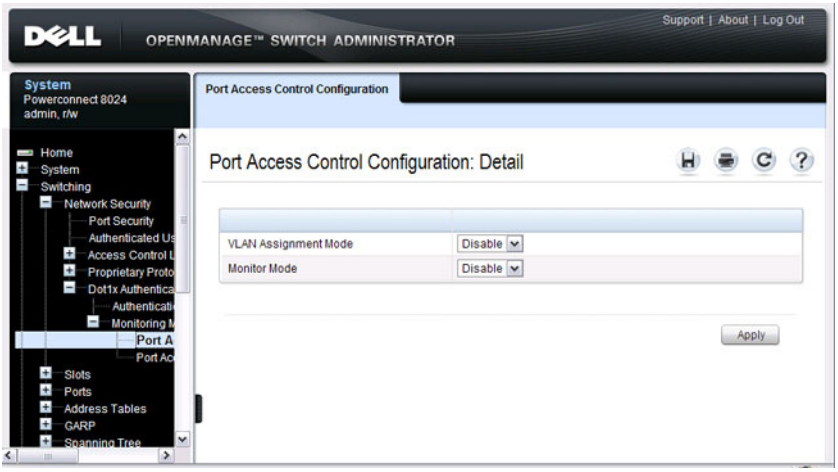
Port Access Control Configuration

Use the **Port Access Control Configuration** page to globally enable or disable RADIUS-assigned VLANs and to enable Monitor Mode to help troubleshoot 802.1X configuration issues.

NOTE: The VLAN Assignment Mode field is the same as the Admin Mode field on the **System** → **Management Security** → **Authorization Network RADIUS** page.

To display the **Port Access Control Configuration** page, click **Switching** → **Network Security** → **Dot1x Authentication** → **Monitor Mode** → **Port Access Control Configuration** in the navigation panel.

Figure 20-5. Port Access Control Configuration

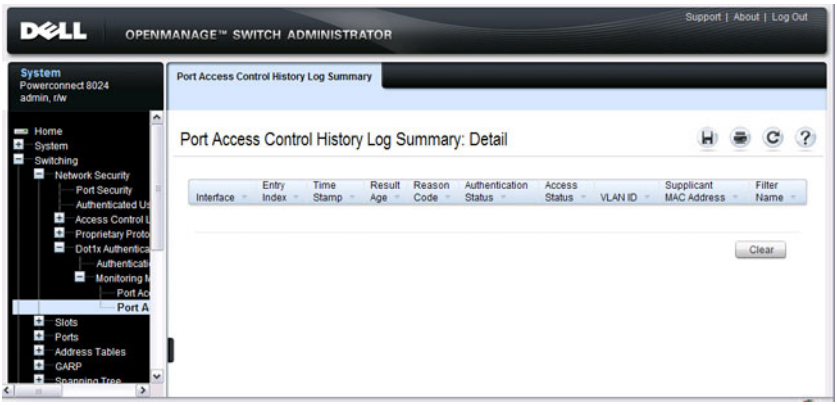


Port Access Control History Log Summary

Use the **Port Access Control History Log Summary** page to view log messages about 802.1X client authentication attempts. The information on this page can help you troubleshoot 802.1X configuration issues.

To display the **Port Access Control History Log Summary** page, click **Port Access Control Configuration** page, click **Switching** → **Network Security** → **Dot1x Authentication** → **Monitor Mode** → **Port Access Control History Log Summary** in the navigation panel.

Figure 20-6. Port Access Control History Log Summary

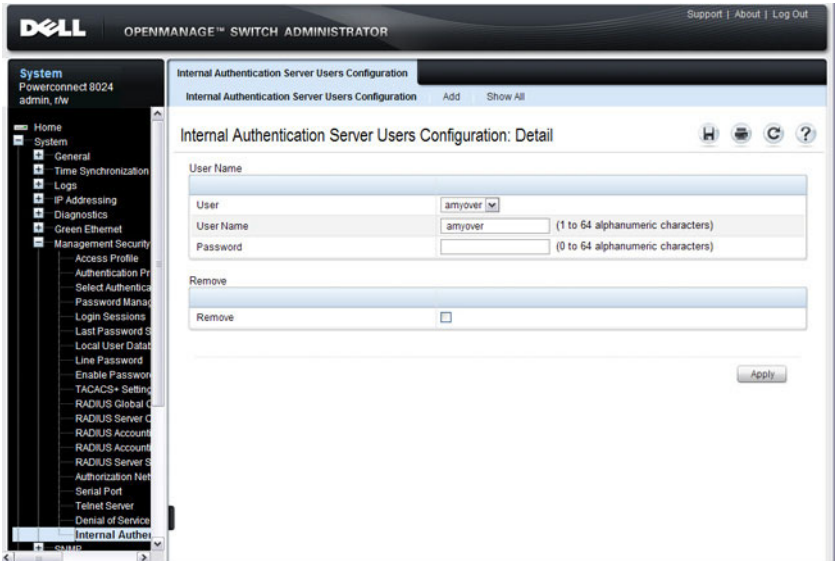


Internal Authentication Server Users Configuration

Use the **Internal Authentication Server Users Configuration** page to add users to the local IAS database and to view the database entries.

To display the **Internal Authentication Server Users Configuration** page, click **System** → **Management Security** → **Internal Authentication Server Users Configuration** in the navigation panel.

Figure 20-7. Internal Authentication Server Users Configuration



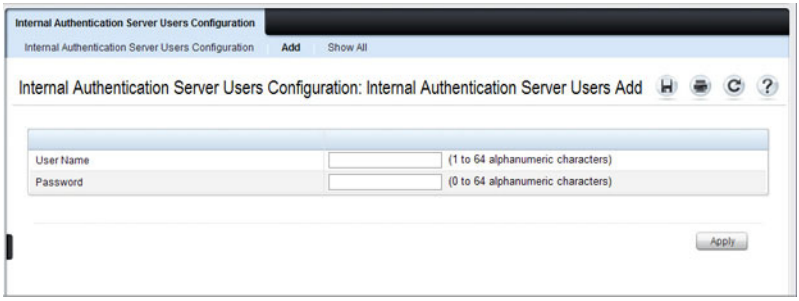
NOTE: If no users exist in the IAS database, the IAS Users Configuration Page does not display the fields shown in the image.

Adding Users to the IAS Database

To add IAS users:

1. Open the **Internal Authentication Server Users Configuration** page.
2. Click **Add** to display the **Internal Authentication Server Users Add** page.
3. Specify a username and password in the appropriate fields.

Figure 20-8. Adding an IAS User



4 Click **Apply**.

To view the Internal Authentication Server Users Table page, click **Show All**.

Removing an IAS User

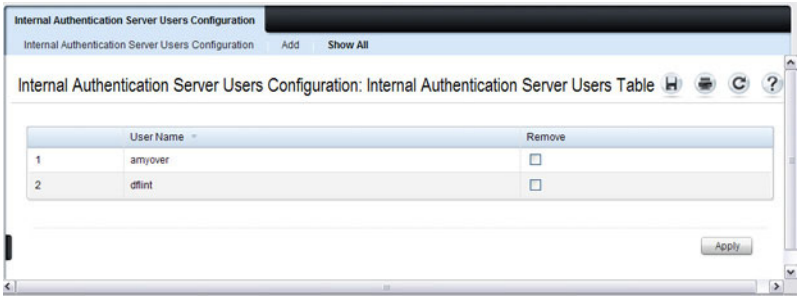
To delete an IAS user:

1 Open the **Internal Authentication Server Users Configuration** page.

2 From the User menu, select the user to remove, select the user to remove.

3 Select the **Remove** check box.

Figure 20-9. Removing an IAS User



4 Click **Apply**.

Configuring IEEE 802.1X (CLI)

This section provides information about commands you use to configure 802.1X and Port Security settings. For additional information about the commands in this section, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Configuring Basic 802.1X Authentication Settings

Beginning in Privileged EXEC mode, use the following commands to enable and configure 802.1X authentication on the switch.

Command	Purpose
configure	Enter Global Configuration mode.
aaa accounting dot1x default	Sets 802.1X accounting to the default operational mode
aaa authentication dot1x default *method1*	Specify the authentication method to use to authenticate 802.1X clients that connect to the switch. *method1*—The method keyword can be **radius**, **none, or ias**.
dot1x system-auth-control	Globally enable 802.1X authentication on the switch.
interface *interface*	Enter interface configuration mode for the specified interface. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**. You can also specify a range of interfaces with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12.

Command	Purpose
dot1x port-control {force-authorized | force-unauthorized | auto | mac-based}	Specify the 802.1X mode for the port. **NOTE:** For standard 802.1X implementations in which one client is connected to one port, use the **dot1x port-control auto** command to enable 802.1X authentication on the port. • **auto** — Enables 802.1X authentication on the interface and causes the port to transition to the authorized or unauthorized state based on the 802.1X authentication exchange between the switch and the client. • **force-authorized** — Disables 802.1X authentication on the interface and causes the port to transition to the authorized state without any authentication exchange required. The port sends and receives normal traffic without 802.1X-based authentication of the client. • **force-unauthorized** — Denies all access through this interface by forcing the port to transition to the unauthorized state, ignoring all attempts by the client to authenticate. The switch cannot provide authentication services to the client through the interface. • **mac-based** — Enables 802.1X authentication on the interface and allows multiple hosts to authenticate on a single port. The hosts are distinguished by their MAC addresses.
dot1x mac-auth-bypass	If the 802.1X mode on the interface is **mac-based**, you can optionally use this command to enable MAB on an interface.
CTRL + Z	Exit to Privileged EXEC mode.
show dot1x	View the current 802.1X configuration.
show dot1x clients {all | *interface*}	View information about 802.1X clients that have successfully authenticated and are connected to the switch. The *interface* variable includes the interface type and number.
show dot1x users [**username** *username*]	View the 802.1X authenticated users for the switch.

NOTE: To enable 802.1X Monitor Mode to help troubleshoot authentication issues, use the **dot1x system-auth-control monitor** command in Global Configuration mode. To view 802.1X authentication events and information, use the **show dot1x authentication-history** {<*interface*> | **all**} [**failed-auth-only**] [**detail**] command in Privileged EXEC mode. To clear the history, use the **clear dot1x authentication-history** command.

Configuring Additional 802.1X Interface Settings

Beginning in Privileged EXEC mode, use the following commands to configure 802.1X interface settings such as the reauthentication period and switch-to-client retransmission time.

Command	Purpose
configure	Enter Global Configuration mode.
interface *interface*	Enter interface configuration mode for the specified interface. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**.
	You can also specify a range of interfaces with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12.
dot1x reauthentication	Enable periodic re-authentication of the client.
dot1x timeout re-authperiod *seconds*	Set the number of seconds between re-authentication attempts.
dot1x timeout server-timeout *seconds*	Set the time that the switch waits for a response from the authentication server.
dot1x timeout tx-period *seconds*	Set the number of seconds that the switch waits for a response to an Extensible Authentication Protocol (EAP)-request/identity frame from the client before resending the request.
dot1x timeout quiet-period *seconds*	Set the number of seconds that the switch remains in the quiet state following a failed authentication exchange (for example, the client provided an invalid password).

Command	Purpose
dot1x timeout supp-timeout *seconds*	Set the time that the switch waits for a response before retransmitting an Extensible Authentication Protocol (EAP)-request frame to the client.
dot1x max-req *count*	Set the maximum number of times that the switch sends an Extensible Authentication Protocol (EAP)-request frame (assuming that no response is received) to the client before restarting the authentication process.
dot1x max-users *users*	Set the maximum number of clients supported on the port when MAC-based 802.1X authentication is enabled on the port.
CTRL + Z	Exit to Privileged EXEC mode.
dot1x re-authenticate [*interface*]	Manually initiate the re-authentication of all 802.1X-enabled ports or on the specified 802.1X-enabled port. The *interface* variable includes the interface type and number.
dot1x initialize [*interface*]	Start the initialization sequence on all ports or on the specified port. **NOTE:** This command is valid only if the port-control mode for the specified port is auto or MAC-based.
show dot1x [interface *interface]*	View 802.1X settings for the switch or for the specified interface.
show dot1x interface *interface* **statistics**	View 802.1X statistics for the specified interface.

Configuring 802.1X Settings for RADIUS-Assigned VLANs

Beginning in Privileged EXEC mode, use the following commands to configure 802.1X settings that affect the RADIUS-assigned VLAN.

Command	Purpose
configure	Enter Global Configuration mode.
aaa authorization network default radius	Allow the RADIUS server to assign VLAN IDs to clients.

Command	Purpose
dot1x dynamic-vlan enable	If the RADIUS assigned VLAN does not exist on the switch, allow the switch to dynamically create the assigned VLAN.
interface *interface*	Enter interface configuration mode for the specified interface. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**. You can also specify a range of interfaces with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12.
dot1x guest-vlan *vlan-id*	Specify the guest VLAN.
dot1x unauth-vlan *vlan-id*	Specify the unauthenticated VLAN. The VLAN must already have been created.
CTRL + Z	Exit to Privileged EXEC mode.
show dot1x advanced *interface*	View the current 802.1X configuration.

NOTE: When dynamically creating VLANs, the uplink port should be in trunk mode so that it will automatically participate in all dynamically-created VLANs. Otherwise, the supplicant may be placed in a VLAN that does not go beyond the switch because no other ports are participating.

Configuring Internal Authentication Server Users

Beginning in Privileged EXEC mode, use the following commands to add users to the IAS database and to use the database for 802.1X authentication.

Command	Purpose
configure	Enter Global Configuration mode.
aaa ias-user username *user*	Add a user to the IAS user database. This command also changes the mode to the AAA User Config mode.
password *password* [**encrypted**]	Configure the password associated with the user.
CTRL + Z	Exit to Privileged EXEC mode.
show aaa ias-users	View all configured IAS users.
clear aaa ias-users	Delete all IAS users from the database.

IEEE 802.1X Configuration Examples

This section contains the following examples:

- Configuring 802.1X Authentication
- Controlling Authentication-Based VLAN Assignment
- Allowing Dynamic VLAN Creation of RADIUS-Assigned VLANs
- Configuring Authentication Server DiffServ Filter Assignments

Configuring 802.1X Authentication

The network in this example requires clients to use 802.1X authentication to access the network through the switch ports. The administrator must configure the following settings on systems other than the switch before configuring the switch:

1. Add the users to the client database on the Authentication Server, such as a RADIUS server with Cisco® Secure Access Control Server (ACS) software.
2. Configure the settings on the client, such a PC running Microsoft® Windows, to require 802.1X authentication.

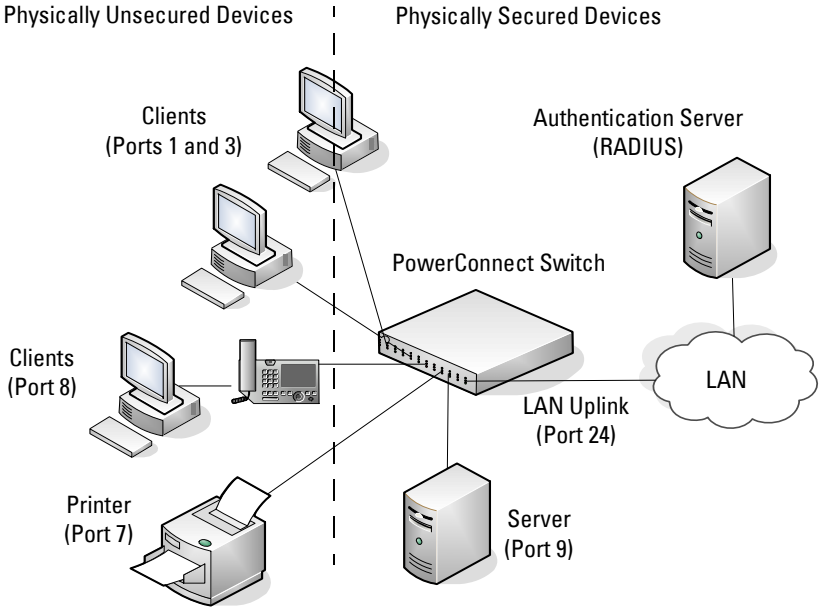
The switch uses the Authentication Server with an IP address of 10.10.10.10 to authenticate clients. Port 7 is connected to a printer in the unsecured area. The printer is an 802.1X unaware client, so Port 7 is configured to use MAC-based authentication with MAB.

NOTE: The printer requires an entry in the client database that uses the printer MAC address as the username.

An IP phone is directly connected to Port 8, and a PC is connected to the IP phone. Both devices are authenticated through MAC-based authentication, which allows multiple hosts to authenticate on a single port. The hosts are distinguished by their MAC addresses, and hosts authenticate separately with the RADIUS server.

Port 9 is connected to a server in a part of the network that has secure physical access (i.e. the doors to the wiring closet and data center are locked), so this port is set to the Authorized state, meaning that the device connected to this port does not need to authenticate using 802.1X. Port 24 is the uplink to a router and is also in the Authorized state.

Figure 20-10. 802.1X Example



The following example shows how to configure the example shown in Figure 20-10.

1. Configure the RADIUS server IP address and shared secret (*secret*).

```
console#configure
console(config)#radius-server host 10.10.10.10
console(Config-radius)#exit
console(config)#radius-server key secret
console(config)#exit
```

2. Enable 802.1X port-based access control on the switch.

```
console(config)#dot1x system-auth-control
```

3. Configure ports 9 and 24 to be in the Authorized state, which allows the devices to connect to these ports to access the switch services without authentication.

```
console(config)#interface range te1/0/9,te1/0/24
```

```
console(config-if)#dot1x port-control force-
authorized
console(config-if)#exit
```

4 Configure Port 7 to require MAC-based authentication with MAB.

```
console(config)#interface te1/0/7
console(config-if-Te1/0/7)#dot1x port-control mac-
based
console(config-if-Te1/0/7)#dot1x mac-auth-bypass
```

5 Set the port to an 802.1Q VLAN. The port must be in general mode in order to enable MAC-based 802.1X authentication.

```
console(config-if-Te1/0/7)#switchport mode general
console(config-if-Te1/0/7)#exit
```

6 Enable MAC-based authentication on port 8 and limit the number of devices that can authenticate on that port to 2.

```
console(config)#interface te1/0/8
console(config-if-Te1/0/8)#dot1x port-control mac-
based
console(config-if-Te1/0/8)#dot1x max-users 2
```

7 Set Port 8 to switchport mode general. The port must be in general mode in order to enable MAC-based 802.1X authentication.

```
console(config-if-Te1/0/8)#switchport mode general
console(config-if-Te1/0/8)#exit
console(config)#exit
```

8 View the client connection status.

When the clients on Ports 1, 3, and 7(supplicants), attempt to communicate via the switch, the switch challenges the supplicants for 802.1X credentials. The switch encrypts the provided information and transmits it to the RADIUS server. If the RADIUS server grants access, the system sets the 802.1X port state of the interface to authorized and the supplicants are able to access network resources.

```
console#show dot1x clients all

Interface..................................... Te1/0/1
User Name..................................... aoversmit
Supp MAC Address.............................. 0012.1753.031A
Session Time.................................. 756
```

```
Filter Id.....................................
VLAN Assigned................................ 1 (Default)

Interface..................................... Te1/0/3
User Name..................................... dflint
Supp MAC Address.............................. 0004.5A55.EFAD
Session Time.................................. 826
Filter Id.....................................
VLAN Assigned................................ 1 (Default)

Interface..................................... Te1/0/7
User Name..................................... 0006.6B33.06BA
Supp MAC Address.............................. 0006.6B33.06BA
Session Time.................................. 826
Filter Id.....................................
VLAN Assigned................................ 1 (Default)
```

9 View a summary of the port status.

```
console#show dot1x
Administrative Mode............... Enabled

Port     Admin                Oper         Reauth    Reauth
         Mode                 Mode         Control   Period
-------  ---------------- ------------  --------  ----------
Te1/0/1  auto             Authorized     FALSE     3600
Te1/0/2  auto             N/A            FALSE     3600
Te1/0/3  auto             Authorized     FALSE     3600
Te1/0/4  auto             N/A            FALSE     3600
Te1/0/5  auto             N/A            FALSE     3600
Te1/0/6  auto             N/A            FALSE     3600
Te1/0/7  mac-based        Authorized     FALSE     3600
Te1/0/8  mac-based        N/A            FALSE     3600
Te1/0/9  force-authorized Authorized     FALSE     3600
Te1/0/10 force-authorized Authorized     FALSE     3600
Te1/0/11 auto             N/A            FALSE     3600
--More-- or (q)uit
```

10 View 802.1X information about Port 8.

```
console#show dot1x interface te1/0/8

Administrative Mode.............. Enabled
Dynamic VLAN Creation Mode........ Enabled
Monitor Mode..................... Disabled

Port     Admin               Oper         Reauth    Reauth
         Mode                Mode         Control   Period
-------  ---------------- ------------  --------  ----------
Te1/0/8  mac-based           Authorized   FALSE     3600

Quiet Period.................................. 60
Transmit Period............................... 30
Maximum Requests.............................. 2
Max Users..................................... 2
VLAN Assigned................................. 1 (Default)
Supplicant Timeout............................ 30
Guest-vlan Timeout............................ 90
Server Timeout (secs)......................... 30
MAB mode (configured)......................... Disabled
MAB mode (operational)........................ Disabled
```

Controlling Authentication-Based VLAN Assignment

The network in this example uses three VLANs to control access to network resources. When a client connects to the network, it is assigned to a particular VLAN based on one of the following events:

- It attempts to contact the 802.1X server and is authenticated.
- It attempts to contact the 802.1X server and fails to authenticate.
- It does not attempt to contact the 802.1X server.

The following table describes the three VLANs:

VLAN ID	VLAN Name	VLAN Purpose
100	Authorized	Data from authorized clients
200	Unauthorized	Data traffic from clients that fail the authentication with the RADIUS server
300	Guest	Data traffic from clients that do not attempt to authenticate with the RADIUS server

NOTE: Dynamic VLAN creation applies only to authorized ports. The VLANs for unauthorized and guest users must be configured on the switch and cannot be dynamically created based on RADIUS-based VLAN assignment.

The commands in this example show how to configure the switch to control VLAN assignment for the example network. This example also contains commands to configure the uplink, or trunk, port (a port connected to a router or the internal network), and to configure the downlink, or access, ports (ports connected to one or more hosts). Ports 1–23 are downstream ports. Port 24 is an uplink port. An external RADIUS server handles the VLAN assignment.

NOTE: The configuration to control the VLAN assignment for authorized users is done on the external RADIUS server.

To configure the switch:

1. Create the VLANs and configure the VLAN names.

```
console(config)#vlan 100
console(config-vlan100)#name Authorized
console(config-vlan100)#exit

console(config)#vlan 200
console(config-vlan200)#name Unauthorized
console(config-vlan200)#exit

console(config)#vlan 300
console(config-vlan300)#name Guest
console(config-vlan300)#exit
```

2. Configure information about the external RADIUS server the switch uses to authenticate clients. The RADIUS server IP address is 10.10.10.10, and the shared secret is qwerty123.

```
console(config)#radius-server key qwerty123
console(config)#radius-server host 10.10.10.10
console(Config-auth-radius)#exit
```

3. Enable 802.1X on the switch.

```
console(config)#dot1x system-auth-control
```

4. Create a default authentication login list and use the RADIUS server for port-based authentication for connected clients.

```
console(config)#aaa authentication dot1x default
radius
```

5. Allow the switch to accept VLAN assignments by the RADIUS server.

```
console(config)#aaa authorization network default
radius
```

6. Enter interface configuration mode for the downlink ports.

```
console(config)#interface range te1/0/1-23
```

7. Set the downlink ports to the access mode because each downlink port connects to a single host that belongs to a single VLAN.

```
console(config-if)#switchport mode access
```

8 Enable periodic reauthentication of the client on the ports and set the number of seconds to wait between reauthentication attempts to 300 seconds. Reauthentication is enabled to increase security. If the client information is removed from the RADIUS server after it has been authenticated, the client will be denied access when it attempts to reauthenticate.

```
console(config-if)#dot1x reauthentication
console(config-if)#dot1x timeout re-authperiod 300
```

9 Set the unauthenticated VLAN on the ports to VLAN 200 so that any client that connects to one of the ports and fails the 802.1X authentication is placed in VLAN 200.

```
console(config-if)#dot1x unauth-vlan 200
```

10 Set the guest VLAN on the ports to VLAN 300. This command automatically enables the Guest VLAN Mode on the downlink ports. Any client that connects to the port and does not attempt to authenticate is placed on the guest VLAN.

```
console(config-if)#dot1x guest-vlan 300
console(config-if)#exit
```

11 Enter Interface Configuration mode for port 24, the uplink (trunk) port.

```
console(config)#interface te1/0/24
```

12 Disable 802.1X authentication on the interface. This causes the port to transition to the authorized state without any authentication exchange required. This port does not connect to any end-users, so there is no need for 802.1X-based authentication.

```
console(config-if-Te1/0/24)#dot1x port-control
force-authorized
```

13 Set the uplink port to trunk mode so that it accepts tagged traffic and transmits it to the connected device (another switch or router).

```
console(config-if-Te1/0/24)#switchport mode trunk
```

Allowing Dynamic VLAN Creation of RADIUS-Assigned VLANs

The network in this example uses a RADIUS server to provide VLAN assignments to host that connect to the switch. In this example, the VLANs are not configured on the switch. Instead, the switch is configured to allow the dynamic creation of VLANs when a RADIUS-assigned VLAN does not already exist on the switch.

In this example, Ports 1–23 are configured as downlink, or access, ports, and Port 24 is the trunk port. As a trunk port, Port 24 is automatically added as a member to all VLANs that are statically or dynamically configured on the switch. However, the network administrator in this example has determined that traffic in VLANs 1000–2000 should not be forwarded on the trunk port, even if the RADIUS server assigns a connected host to a VLAN in this range, and the switch dynamically creates the VLAN.



NOTE: The configuration to control the VLAN assignment for hosts is done on the external RADIUS server.

To configure the switch:

1. Configure information about the external RADIUS server the switch uses to authenticate clients. The RADIUS server IP address is 10.10.10.10, and the shared secret is qwerty123.

   ```
   console(config)#radius-server key qwerty123
   console(config)#radius-server host 10.10.10.10
   console(Config-auth-radius)#exit
   ```

2. Enable 802.1X on the switch.

   ```
   console(config)#dot1x system-auth-control
   ```

3. Create a default authentication login list and use the RADIUS server for port-based authentication for connected clients.

   ```
   console(config)#aaa authentication dot1x default
   radius
   ```

4. Allow the switch to accept VLAN assignments by the RADIUS server.

   ```
   console(config)#aaa authorization network default
   radius
   ```

5 Allow the switch to dynamically create VLANs when a RADIUS-assigned VLAN does not exist on the switch.

```
console(config)#dot1x dynamic-vlan enable
```

6 Enter interface configuration mode for the downlink ports.

```
console(config)#interface range te1/0/1-23
```

7 Set the downlink ports to the access mode because each downlink port connects to a single host that belongs to a single VLAN.

```
console(config-if)#switchport mode access
console(config-if)#exit
```

8 Enter Interface Configuration mode for port 24, the uplink (trunk) port.

```
console(config)#interface te1/0/24
```

9 Disable 802.1X authentication on the interface. This causes the port to transition to the authorized state without any authentication exchange required. This port does not connect to any end-users, so there is no need for 802.1X-based authentication.

```
console(config-if-Te1/0/24)#dot1x port-control
force-authorized
```

10 Set the uplink port to trunk mode so that it accepts tagged traffic and transmits it to the connected device (another switch or router).

```
console(config-if-Te1/0/24)#switchport mode trunk
```

11 Forbid the trunk from forwarding traffic that has VLAN tags for any VLAN from 1000–2000, inclusive.

```
console(config-if-Gi1/0/24)#switchport trunk
allowed vlan remove 1000-2000
console(config-if-Gi1/0/24)#exit
```

Configuring Authentication Server DiffServ Filter Assignments

To enable DiffServ filter assignment by an external server, the following conditions must be true:

- The port that the host is connected to must be enabled for MAC-based port access control by using the following command in Interface Config mode:

  ```
  dot1x port-control mac-based
  ```

- The RADIUS or 802.1X server must specify the policy to assign.

  For example, if the DiffServ policy to assign is named internet_access, include the following attribute in the RADIUS or 802.1X server configuration:

  Filter-id = “internet_access”

- The DiffServ policy specified in the attribute must already be configured on the switch, and the policy names must be identical.

  For information about configuring a DiffServ policy, see "DiffServ Configuration Examples" on page 1126. The example "Providing Subnets Equal Access to External Network" on page 1126, describes how to configure a policy named internet_access.

If you use an authentication server to assign DiffServ policies to an authenticated user, note the following guidelines:

- If the policy specified within the server attribute does not exist on the switch, authentication will fail.
- Do not delete policies used as the filter ID in the RADIUS server while 802.1X is enabled.
- Do not use the DiffServ **service-policy** command to apply the filter to an interface if you configure the RADIUS server or 802.1X authenticator to assign the DiffServ filter.

In the following example, Company XYZ uses IEEE 802.1X to authenticate all users. Contractors and temporary employees at Company XYZ are not permitted to have access to SSH ports, and data rates for Web traffic is limited. When a contractor is authenticated by the RADIUS server, the server assigns a DiffServ policy to control the traffic restrictions.

The network administrator configures two DiffServ classes: *cl-ssh* and *cl-http*. The class *cl-ssh* matches all incoming SSH packets. The class *cl-http* matches all incoming HTTP packets. Then, the administrator configures a traffic policy called *con-pol* and adds the *cl-ssh* and *cl-http*. The policy is configured so that that SSH packets are to be dropped, and HTTP data rates are limited to 1 MB with a burst size of 64 Kbps. HTTP traffic that exceeds the limit is dropped. The host ports, ports 1–23, are configured to use MAC-based dot1X authentication to allow the DiffServ policy to be applied. Finally, the administrator configures the RADIUS server with the attribute Filter-id = “con-pol”.

To configure the switch :

1. Configure the DiffServ traffic class that matches SSH traffic.

```
console#configure
console(config)#class-map match-all cl-ssh
console(config-classmap)#match srcl4port 23
console(config-classmap)#exit
```

2. Configure the DiffServ traffic class that matches HTTP traffic.

```
console(config)#class-map match-all cl-http
console(config-classmap)#match srcl4port 80
console(config-classmap)#exit
```

3. Configure the DiffServ policy.

```
console(config)#policy-map con-pol in
console(config-policy-map)#class cl-ssh
console(config-policy-classmap)#drop
console(config-policy-classmap)#exit
console(config-policy-map)#class cl-http
console(config-policy-classmap)#police-simple
1000000 64 conform-action transmit violate-action
drop
console(config-policy-classmap)#exit
console(config-policy-map)#exit
```

4. Enable DiffServ on the switch.

```
console(config)#diffserv
```

5. Configure information about the external RADIUS server the switch uses to authenticate clients. The RADIUS server IP address is 10.10.10.10, and the shared secret is qwerty123.

```
console(config)#radius-server key qwerty123
console(config)#radius-server host 10.10.10.10
console(Config-auth-radius)#exit
```

6. Enable 802.1X on the switch.

```
console(config)#dot1x system-auth-control
```

7. Create a default authentication login list and use the RADIUS server for port-based authentication for connected clients.

```
console(config)#aaa authentication dot1x default
radius
```

8 Enter Interface Configuration mode for ports 1–23 and enable MAC-based authentication.

```
console(config)#interface range te1/0/1-23
console(config-if)#dot1x port-control mac-based
```

9 Set the ports to an 802.1Q VLAN. The ports must be in general mode in order to enable MAC-based 802.1X authentication.

```
console(config-if)#switchport mode general
console(config-if)#exit
console(config)#exit
```

Port Security (Port-MAC Locking)

The Port Security feature allows you to limit the number of source MAC addresses that can be learned on a port. If a port reaches the configured limit, any other addresses beyond that limit are not learned and the frames are discarded. Frames with a source MAC address that has already been learned will be forwarded.

The purpose of this feature, which is also known as port-MAC locking, is to help secure the network by preventing unknown devices from forwarding packets into the network. For example, to ensure that only a single device can be active on a port, you can set the number of allowable dynamic addresses to one. After the MAC address of the first device is learned, no other devices will be allowed to forward frames into the network.

When link goes down on a port, all of the dynamically locked addresses are cleared from the source MAC address table the feature maintains. When the link is restored, that port can once again learn addresses up to the specified limit.

The port can learn MAC addresses dynamically, and you can manually specify a list of static MAC addresses for a port.

Default 802.1X Values

Table 20-2 lists the default values for the Port Security feature.

Table 20-3. Default Port Security Values

Feature	Description
Port security	Unlocked
Port security traps	Disabled
Maximum learned MAC addresses	100 (when locked)
Monitor mode	Disabled

Configuring Port Security Configuration (Web)

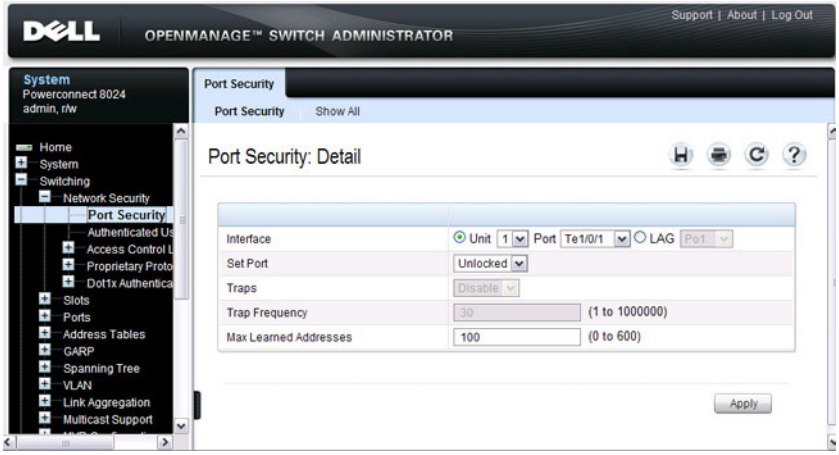
This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring the IEEE 802.1X features and Port Security on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

Port Security

Use the **Port Security** page to enable MAC locking on a per-port basis. When a port is locked, you can limit the number of source MAC addresses that are allowed to transmit traffic on the port.

To display the **Port Security** page, click **Switching** → **Network Security** → **Port Security** in the navigation panel.

Figure 20-11. Network Security Port Security



Configuring Port Security Settings on Multiple Ports

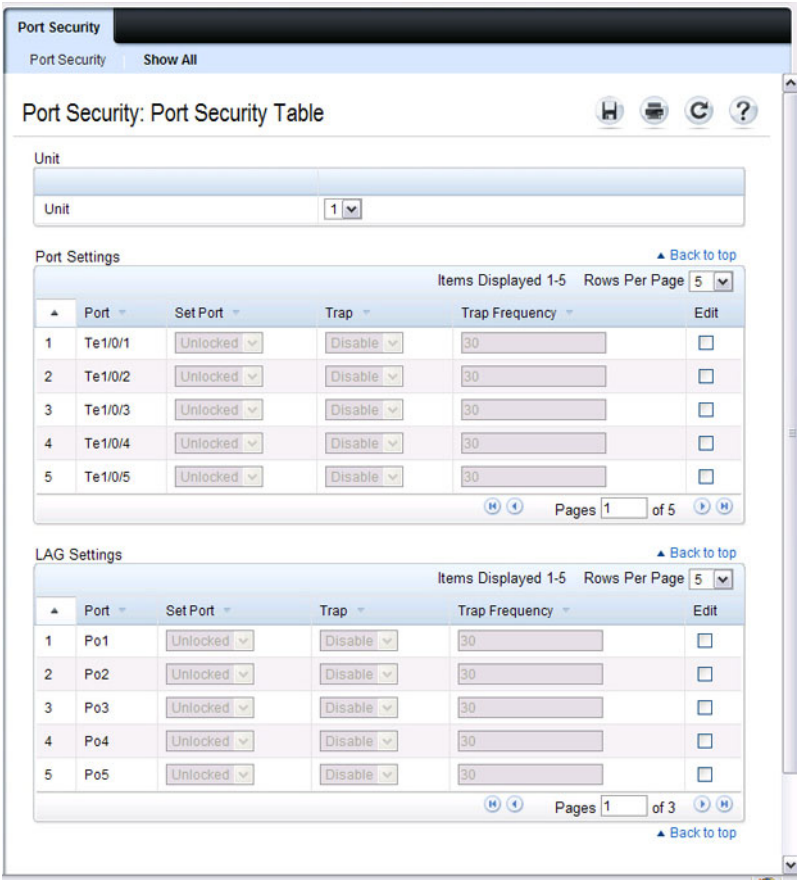
To configure port security on multiple ports:

1. Open the **Port Security** page.
2. Click **Show All** to display the **Port Security Table** page.

3 In the Ports list, select the check box in the **Edit** column for the port to configure.

4 Select the desired settings for all ports that are selected for editing.

Figure 20-12. Configure Port Security Settings



5 Click **Apply**.

Configuring Port Security (CLI)

Beginning in Privileged EXEC mode, use the following commands to enable port security on an interface to limit the number of source MAC addresses that can be learned.

Command	Purpose
configure	Enter Global Configuration mode.
interface *interface*	Enter interface configuration mode for the specified interface. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**. You can also specify a range of interfaces with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12.
port security [**discard**] [**trap** *seconds*]	Enable port security on the port. This prevents the switch from learning new addresses on this port after the maximum number of addresses has been learned. • **discard** — Discards frames with unlearned source addresses. This is the default if no option is indicated. • **trap** *seconds* — Sends SNMP traps and defines the minimal amount of time in seconds between two consecutive traps. (Range: 1–1000000)
port security max *max-addr*	Set the maximum number of MAC addresses that can be learned on the port while port security is enabled.
CTRL + Z	Exit to Privileged EXEC mode.
show ports security [*interface*]	View port security settings on all interfaces or the specified interface.
show ports security addresses [*interface*]	View the current MAC addresses that have been learned on all ports or the specified port.

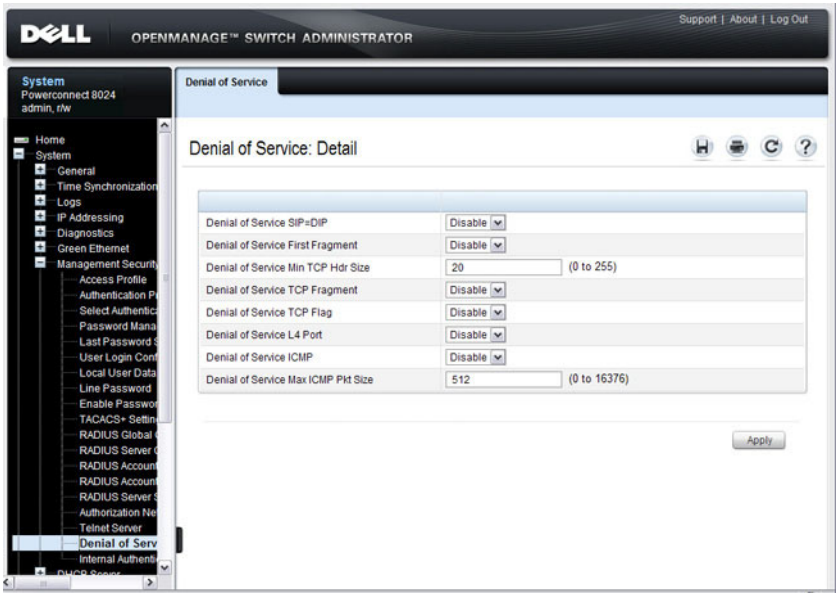
Denial of Service

Denial of Service (DoS) refers to the exploitation of a variety of vulnerabilities which would interrupt the service of a host or make a network unstable. Use the **Denial of Service** page to configure settings to help prevent DoS attacks.

DoS protection is disabled by default.

To display the **Denial of Service** page, click **System** → **Management Security** → **Denial of Service** in the navigation panel.

Figure 20-13. Denial of Service



21

Configuring Access Control Lists

This chapter describes how to configure Access Control Lists (ACLs), including IPv4, IPv6, and MAC ACLs. This chapter also describes how to configure time ranges that can be applied to any of the ACL types.

The topics covered in this chapter include:

- ACL Overview
- Configuring ACLs (Web)
- Configuring ACLs (CLI)
- ACL Configuration Examples

ACL Overview

Access Control Lists (ACLs) are a collection of permit and deny conditions, called rules, that provide security by blocking unauthorized users and allowing authorized users to access specific resources.

ACLs can also provide traffic flow control, restrict contents of routing updates, and decide which types of traffic are forwarded or blocked. ACLs can reside in a firewall router, a router connecting two internal networks, or a Layer 3 switch, such as a PowerConnect 8000-series and 8100-series switches.

You can also create an ACL that limits access to the management interfaces based on the connection method (for example, Telnet or HTTP) and/or the source IP address.

The PowerConnect 8000-series and 8100-series switches support ACL configuration in both the ingress and egress direction. Egress ACLs provide the capability to implement security rules on the egress flows (traffic leaving a port) rather than the ingress flows (traffic entering a port). Ingress and egress ACLs can be applied to any physical port, port-channel (LAG), or VLAN routing port.

Depending on whether an ingress or egress ACL is applied to a port, when the traffic enters (ingress) or leaves (egress) a port, the ACL compares the criteria configured in its rules, in order, to the fields in a packet or frame to check for matching conditions. The ACL forwards or blocks the traffic based on the rules.

NOTE: Every ACL is terminated by an implicit **deny all** rule, which covers any packet not matching a preceding explicit rule.

You can set up ACLs to control traffic at Layer 2, Layer 3, or Layer 4. MAC ACLs operate on Layer 2. IP ACLs operate on Layers 3 and 4. PowerConnect 8000-series and 8100-series switches support both IPv4 and IPv6 ACLs.

What Are MAC ACLs?

MAC ACLs are Layer 2 ACLs. You can configure the rules to inspect the following fields of a packet:

- Source MAC address
- Source MAC mask
- Destination MAC address
- Destination MAC mask
- VLAN ID
- Class of Service (CoS) (802.1p)
- EtherType

L2 ACLs can apply to one or more interfaces.

Multiple access lists can be applied to a single interface; sequence number determines the order of execution.

You can assign packets to queues using the assign queue option.

What Are IP ACLs?

IP ACLs classify for Layers 3 and 4 on IPv4 or IPv6 traffic.

Each ACL is a set of up to ten rules applied to inbound traffic. Each rule specifies whether the contents of a given field should be used to permit or deny access to the network, and may apply to one or more of the following fields within a packet:

- Destination IP with wildcard mask
- Destination L4 Port
- Every Packet
- IP DSCP
- IP Precedence
- IP TOS
- Protocol
- Source IP with wildcard mask
- Source L4 port
- Destination Layer 4 port

What Is the ACL Redirect Function?

The redirect function allows traffic that matches a permit rule to be redirected to a specific physical port or LAG instead of processed on the original port. The redirect function and mirror function are mutually exclusive. In other words, you cannot configure a given ACL rule with mirror and redirect attributes.

What Is the ACL Mirror Function?

ACL mirroring provides the ability to mirror traffic that matches a permit rule to a specific physical port or LAG. Mirroring is similar to the redirect function, except that in flow-based mirroring a copy of the permitted traffic is delivered to the mirror interface while the packet itself is forwarded normally through the device. You cannot configure a given ACL rule with both mirror and redirect attributes.

Using ACLs to mirror traffic is considered to be flow-based mirroring since the traffic flow is defined by the ACL classification rules. This is in contrast to port mirroring, where all traffic encountered on a specific interface is replicated on another interface.

What Is ACL Logging

ACL Logging provides a means for counting the number of “hits” against an ACL rule. When you configure ACL Logging, you augment the ACL deny rule specification with a "log" parameter that enables hardware hit count collection and reporting. The switch uses a fixed five minute logging interval, at which time trap log entries are written for each ACL logging rule that accumulated a non-zero hit count during that interval. You cannot configure the logging interval.

What Are Time-Based ACLs?

The time-based ACL feature allows the switch to dynamically apply an explicit ACL rule within an ACL for a predefined time interval by specifying a time range on a per-rule basis within an ACL, so that the time restrictions are imposed on the ACL rule.

With a time-based ACL, you can define when and for how long an individual rule of an ACL is in effect. To apply a time to an ACL, first you define a specific time interval and then apply it to an individual ACL rule so that it is operational only during the specified time range, for example, during a specified time period or on specified days of the week.

A time range can be absolute (specific time) or periodic (recurring). If an absolute and periodic time range entry are defined within the same time range, the periodic timer is active only when the absolute timer is active.

NOTE: Adding a conflicting periodic time range to an absolute time range will cause the time range to become inactive. For example, consider an absolute time range from 8:00 AM Tuesday March 1st 2011 to 10 PM Tuesday March 1st 2011. Adding a periodic entry using the 'weekend' keyword will cause the time-range to become inactive because Tuesdays are not on the weekend.

A named time range can contain up to 10 configured time ranges. Only one absolute time range can be configured per time range. During the ACL configuration, you can associate a configured time range with the ACL to provide additional control over permitting or denying a user access to network resources.

Benefits of using time-based ACLs include:

- Providing more control over permitting or denying a user access to resources, such as an application (identified by an IP address/mask pair and a port number).
- Providing control of logging messages. Individual ACL rules defined within an ACL can be set to log traffic only at certain times of the day so you can simply deny access without needing to analyze many logs generated during peak hours.

What Are the ACL Limitations?

The following limitations apply to ingress and egress ACLs.

- Maximum of 100 ACLs.
- Maximum rules per ACL is a maximum of 1023 rules, with 1023 ingress and 511 egress IPv4 rules or 509 ingress and 253 egress IPv6 rules.
- You can configure mirror or redirect attributes for a given ACL rule, but not both.
- The PowerConnect 8000-series and 8100-series switches support a limited number of counter resources, so it may not be possible to log every ACL rule. You can define an ACL with any number of logging rules, but the number of rules that are actually logged cannot be determined until the ACL is applied to an interface. Furthermore, hardware counters that become available after an ACL is applied are not retroactively assigned to rules that were unable to be logged (the ACL must be un-applied then re-applied). Rules that are unable to be logged are still active in the ACL for purposes of permitting or denying a matching packet. If console logging is enabled and the severity is set to Info (6) or a lower severity, a log entry may appear on the screen.
- The order of the rules is important: when a packet matches multiple rules, the first rule takes precedence. Also, once you define an ACL for a given port, all traffic not specifically permitted by the ACL is denied access.

NOTE: The actual number of ACLs and rules supported depends on the resources consumed by other processes and configured features running on the switch.

How Are ACLs Configured?

To configure ACLs, follow these steps:

1. Create a MAC ACL by specifying a name.
2. Create an IP ACL by specifying a number.
3. Add new rules to the ACL.
4. Configure the match criteria for the rules.
5. Apply the ACL to one or more interfaces.

Preventing False ACL Matches

Be sure to specify ACL access-list, permit, and deny rule criteria as fully as possible to avoid false matches. This is especially important in networks with protocols such as FCoE that have newly-introduced EtherType values. For example, rules that specify a TCP or UDP port value should also specify the TCP or UDP protocol and the IPv4 or IPv6 EtherType. Rules that specify an IP protocol should also specify the EtherType value for the frame.

In general, any rule that specifies matching on an upper-layer protocol field should also include matching constraints for each of the lower-layer protocols. For example, a rule to match packets directed to the well-known UDP port number 22 (SSH) should also include matching constraints on the IP protocol field (protocol=0x11 or UDP) and the EtherType field (EtherType= 0x0800 or IPv4). Figure 21-1 lists commonly-used EtherTypes numbers:

Table 21-1. Common EtherType Numbers

EtherType	Protocol
0x0800	Internet Protocol version 4 (IPv4)
0x0806	Address Resolution Protocol (ARP)
0x0842	Wake-on LAN Packet
0x8035	Reverse Address Resolution Protocol (RARP)
0x8100	VLAN tagged frame (IEEE 802.1Q)

Table 21-1. Common EtherType Numbers (Continued)

EtherType	Protocol
0x86DD	Internet Protocol version 6 (IPv6)
0x8808	MAC Control
0x8809	Slow Protocols (IEEE 802.3)
0x8870	Jumbo frames
0x888E	EAP over LAN (EAPOL – 802.1x)
0x88CC	Link Layer Discovery Protocol
0x8906	Fibre Channel over Ethernet
0x8914	FCoE Initialization Protocol
0x9100	Q in Q

Figure 21-2 lists commonly-used IP protocol numbers:

\

Table 21-2. Common IP Protocol Numbers

IP Protocol Number	Protocol
0x00	IPv6 Hop-by-hop option
0x01	ICMP
0x02	IGMP
0x06	TCP
0x08	EGP
0x09	IGP
0x11	UDP

Configuring ACLs (Web)

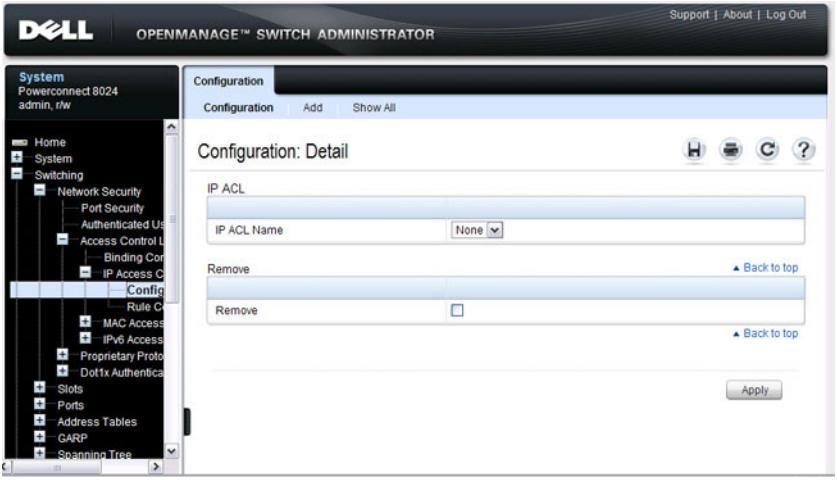
This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring ACLs on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

IP ACL Configuration

Use the **IP ACL Configuration** page to add or remove IP-based ACLs.

To display the **IP ACL Configuration** page, click **Switching** → **Network Security** → **Access Control Lists** → **IP Access Control Lists** → **Configuration** in the navigation panel.

Figure 21-1. IP ACL Configuration

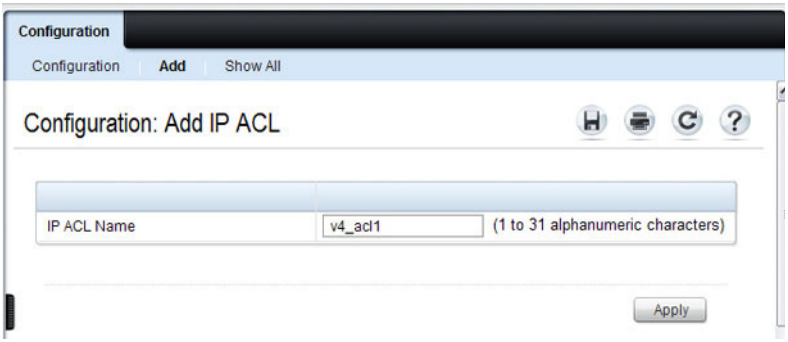


Adding an IPv4 ACL

To add an IPv4 ACL:

1. Open the **IP ACL Configuration** page.
2. Click **Add** to display the **Add IP ACL** page.
3. Specify an ACL name.

Figure 21-2. Add IP ACL



4 Click **Apply**.

Removing IPv4 ACLs

To delete an IPv4 ACL:

1 From the **IP ACL Name** menu on the **IP ACL Configuration** page, select the ACL to remove.

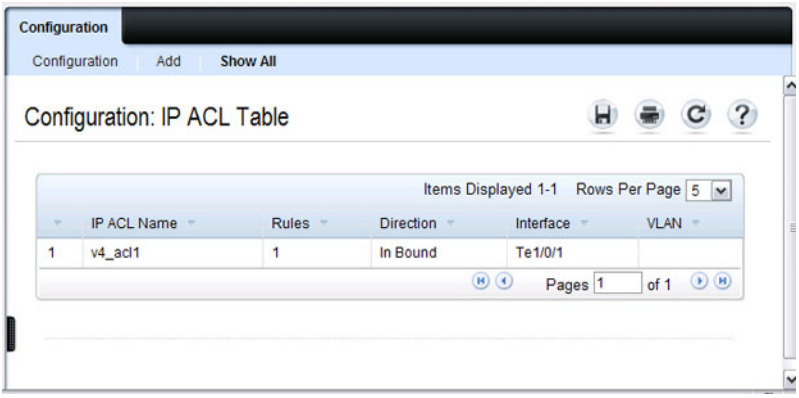
2 Select the **Remove** checkbox.

3 Click **Apply**.

Viewing IPv4 ACLs

To view configured ACLs, click **Show All** from the **IP ACL Configuration** page.

Figure 21-3. View IPv4 ACLs



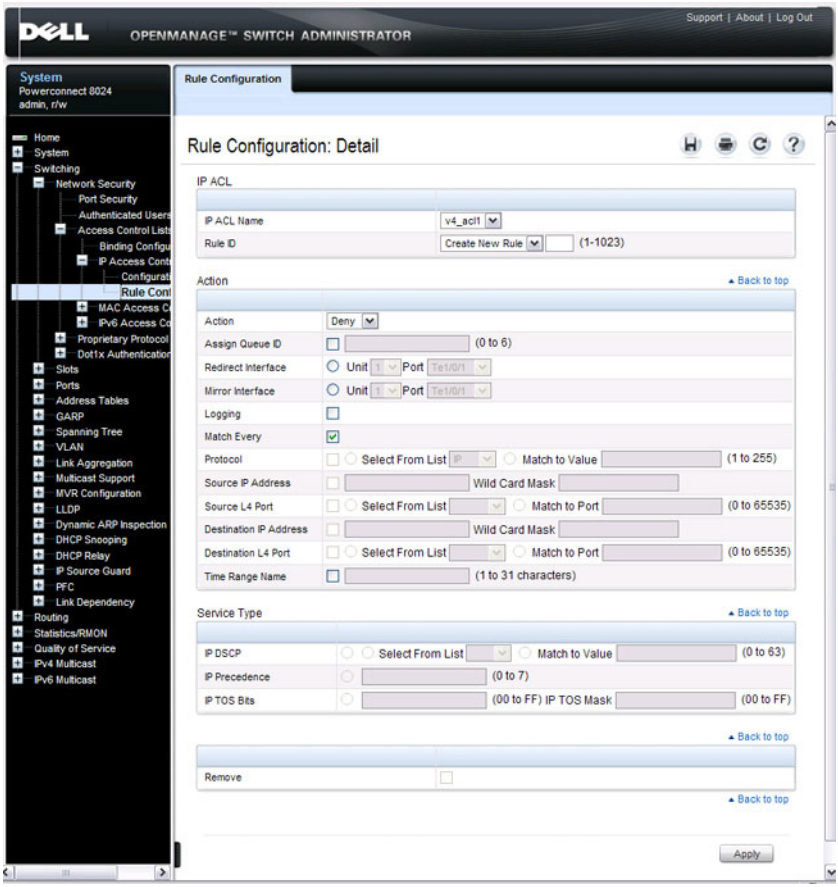
IP ACL Rule Configuration

Use the **IP ACL Rule Configuration** page to define rules for IP-based ACLs. The access list definition includes rules that specify whether traffic matching the criteria is forwarded normally or discarded. Additionally, you can specify to assign traffic to a particular queue, filter on some traffic, change VLAN tag, shut down a port, and/or redirect the traffic to a particular port.

NOTE: There is an implicit **deny all** rule at the end of an ACL list. This means that if an ACL is applied to a packet and if none of the explicit rules match, then the final implicit "deny all" rule applies and the packet is dropped.

To display the **IP ACL Rule Configuration** page, click **Switching → Network Security → Access Control Lists → IP Access Control Lists → Rule Configuration** in the navigation panel.

Figure 21-4. IP ACL - Rule Configuration



Removing an IP ACL Rule

To delete an IP ACL rule:

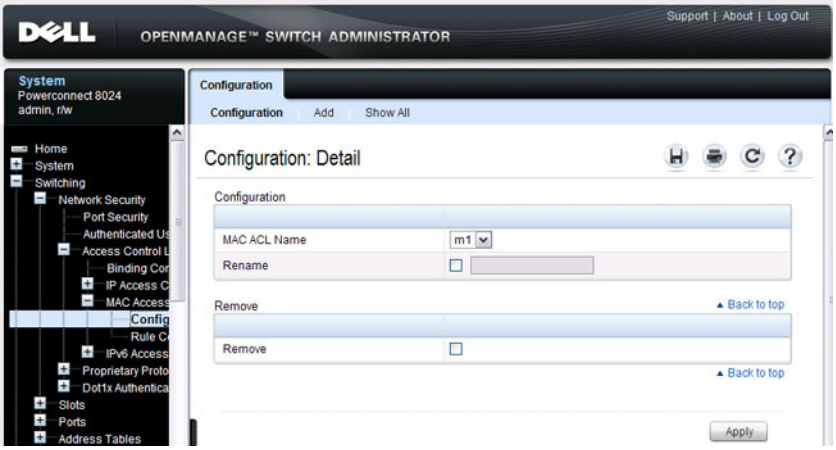
1. From the **Rule ID** menu, select the ID of the rule to delete.
2. Select the **Remove** option near the bottom of the page.
3. Click **Apply** to remove the selected rule.

MAC ACL Configuration

Use the **MAC ACL Configuration** page to define a MAC-based ACL.

To display the **MAC ACL Configuration** page, click **Switching → Network Security → Access Control Lists → MAC Access Control Lists → Configuration** in the navigation panel.

Figure 21-5. MAC ACL Configuration

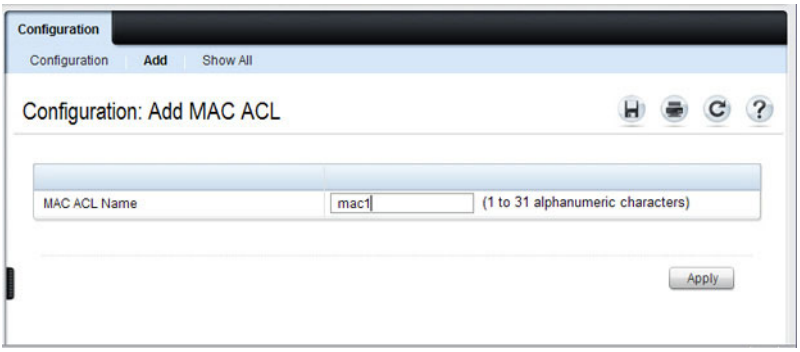


Adding a MAC ACL

To add a MAC ACL:

1. Open the **MAC ACL Configuration** page.
2. Click **Add** to display the **Add MAC ACL** page.
3. Specify an ACL name.

Figure 21-6. Add MAC ACL



4 Click **Apply**.

Renaming or Removing MAC ACLs

To rename or delete a MAC ACL:

1 From the **MAC ACL Name** menu on the **MAC ACL Configuration** page, select the ACL to rename or remove.

2 To rename the ACL, select the **Rename** checkbox and enter a new name in the associated field.

3 To remove the ACL, select the **Remove** checkbox.

4 Click **Apply**.

Viewing MAC ACLs

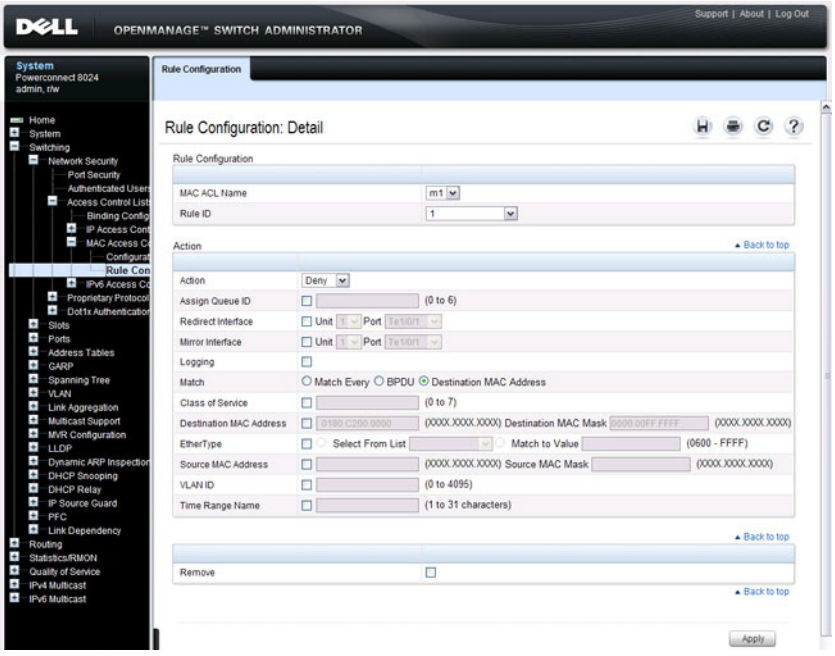
To view configured ACLs, click **Show All** from the **MAC ACL Configuration** page.

MAC ACL Rule Configuration

Use the **MAC ACL Rule Configuration** page to define rules for MAC-based ACLs. The access list definition includes rules that specify whether traffic matching the criteria is forwarded normally or discarded. A default **deny all** rule is the last rule of every list.

To display the **MAC ACL Rule Configuration** page, click **Switching → Network Security → Access Control Lists → MAC Access Control Lists → Rule Configuration** in the navigation panel.

Figure 21-7. MAC ACL Rule Configuration



Removing a MAC ACL Rule

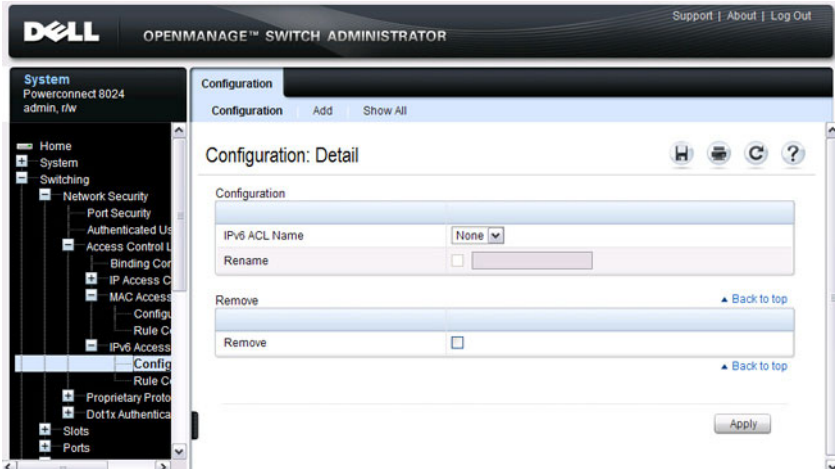
To delete a MAC ACL rule:

1. From the **Rule ID** menu, select the ID of the rule to delete.
2. Select the **Remove** option near the bottom of the page.
3. Click **Apply** to remove the selected rule.

IPv6 ACL Configuration

Use the **IPv6 ACL Configuration** page to add or remove IP-based ACLs. To display the IP ACL Configuration page, click **Switching → Network Security → Access Control Lists → IPv6 Access Control Lists → IPv6 ACL Configuration** in the navigation panel.

Figure 21-8. IPv6 ACL Configuration

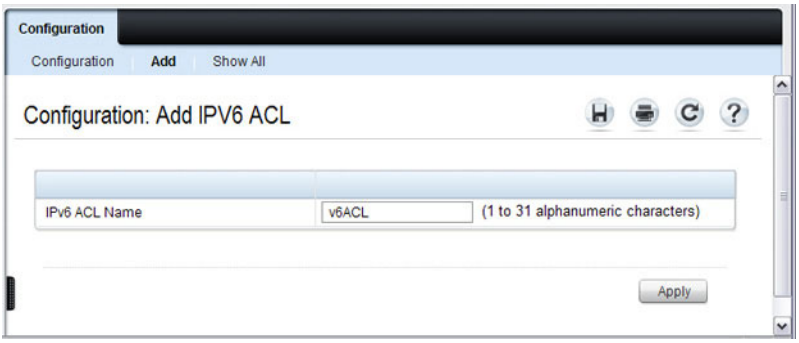


Adding an IPv6 ACL

To add an IPv6 ACL:

1. Open the **IPv6 ACL Configuration** page.
2. Click **Add** to display the **Add IPv6 ACL** page.
3. Specify an ACL name.

Figure 21-9. Add IPv6 ACL



4 Click **Apply**.

Removing IPv6 ACLs

To delete an IPv6 ACL:

1 From the **IPv6 ACL Name** menu on the **IPv6 ACL Configuration** page, select the ACL to remove.

2 Select the **Remove** checkbox.

3 Click **Apply**.

Viewing IPv6 ACLs

To view configured ACLs, click **Show All** from the **IPv6 ACL Configuration** page. The **IPv6 ACL Table** page displays.

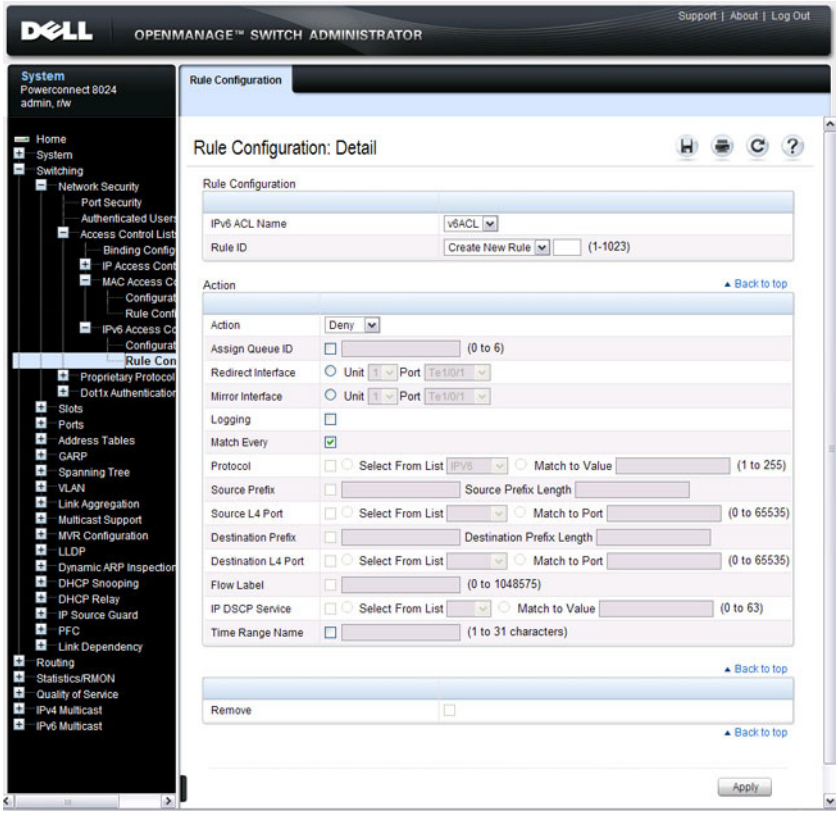
IPv6 ACL Rule Configuration

Use the IPv6 ACL Rule Configuration page to define rules for IPv6-based ACLs. The access list definition includes rules that specify whether traffic matching the criteria is forwarded normally or discarded. Additionally, you can specify to assign traffic to a particular queue, filter on some traffic, change VLAN tag, shut down a port, and/or redirect the traffic to a particular port. By default, no specific value is in effect for any of the IPv6 ACL rules.

There is an implicit **deny all** rule at the end of an ACL list. This means that if an ACL is applied to a packet and if none of the explicit rules match, then the final implicit **deny all** rule applies and the packet is dropped.

To display the IPv6 ACL Rule Configuration page, click **Switching** → **Network Security** → **Access Control Lists** → **IPv6 Access Control Lists** → **Rule Configuration** in the navigation menu.

Figure 21-10. IPv6 ACL - Rule Configuration



Removing an IPv6 ACL Rule

To delete an IPv6 ACL rule:

1. From the **Rule ID** menu, select the ID of the rule to delete.
2. Select the **Remove** option near the bottom of the page.
3. Click **Apply** to remove the selected rule.

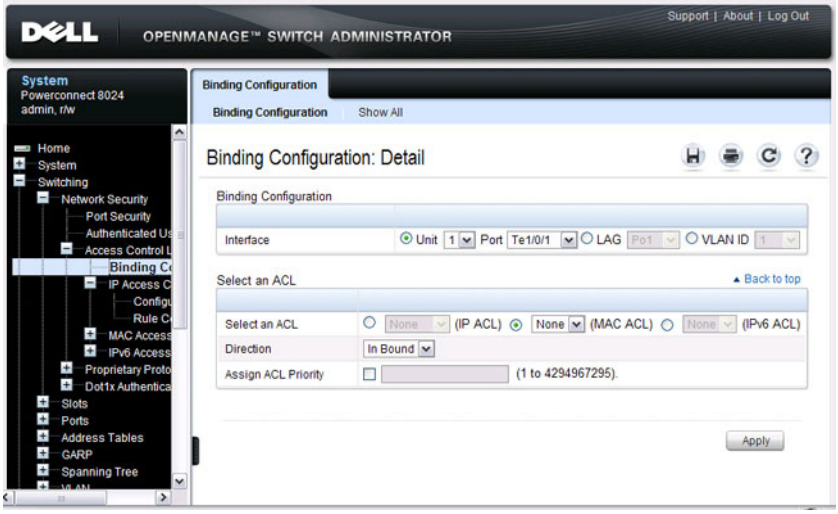
ACL Binding Configuration

When an ACL is bound to an interface, all the rules that have been defined are applied to the selected interface. Use the **ACL Binding Configuration** page to assign ACL lists to ACL Priorities and Interfaces.

From the web interface, you can configure the ACL rule in the ingress or egress direction so that the ACLs implement security rules for packets entering or exiting the port. You can apply ACLs to any physical (including 10 Gb) interface, LAG, or routing port.

To display the **ACL Binding Configuration** page, click **Switching → Network Security → Access Control Lists → Binding Configuration** in the navigation panel.

Figure 21-11. ACL Binding Configuration

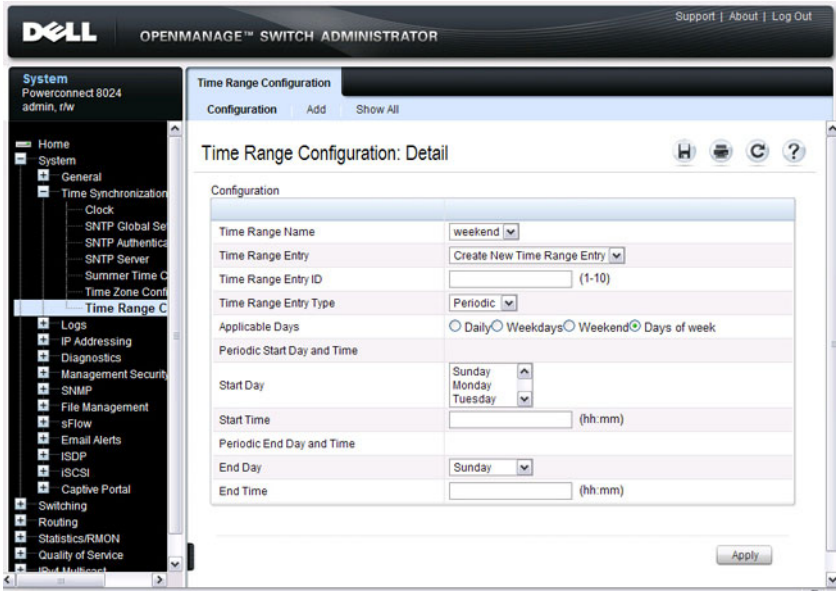


Time Range Entry Configuration

Use the **Time Range Entry Configuration** page to define time ranges to associate with ACL rules.

To display the **Time Range Entry Configuration** page, click **System → Time Synchronization → Time Range Configuration** in the navigation panel. The following image shows the page after at least one time range has been added. Otherwise, the page indicates that no time ranges are configured, and the time range configuration fields are not displayed.

Figure 21-12. Time Range Configuration

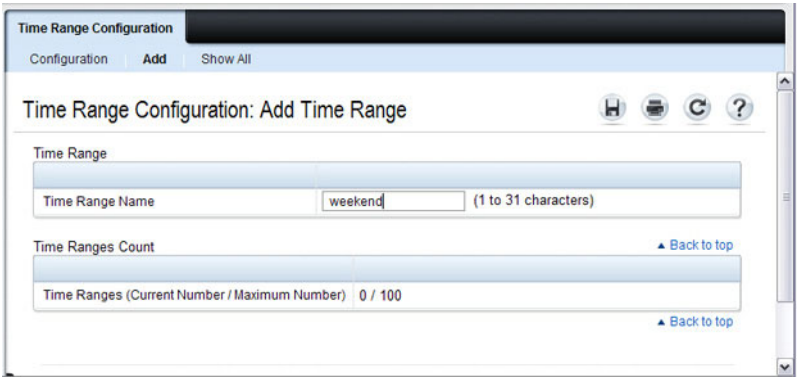


Adding a Time Range

To configure a time range:

1. From the **Time Range Entry Configuration** page, click **Add**.
2. Specify a name to identify the time range.

Figure 21-13. Add a Time Range



3 Click **Apply**.

4 Click **Configuration** to return to the **Time Range Entry Configuration** page.

5 In the **Time Range Name** field, select the name of the time range to configure.

6 Specify an ID for the time range. You can configure up to 10 different time range entries to include in the named range. However, only one absolute time entry is allowed per time range.

7 Configure the values for the time range entry.

8 Click **Apply**.

9 To add additional entries to the named time range, repeat step 5 through step 8.

Configuring ACLs (CLI)

This section provides information about the commands you use to create and configure ACLs. For more information about the commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Configuring an IPv4 ACL

Beginning in Privileged EXEC mode, use the following commands to create an IPv4 ACL, configure rules for the ACL, and bind the ACL to an interface.

NOTE: The **ip access-group** command can be issued in Global Configuration mode or Interface configuration mode. If it is applied in Global Configuration mode, the ACL binding is applied to all interfaces. If it is applied in Interface Configuration mode, it is applied only to the specified interfaces within the mode.

Command	Purpose
configure	Enter global configuration mode.
access-list *name* {**deny** | **permit**} {**every** | {[**icmp** | **igmp** | **ip** | **tcp** | **udp** | *number*] {*srcip srcmask* | **any**} [**eq** [*portkey* | *portvalue*]] {**dstip** *dstmask* | **any**} [**eq** [**portkey** | **portvalue**]] [**precedence** *precedence* | **tos** *tos tosmask* | **dscp** *dscp*] [**log**] [**time-range** *time-range-name*] [**assign-queue** *queue-id*] [**redirect** *interface* | **mirror** *interface*]}}	Create a named ACL (if it does not already exist) and create a rule for the named ACL. If the ACL already exists, this command creates a new rule for the ACL. • *list-name* — Access-list name up to 31 characters in length. • **deny** | **permit** — Specifies whether the IP ACL rule permits or denies an action. • **every** — Allows all protocols. • **eq** — Equal. Refers to the Layer 4 port number being used as match criteria. The first reference is source match criteria, the second is destination match criteria. • *number* — Standard protocol number. Protocol keywords **icmp**, **igmp**, **ip**, **tcp**, **udp**. • *srcip* — Source IP address. • *srcmask* — Source IP mask. • *dstip* — Destination IP address. • *dstmask* — Destination IP mask.

Command	Purpose
(continued)	• *portvalue* — The source layer 4 port match condition for the ACL rule is specified by the port value parameter (Range: 0–65535). • *portkey* — Or you can specify the *portkey*, which can be one of the following keywords: domain, echo, ftp, ftpdata, http, smtp, snmp, telnet, tftp, and www. • **log** — Specifies that this rule is to be logged. • *time-range-name* — Specifies the named time range to associate with the ACL rule. • **assign-queue** *queue-id* — Specifies the particular hardware queue for handling traffic that matches the rule. (Range: 0-6) • **mirror** *interface* — Allows the traffic matching this rule to be copied to the specified interface. • **redirect** *interface* — This parameter allows the traffic matching this rule to be forwarded to the specified interface.
interface *interface*	(Optional) Enter interface configuration mode for the specified interface. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**. You can also specify a range of interfaces with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12.
ip access-group *name direction seqnum*	Bind the specified ACL to an interface. **NOTE:** To apply this ACL to all interfaces, issue the command in Global Configuration mode. • *name* — Access list name. (Range: Valid IP access-list name up to 31 characters in length) • *direction* — Direction of the ACL. (Range: In or out. Default is *in*.) • *seqnum* — Precedence for this interface and direction. A lower sequence number has higher precedence. Range: 1 – 4294967295. Default is1.

Command	Purpose
CTRL + Z	Exit to Privileged EXEC mode.
show ip access-lists [*name*]	Display all IPv4 access lists and all of the rules that are defined for the IPv4 ACL. Use the optional *name* parameter to identify a specific IPv4 ACL to display.

Configuring a MAC ACL

Beginning in Privileged EXEC mode, use the following commands to create an MAC ACL, configure rules for the ACL, and bind the ACL to an interface.

Command	Purpose
configure	Enter global configuration mode.
mac access-list extended *name*	Create a named MAC ACL. This command also enters MAC Access List Configuration mode. If a MAC ACL with this name already exists, this command enters the mode to update the existing ACL.
{deny | permit} {*srcmac srcmacmask* | **any**} {*dstmac dstmacmask* | **any** | **bpdu** } [{*ethertypekey* | *0x0600-0xFFFF* }] [**vlan eq** *0-4095*] [**cos** *0-7*] [**secondary-vlan eq** *0-4095*] [**secondary-cos** *0-7*] [**log**] [**time-range** *time-range-name*] [**assign-queue** *queue-id*] [{**mirror** |**redirect**} *interface*]	Specify the rules (match conditions) for the MAC access list. • *srcmac* — Valid source MAC address in format xxxx.xxxx.xxxx. • *srcmacmask* — Valid MAC address bitmask for the source MAC address in format xxxx.xxxx.xxxx. • **any** — Packets sent to or received from any MAC address • *dstmac* — Valid destination MAC address in format xxxx.xxxx.xxxx. • *destmacmask* — Valid MAC address bitmask for the destination MAC address in format xxxx.xxxx.xxxx. • **bpdu** — Bridge protocol data unit • *ethertypekey* — Either a keyword or valid four-digit hexadecimal number. (Range: Supported values are appletalk, arp, ibmsna, ipv4, ipv6, ipx, mplsmcast, mplsucast, Netbios, novell, pppoe, rarp.) • *0x0600-0xFFFF* — Specify custom EtherType value (hexadecimal range 0x0600-0xFFFF)

Command	Purpose
(Continued)	• **vlan eq** — VLAN number. (Range 0-4095) • **cos** — Class of service. (Range 0-7) • **log** — Specifies that this rule is to be logged. • *time-range-name* — Specifies the named time range to associate with the ACL rule. • **assign-queue** — Specifies particular hardware queue for handling traffic that matches the rule. • *queue-id* — 0-6, where n is number of user configurable queues available for that hardware platform. • **mirror** *interface* — Allows the traffic matching this rule to be copied to the specified interface. • **redirect** *interface* — This parameter allows the traffic matching this rule to be forwarded to the specified interface.
interface *interface*	(Optional) Enter interface configuration mode for the specified interface. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**. You can also specify a range of interfaces with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12.
mac access-group *name direction seqnum*	Bind the specified MAC ACL to an interface. **NOTE:** To apply this ACL to all interfaces, issue the command in Global Configuration mode. • *name* — Access list name. (Range: Valid MAC access-list name up to 31 characters in length) • *direction* — Direction of the ACL. (Range: In or out. Default is *in*.) • *seqnum* — Precedence for this interface and direction. A lower sequence number has higher precedence. Range: 1 – 4294967295. Default is1.
CTRL + Z	Exit to Privileged EXEC mode.

Command	Purpose
show mac access-lists [*name*]	Display all MAC access lists and all of the rules that are defined for the MAC ACL. Use the optional *name* parameter to identify a specific MAC ACL to display.

Configuring an IPv6 ACL

Beginning in Privileged EXEC mode, use the following commands to create an IPv6 ACL, configure rules for the ACL, and bind the ACL to an interface.

Command	Purpose
configure	Enter global configuration mode.
ipv6 access-list *name*	Create a named IPv6 ACL. This command also enters IPv6 Access List Configuration mode. If an IPv6 ACL with this name already exists, this command enters the mode to update the existing ACL.
{**permit** | **deny**} {**every** | {{**icmp** | **igmp** | **ipv6** | **tcp** | **udp** | *number*} {**any** | *source ipv6 prefix/prefix length*} [**eq** {*portkey* | *portvalue*}] {**any** | *destination ipv6 prefix/prefix length*} [**eq** {*portkey* | *portvalue*}] [flow-label *value*] [**dscp** *dscp*]}} [**log**] [**time-range** *time-range-name*] [**assign-queue** *queue-id*] [{**mirror** | **redirect**} *interface*]	Specify the match conditions for the IPv6 access list. • **deny** | **permit** — Specifies whether the IP ACL rule permits or denies an action. • **every** — Allows all protocols. • *number* — Standard protocol number or protocol keywords **icmp**, **igmp**, **ipv6**, **tcp**, **udp**. • *source ipv6 prefix* — IPv6 prefix in IPv6 global address format. • *prefix-length* — IPv6 prefix length value. • **eq** — Equal. Refers to the Layer 4 port number being used as a match criteria. The first reference is source match criteria, the second is destination match criteria. • *portkey* — Or you can specify the portkey, which can be one of the following keywords: **domain**, **echo**, **efts**, **ftpdata**, **http**, **smtp**, **snmp**, **telnet**, **tftp**, and **www**. • *portvalue* — The source layer 4 port match condition for the ACL rule is specified by the port value parameter. (Range: 0–65535).

Command	Purpose
(Continued)	• *destination ipv6 prefix* — IPv6 prefix in IPv6 global address format. • **flow label** *value* — The value to match in the Flow Label field of the IPv6 header (Range 0–1048575). • **dscp** *dscp* — Specifies the TOS for an IPv6 ACL rule depending on a match of DSCP values using the parameter dscp. • **log** — Specifies that this rule is to be logged. • *time-range-name* — Specifies the named time range to associate with the ACL rule. • **assign-queue** *queue-id* — Specifies particular hardware queue for handling traffic that matches the rule. • **mirror** *interface* — Allows the traffic matching this rule to be copied to the specified interface. • **redirect** *interface* — This parameter allows the traffic matching this rule to be forwarded to the specified interface.
interface *interface*	(Optional) Enter interface configuration mode for the specified interface. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**. You can also specify a range of interfaces with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12.
ipv6 traffic-filter *name direction* [**sequence** *seq-num*]	Bind the specified IPv6 ACL to an interface. **NOTE:** To apply this ACL to all interfaces, issue the command in Global Configuration mode. • *name* — Access list name. (Range: Valid IPv6 access-list name up to 31 characters in length) • *direction* — Direction of the ACL. (Range: In or out. Default is *in*.) • *seqnum* — Precedence for this interface and direction. A lower sequence number has higher precedence. Range: 1 – 4294967295. Default is1.

Command	Purpose
CTRL + Z	Exit to Privileged EXEC mode.
show ipv6 access-lists [*name*]	Display all IPv6 access lists and all of the rules that are defined for the IPv6 ACL. Use the optional *name* parameter to identify a specific IPv6 ACL to display.

Configuring a Time Range

Beginning in Privileged EXEC mode, use the following commands to create a time range and configure time-based entries for the time range.

Command	Purpose
configure	Enter global configuration mode.
time-range *name*	Create a named time range and enter the Time-Range Configuration mode for the range.
absolute {[**start** *time date*] [**end** *time date*]}	Configure a nonrecurring time entry for the named time range. • **start** *time date* — Time and date the ACL rule starts going into effect. The time is expressed in a 24-hour clock, in the form of hours:minutes. For example, 8:00 is 8:00 am and 20:00 is 8:00 pm. The date is expressed in the format day month year. If no start time and date are specified, the configuration statement is in effect immediately. • **end** *time date* — Time and date the ACL rule is no longer in effect.

Command	Purpose
periodic {*days-of-the-week time*} **to** {[*days-of-the-week*] *time*}	Configure a recurring time entry for the named time range. • *days-of-the-week* —The first occurrence indicates the starting day(s) the ACL goes into effect. The second occurrence is the ending day(s) when the ACL rule is no longer in effect. If the end *days-of-the-week* are the same as the start, they can be omitted This variable can be any single day or combinations of days: **Monday**, **Tuesday**, **Wednesday**, **Thursday**, **Friday**, **Saturday**, **Sunday**. Other possible values are: – **daily** -- Monday through Sunday – **weekdays** -- Monday through Friday – **weekend** -- Saturday and Sunday • *time* — Time the ACL rule starts going into effect (first occurrence) or ends (second occurrence). The time is expressed in a 24-hour clock, in the form of hours:minutes.
CTRL + Z	Exit to Privileged EXEC mode.
show time-range [*name*]	View information about all configured time ranges, including the absolute/periodic time entries that are defined for each time range. Use the *name* variable to view information about the specified time range.

ACL Configuration Examples

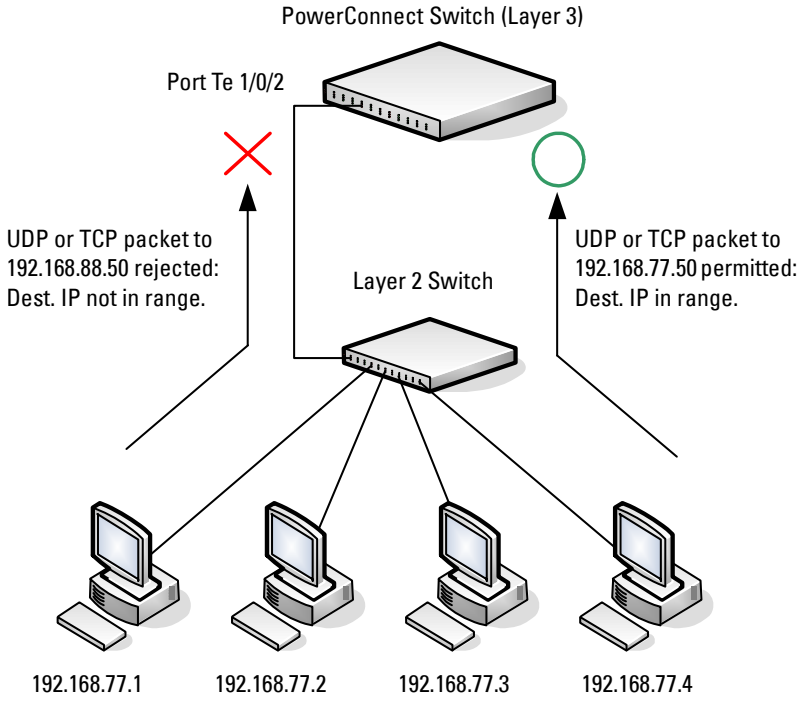
This section contains the following examples:

- Configuring an IP ACL
- Configuring a MAC ACL
- Configuring a Time-Based ACL
- Configuring a Management Access List

Configuring an IP ACL

The commands in this example set up an IP ACL that permits hosts in the 192.168.77.0/24 subnet to send TCP and UDP traffic only to the host with an IP address of 192.168.77.50. The ACL is applied to port 2 on the PowerConnect switch.

Figure 21-14. IP ACL Example Network Diagram



To configure the switch:

1. Create an ACL named list1 and configures a rule for the ACL that permits packets carrying TCP traffic that matches the specified Source IP address (192.168.77.0/24), and sends these packets to the specified Destination IP address (192.168.77.50).

   ```
   console#config
   console(config)#access-list list1 permit tcp
   192.168.77.0 0.0.0.255 192.168.77.50 0.0.0.0
   ```

2. Define the rule to set similar conditions for UDP traffic as for TCP traffic.

   ```
   console(config)#access-list list1 permit udp
   192.168.77.0 0.0.0.255 192.168.77.3 0.0.0.255
   console(config)#exit
   ```

3 Apply the rule to inbound (ingress) traffic on 10-Gigabit Ethernet Port 2. Only traffic matching the criteria will be accepted on this port.

```
console(config)#interface te1/0/2
console(config-if-Te1/0/2)#ip access-group list1
in
console(config-if-Te11/0/2)#exit
```

Configuring a MAC ACL

The following example creates a MAC ACL named mac1 that denies all IPX traffic on all ports. All other type of traffic is permitted.

To configure the switch:

1 Create a MAC Access List named mac1

```
console#config
console(config)#mac access-list extended mac1
```

2 Configure a rule to deny all IPX traffic, regardless of the source or destination MAC address.

```
console(config-mac-access-list)#deny any any ipx
```

3 Configure a rule to permit all other types of traffic, regardless of the source or destination MAC address.

```
console(config-mac-access-list)#permit any any
console(config-mac-access-list)#exit
```

4 Bind the ACL to all ports.

```
console(config)#mac access-group mac1 in
console(config)#exit
```

5 View information about the configured ACL.

```
console#show mac access-lists

Current number of all ACLs: 1  Maximum number of
all ACLs: 100

MAC ACL Name  Rules      Interface(s) Direction
------------- --------   ------------ ---------
```

```
mac1            2           ch1-12,       Inbound
                            Te1/0/1-
                            Te1/0/24

console#show mac access-lists mac1

MAC ACL Name: mac1

Inbound Interface(s):
ch1-12,Te1/0/1-Te1/0/24


Rule Number: 1
Action.................................. deny
Ethertype.................................ipx

Rule Number: 2
Action.................................. permit
Match All............................... TRUE
```

Configuring a Time-Based ACL

The following example configures an ACL that denies HTTP traffic from 8:00 pm to 12:00 pm and 1:00 pm to 6:00 pm on weekdays and from 8:30 am to 12:30 pm on weekends. The ACL affects all hosts connected to ports that are members of VLAN 100. The ACL permits VLAN 100 members to browse the Internet only during lunch and after hours.

To configure the switch:

1. Create a time range called *work-hours*.

   ```
   console#config
   console(config)#time-range work-hours
   ```

2. Configure an entry for the time range that applies to the morning shift Monday through Friday.

   ```
   console(config-time-range)#periodic weekdays 8:00
   to 12:00
   ```

3. Configure an entry for the time range that applies to the afternoon shift Monday through Friday.

   ```
   console(config-time-range)#periodic weekdays 13:00
   to 18:00
   ```

4. Configure an entry for the time range that applies to Saturday and Sunday.

   ```
   console(config-time-range)#periodic weekend 8:30
   to 12:30
   console(config-time-range)#exit
   ```

5. Create an ACL named *web-limit* that denies HTTP traffic during the *work-hours* time range.

   ```
   console(config)#access-list web-limit deny tcp any
   any eq http time-range work-hours
   ```

6. Enter interface configuration mode for VLAN 100 and apply the ACL to ingress traffic.

   ```
   console(config)#interface vlan 100
   console(config-if-vlan100)#ip access-group web-
   limit in
   console(config-if-vlan100)#exit
   console(config)#exit
   ```

7 Verify the configuration.

```
console#show ip access-lists web-limit

IP ACL Name: web-limit

Inbound VLAN(s):
 100

Rule Number: 1
Action..............................  deny
Match All...........................  FALSE
Protocol............................  6(tcp)
Source IP Address...................  any
Destination IP Address..............  any
Destination L4 Port Keyword.........  80(www/http)ip
Time Range Name.....................work-hours
Rule Status.........................  inactive
```

Configuring a Management Access List

You can create a management access list that contains rules to apply to one or more in-band ports, LAGs, or VLANs to limit management access by method (for example, Telnet or HTTP) and/or source IP address.

NOTE: Management ACLs cannot be applied to the OOB port.

NOTE: Management ACLs can be applied only to in-band ports and cannot be applied to the OOB port.

Management Access List Commands

Beginning in Privileged EXEC mode, use the following commands to create a management access list. There is an implicit deny-all rule at the end of every management ACL. This means that any host that does not meet the criteria defined in a **permit** command is denied access to the management interface.

Command	Purpose
configure	Enter Global Configuration mode.

Command	Purpose
management access-list *name*	Define an access list for management, and enter the access-list for configuration.
permit ip-source *ip-address* [**mask** *mask* | *prefix-length*] [*interface-type interface-number*] [**service** *service*] [**priority** *priority-value*]	Allow access to the management interface from hosts that meet the specified IP address value and other optional criteria. • *interface-type interface-number* — A valid port, LAG, or VLAN interface, for example te1/0/13, port-channel 3, or vlan 200. • *ip-address* — Source IP address. • **mask** *mask* — Specifies the network mask of the source IP address. • **mask** *prefix-length* — Specifies the number of bits that comprise the source IP address prefix. The prefix length must be preceded by a forward slash (/). (Range: 0–32) • **service** *service* — Indicates service type. Can be one of the following: telnet, ssh, http, https, tftp, snmp, sntp, or any. • **priority** *priority-value* — Priority for the rule. (Range: 1 – 64)
permit {*interface-type interface-number*} [**service** *service*] [**priority** *priority-value*]	Permit access to the management interface from the specified port, VLAN, or LAG and meet the other optional criteria.
permit service *service* [**priority** *priority-value*]	Permit access to the management interface from the specified service.
exit	Exit to Global Configuration mode.
management access-class {**console-only** | *name*}	Activate the management ACL or restrict access so that it is available only through the console port.
exit	Exit to Privileged EXEC mode.
show management access-class	Display information about the active management access list.
show management access-list [*name*]	Display information about the configured management ACL and its rules.

Management Access List Example

The commands in this example create a management ACL that permits access to the switch through the in-band switch ports on VLAN 1 and on port 9 from hosts with an IP address in the 10.27.65.0 subnet. Attempts to access the management interfaces from any other hosts and on any other interfaces is denied.

To configure the switch:

1. Create a management ACL and enter the configuration mode for the ACL.

```
console#configure
console(config)#management access-list mgmt_ACL
```

2. Create a rule that allows access from hosts in the 10.27.65.0 network on VLAN 1 and assign a priority of 1 to the rule.

```
console(config-macl)#permit ip-source 10.27.65.0
mask 255.255.255.0 vlan 1 priority 1
```

3. Create a rule that allows access from hosts in the 10.27.65.0 network on connected to port 9 and assign a priority of 2 to the rule.

```
console(config-macl)#permit ip-source 10.27.65.0
mask 255.255.255.0 Te1/0/9 priority 2
console(config-macl)#exit
```

4. Activate the ACL.

```
console(config)#management access-class mgmt_ACL
console(config)#exit
```

5. Verify the management ACL configuration.

```
console#show management access-list

mgmt_ACL
--------
permit ip-source 10.27.65.0 mask 255.255.255.0
vlan 1 priority 1
permit ip-source 10.27.65.0 mask 255.255.255.0
Te1/0/9 priority 2
! (Note: all other access implicitly denied)
```

6. Verify that the configured management ACL is in use.

```
console#show management access-class
Management access-class is enabled, using access
list mgmt_ACL.
```

22

Configuring VLANs

This chapter describes how to configure VLANs, including port-based VLANs, protocol-based VLANs, double-tagged VLANs, subnet-based VLANs, and Voice VLANs.

The topics covered in this chapter include:

- VLAN Overview
- Default VLAN Behavior
- Configuring VLANs (Web)
- Configuring VLANs (CLI)
- VLAN Configuration Examples

VLAN Overview

By default, all switchports on a PowerConnect 8000-series and 8100-series switches are in the same broadcast domain. This means when one host connected to the switch broadcasts traffic, every device connected to the switch receives that broadcast. All ports in a broadcast domain also forward multicast and unknown unicast traffic to the connected host. Large broadcast domains can result in network congestion, and end users might complain that the network is slow. In addition to latency, large broadcast domains are a greater security risk since all hosts receive all broadcasts.

Virtual Local Area Networks (VLANs) allow you to divide a broadcast domain into smaller, logical networks. Like a bridge, a VLAN switch forwards traffic based on the Layer 2 header, which is fast, and like a router, it partitions the network into logical segments, which provides better administration, security, and management of multicast traffic.

Network administrators have many reasons for creating logical divisions, such as department or project membership. Because VLANs enable logical groupings, members do not need to be physically connected to the same switch or network segment. Some network administrators use VLANs to segregate traffic by type so that the time-sensitive traffic, like voice traffic, has

priority over other traffic, such as data. Administrators also use VLANs to protect network resources. Traffic sent by authenticated clients might be assigned to one VLAN, while traffic sent from unauthenticated clients might be assigned to a different VLAN that allows limited network access.

When one host in a VLAN sends a broadcast, the switch forwards traffic only to other members of that VLAN. For traffic to go from a host in one VLAN to a host in a different VLAN, the traffic must be forwarded by a layer 3 device, such as a router. VLANs work across multiple switches, so there is no requirement for the hosts to be located near each other to participate in the same VLAN.

NOTE: PowerConnect 8000-series and 8100-series switches support VLAN routing. When you configure VLAN routing, the switch acts as a layer 3 device and can forward traffic between VLANs. For more information, see "What Are VLAN Routing Interfaces?" on page 843.

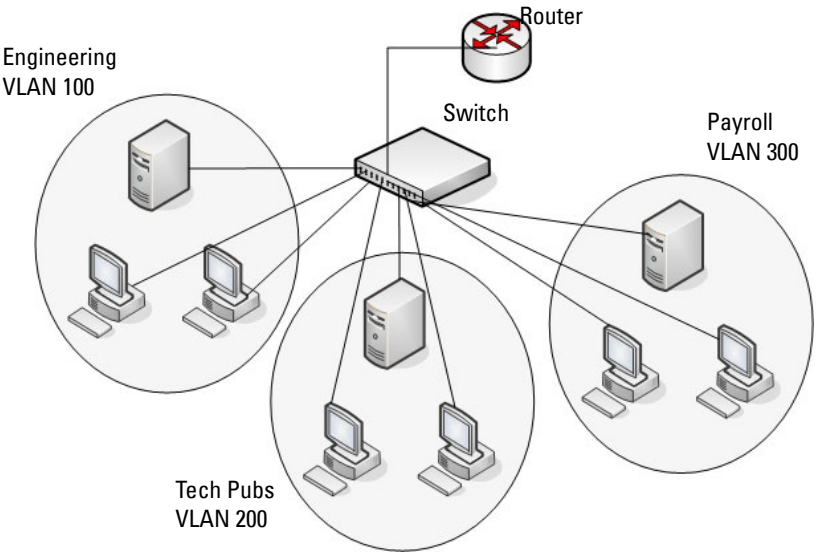
Each VLAN has a unique number, called the VLAN ID. The PowerConnect 8000-series and 8100-series switches support a configurable VLAN ID range of 2–4093. A VLAN with VLAN ID 1 is configured on the switch by default. VLAN 1 is named *default*, which cannot be changed. However, you can associate names with any other VLANs that you create.

In a tagged frame, the VLAN is identified by the VLAN ID in the tag. In an untagged frame, the VLAN identifier is the Port VLAN ID (PVID) specified for the port that received the frame. For information about tagged and untagged frames, see "VLAN Tagging" on page 543.

The PowerConnect 8000-series and 8100-series switches support adding individual ports and Link Aggregation Groups (LAGs) as VLAN members.

Figure 22-1 shows an example of a network with three VLANs that are department-based. The file server and end stations for the department are all members of the same VLAN.

Figure 22-1. Simple VLAN Topology



In this example, each port is manually configured so that the end station attached to the port is a member of the VLAN configured for the port. The VLAN membership for this network is port-based or static.

PowerConnect 8000-series and 8100-series switches also support VLAN assignment based on any of the following criteria:

- MAC address of the end station
- IP subnet of the end station
- Protocol of the packet transmitted by the end station

Table 22-1 provides an overview of the types of VLANs you can use to logically divide the network.

Table 22-1. VLAN Assignment

VLAN Assignment	Description
Port-based (Static)	This is the most common way to assign hosts to VLANs. The port where the traffic enters the switch determines the VLAN membership.
IP Subnet	Hosts are assigned to a VLAN based on their IP address. All hosts in the same subnet are members of the same VLAN.
MAC-Based	The MAC address of the device determines the VLAN assignment. This type of VLAN is useful when a host might not always connect to the network through the same port but needs to be on the same VLAN.
Protocol	Protocol-based VLANs were developed to separate traffic based on the protocol type before IP traffic became the de facto standard in the LAN. Use a protocol-based VLAN on networks where you might have a group of hosts that use IPX or another legacy protocol. With protocol-based VLANs, you can segregate traffic based on the EtherType value in the frame.

Switchport Modes

You can configure each port on a PowerConnect 8000-series and 8100-series switches to be in one of the following modes:

- Access — Access ports are intended to connect end-stations to the system, especially when the end-stations are incapable of generating VLAN tags.
- Trunk — Trunk-mode ports are intended for switch-to-switch links. Trunk ports can receive both tagged and untagged packets. Tagged packets received on a trunk port are forwarded on the VLAN contained in the tag. Untagged packets received on a trunk port are forwarded on the native VLAN. Packets received on another interface belonging to the native VLAN are transmitted untagged on a trunk port.
- General — General ports can be either access or trunk ports.

VLAN membership rules that apply to a port are based on the switchport mode configured for the port. Table 22-2 shows the behavior of the three switchport modes.

Table 22-2. Switchport Mode Behavior

Mode	VLAN Membership	Frames Accepted	Frames Sent	Ingress Filtering
Access	One VLAN	Untagged	Untagged	Always On
Trunk	All VLANs that exist in the system	Tagged	Tagged	Always On
General	As many as desired	Tagged or Untagged	Tagged or Untagged	On or Off

When a port is in General mode, all VLAN features are configurable. When ingress filtering is on, the frame is dropped if the port is not a member of the VLAN identified by the VLAN ID in the tag. If ingress filtering is off, all tagged frames are forwarded. The port decides whether to forward or drop the frame when the port receives the frame.

VLAN Tagging

PowerConnect 8000-series and 8100-series switches support IEEE 802.1Q tagging. Ethernet frames on a tagged VLAN have a 4-byte VLAN tag in the header. VLAN tagging is required when a VLAN spans multiple switches, which is why trunk ports transmit and receive only tagged frames.

Tagging may be required when a single port supports multiple devices that are members of different VLANs. For example, a single port might be connected to an IP phone, a PC, and a printer (the PC and printer are connected via ports on the IP phone). IP phones are typically configured to use a tagged VLAN for voice traffic, while the PC and printers typically use the untagged VLAN.

Trunk ports can receive tagged and untagged traffic. Untagged traffic is tagged internally with the native VLAN. Native VLAN traffic received untagged is transmitted untagged on a trunk port.

By default, trunk ports are members of all existing VLANs and will automatically participate in any newly created VLANs. The administrator can restrict the VLAN membership of a trunk port.

Access ports can receive untagged traffic and traffic tagged with the access port PVID.

GVRP

The GARP VLAN Registration Protocol (GVRP) helps to dynamically manage VLAN memberships on trunk ports. When GARP is enabled, switches can dynamically register (and de-register) VLAN membership information with other switches attached to the same segment.

Information about the active VLANs is propagated across all networking switches in the bridged LAN that support GVRP. You can configure ports to forbid dynamic VLAN assignment through GVRP.

The operation of GVRP relies upon the services provided by the Generic Attribute Registration Protocol (GARP). GVRP can create up to 1024 VLANs. For information about GARP timers, see "What Are GARP and GMRP?" on page 686.

Double-VLAN Tagging

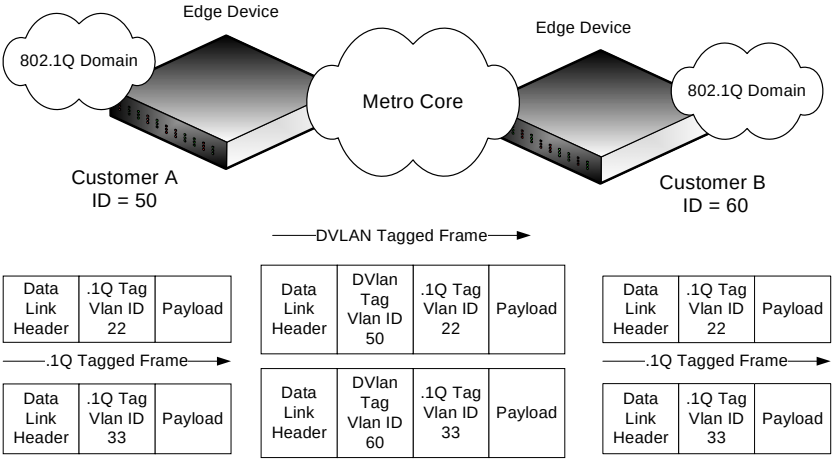
For trunk ports, which are ports that connect one switch to another switch, the PowerConnect 8000-series and 8100-series switches support double-VLAN tagging. This feature allows service providers to create Virtual Metropolitan Area Networks (VMANs). With double-VLAN tagging, service providers can pass VLAN traffic from one customer domain to another through a metro core in a simple and cost-effective manner. By using an additional tag on the traffic, the switch can differentiate between customers in the MAN while preserving an individual customer's VLAN identification when the traffic enters the customer's 802.1Q domain.

With the introduction of this second tag, customers are no longer required to divide the 4-byte VLAN ID space to send traffic on a Ethernet-based MAN. In short, every frame that is transmitted from an interface has a double-VLAN tag attached, while every packet that is received from an interface has a tag removed (if one or more tags are present).

In Figure 22-2, two customers share the same metro core. The service provider assigns each customer a unique ID so that the provider can distinguish between the two customers and apply different rules to each. When the configurable EtherType is assigned to something different than the 802.1Q (0x8100) EtherType, it allows the traffic to have added security from

misconfiguration while exiting the metro core. For example, if the edge device on the other side of the metro core is not stripping the second tag, the packet would never be classified as a 802.1Q tag, so the packet would be dropped rather than forwarded in the incorrect VLAN.

Figure 22-2. Double VLAN Tagging Network Example



Voice VLAN

The Voice VLAN feature enables switch ports to carry voice traffic with defined priority. When multiple devices, such as a PC and an IP phone, are connected to the same port, you can configure the port to use one VLAN for voice traffic and another VLAN for data traffic.

Voice over IP (VoIP) traffic is inherently time-sensitive: for a network to provide acceptable service, the transmission rate is vital. The priority level enables the separation of voice and data traffic coming onto the port.

A primary benefit of using Voice VLAN is to ensure that the sound quality of an IP phone is safeguarded from deteriorating when the data traffic on the port is high. The switch uses the source MAC address of the traffic traveling through the port to identify the IP phone data flow.

The Voice VLAN feature can be enabled on a per-port basis. This feature supports a configurable voice VLAN DSCP value. This value is later retrieved by LLDP when the LLDPDU is transmitted, if LLDP has been enabled on the port and the required TLV is configured for the port.

Identifying Voice Traffic

Some VoIP phones contain full support for IEEE 802.1X. When these phones are connected to a port that uses 802.1X port-based authentication, these phones authenticate and receive their VLAN information from LLDP-MED. However, if a VoIP phone has limited support for 802.1X authentication it might try to authenticate and fail. A phone with no 802.1X support would not attempt to authenticate at all. Instead of placing these phones on an unauthenticated or guest VLAN, the switch can automatically direct the VoIP traffic to the Voice VLAN without manual configuration.

The switch identifies the device as a VoIP phone by one of the following protocols:

- Cisco Discovery Protocol (CDP) or Industry Standard Discovery Protocol (ISDP) for Cisco VoIP phones
- DHCP vendor-specific options for Avaya VoIP phones
- LLDP-MED for most VoIP phones

NOTE: By default, ISDP is enabled globally and per-interface on the switch. LLDP-MED is disabled on each interface by default. Port-based authentication using 802.1X is also disabled on each port by default.

After the VoIP phone receives its VLAN information, all traffic is tagged with the VLAN ID of the Voice VLAN. The phone is considered to be authorized to send traffic but not necessarily authenticated.

Segregating Traffic with the Voice VLAN

You can configure the switch to support Voice VLAN on a port that is connecting the VoIP phone. Both of the following methods segregate the voice traffic and the data traffic in order to provide better service to the voice traffic.

- When a VLAN is associated with the Voice VLAN port, then the VLAN ID information is passed onto the VoIP phone using either the LLDP-MED or the CPD mechanism, depending on how the phone is identified: if it is identified via CDP, then the VLAN assignment is via CDP and if it is identified via LLDP-MED, then the VLAN assignment is via LLDP-MED. By this method, the voice data coming from the VoIP phone is tagged with the exchanged VLAN ID. Untagged data arriving on the switch is given the default PVID of the port, and the voice traffic is received tagged with the predefined VLAN. As a result, both kinds of traffic are segregated in order to provide better service to the voice traffic.
- When a dot1p priority is associated with the Voice VLAN port instead of a VLAN ID, then the priority information is passed onto the VoIP phone using the LLDP-MED or CDP mechanism. By this method, the voice data coming from the VoIP phone is tagged with VLAN 0 and with the exchanged priority; thus regular data arriving on the switch is given the default priority of the port (default 0), and the voice traffic is received with a higher priority.

You can configure the switch to override the data traffic CoS. This feature can override the 802.1 priority of the data traffic packets arriving at the port enabled for Voice VLAN. Therefore, any rogue client that is also connected to the Voice VLAN port does not deteriorate the voice traffic.

Voice VLAN and LLDP-MED

The interactions with LLDP-MED are important for Voice VLAN:

- LLDP-MED notifies the Voice VLAN component of the presence and absence of a VoIP phone on the network.
- The Voice VLAN component interacts with LLDP-MED for applying VLAN ID, priority, and tag information to the VoIP phone traffic.

Private VLANs

Private VLANs partition a standard VLAN domain into two or more subdomains. Each subdomain is defined by a primary VLAN and a secondary VLAN. The primary VLAN ID is the same for all subdomains that belong to a particular private VLAN instance. The secondary VLAN ID differentiates the subdomains from each other and provides layer 2 isolation between ports on the same private VLAN.

The following types of VLANs can be configured in a private VLAN:

- **Primary VLAN**—Forwards the traffic from the promiscuous ports to isolated ports, community ports and other promiscuous ports in the same private VLAN. Only one primary VLAN can be configured per private VLAN. All ports within a private VLAN share the same primary VLAN.
- **Isolated VLAN**—A secondary VLAN. It carries traffic from isolated ports to promiscuous ports. Only one isolated VLAN can be configured per private VLAN.
- **Community VLAN**—A secondary VLAN. It forwards traffic between ports which belong to the same community and to the promiscuous ports. There can be multiple community VLANs per private VLAN.

A port may be designated as one of the following types in a private VLAN:

- **Promiscuous port**—A port associated with a primary VLAN that is able to communicate with all interfaces in the private VLAN, including other promiscuous ports, community ports and isolated ports.
- **Host port**—A port associated with a secondary VLAN that can either communicate with the promiscuous ports in the VLAN and with other ports in the same community (if the secondary VLAN is a community VLAN) or can communicate only with the promiscuous ports (if the secondary VLAN is an isolated VLAN).

Private VLANs may be configured across a stack.

Private VLAN Usage Scenarios

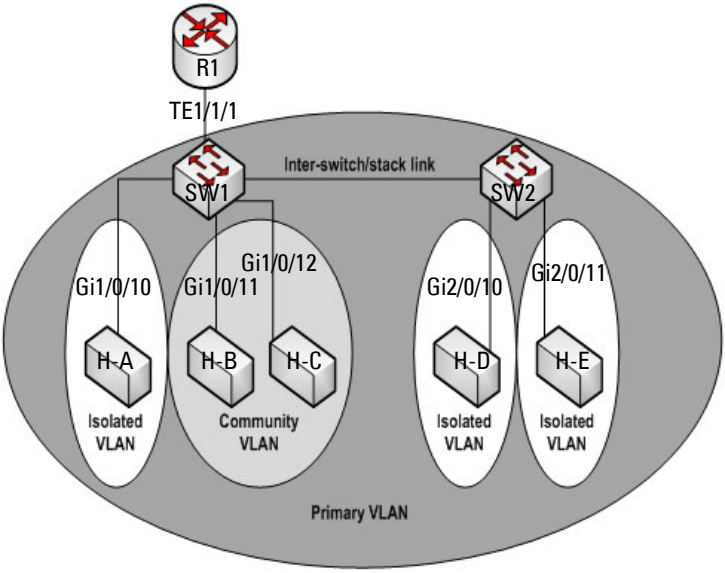
Private VLANs are typically implemented in a DMZ for security reasons. Servers in a DMZ are generally not allowed to communicate with each other but they must communicate to a router, through which they are connected to the users. Such servers are connected to host ports, and the routers are attached to promiscuous ports. Then, if one of the servers is compromised, the intruder cannot use it to attack another server in the same network segment.

The same traffic isolation can be achieved by assigning each port with a different VLAN, allocating an IP subnet for each VLAN, and enabling layer 3 routing between them. In a private VLAN domain, on the other hand, all members can share the common address space of a single subnet, which is

associated with a primary VLAN. So, the advantage of the private VLANs feature is that it reduces the number of consumed VLANs, improves IP addressing space utilization, and helps to avoid layer 3 routing.

Figure 22-3 shows an example Private VLAN scenario, in which five hosts (H-A through H-E) are connected to a stack of switches (SW1, SW2). The switch stack is connected to router R1. Port references shown are with reference to the stack.

Figure 22-3. Private VLAN Domain



Promiscuous Ports

An endpoint connected to a promiscuous port is allowed to communicate with any endpoint within the private VLAN. Multiple promiscuous ports can be defined for a single private VLAN domain.

In the configuration shown in Figure 22-3, the port connected from SW1 to R1 (TE1/1/1) is configured as a promiscuous port.

Isolated Ports

An endpoint connected to an isolated port is allowed to communicate with endpoints connected to promiscuous ports only. Endpoints connected to adjacent isolated ports cannot communicate with each other.

Community Ports

An endpoint connected to a community port is allowed to communicate with the endpoints within a community and can also communicate with any configured promiscuous port. The endpoints that belong to one community cannot communicate with endpoints that belong to a different community, or with endpoints connected to isolated ports.

Private VLAN Operation in the Switch Stack and Inter-switch Environment

The Private VLAN feature is supported in a stacked switch environment. The stack links are transparent to the configured VLANs; thus, there is no need for special private VLAN configuration beyond what would be configured for a single switch. Any private VLAN port can reside on any stack member.

To enable private VLAN operation across multiple switches that are not stacked, trunk ports must be configured between the switches to transport the private VLANs. The trunk ports must be configured with the promiscuous, isolated, and community VLANs. Trunk ports must also be configured on all devices separating the switches.

In regular VLANs, ports in the same VLAN switch traffic at L2. However, for a private VLAN, the promiscuous port forwards received traffic to secondary ports in the VLAN (isolated and community). Community ports forward received traffic to the promiscuous ports and other community ports using the same secondary VLAN. Isolated ports transmit received traffic to the promiscuous ports only.

The ports to which the broadcast traffic is forwarded depend on the type of port on which the traffic was received. If the received port is a host port, traffic is broadcast to all promiscuous and trunk ports. If the received port is a community port, the broadcast traffic is forwarded to all promiscuous, trunk, and community ports in the same secondary VLAN. A promiscuous port broadcasts traffic to other promiscuous ports, isolated ports, and community ports.

Table 22-3. Forwarding Rules for Traffic in Primary VLAN

	To				
From	**promiscuous**	**community 1**	**community 2**	**isolated**	**stack (trunk)**
promiscuous	allow	allow	allow	allow	allow
community 1	N/A	N/A	N/A	N/A	N/A
community 2	N/A	N/A	N/A	N/A	N/A
isolated	N/A	N/A	N/A	N/A	N/A
stack (trunk)	allow	allow	allow	allow	allow

Table 22-4. Forwarding Rules for Traffic in Community 1 VLAN

	To				
From	**promiscuous**	**community 1**	**community 2**	**isolated**	**stack (trunk)**
promiscuous	N/A	N/A	N/A	N/A	N/A
community 1	allow	allow	deny	deny	allow
community 2	N/A	N/A	N/A	N/A	N/A
isolated	N/A	N/A	N/A	N/A	N/A
stack (trunk)	allow	allow	deny	deny	allow

Table 22-5. Forwarding Rules for Traffic in Isolated VLAN

	To				
From	**promiscuous**	**community 1**	**community 2**	**isolated**	**stack (trunk)**
promiscuous	N/A	N/A	N/A	N/A	N/A
community 1	N/A	N/A	N/A	N/A	N/A
community 2	N/A	N/A	N/A	N/A	N/A
isolated	allow	deny	deny	deny	allow
stack (trunk)	allow	deny	deny	deny	Allow

Limitations and Recommendations

- Only a single isolated VLAN can be associated with a primary VLAN. Multiple community VLANs can be associated with a primary VLAN.
- Trunk and general modes are not supported on private VLAN ports.
- Do not configure access ports using the VLANs participating in any of the private VLANs.
- Multiple primary VLANs may be configured. Each primary VLAN must be unique and each defines a separate private VLAN domain. The operator must take care to use only the secondary VLANs associated with the primary VLAN of a domain.
- Private VLANs cannot be enabled on a preconfigured interface. The interface must physically exist in the switch.
- Secondary (community and isolated) VLANS are associated to the same multiple spanning tree instance as the primary VLAN.
- GVRP/MVRP cannot be enabled after the private VLAN is configured. The administrator will need to disable both before configuring the private VLAN.
- DHCP snooping can be configured on the primary VLAN. If it is enabled for a secondary VLAN, the configuration does not take effect if a primary VLAN is already configured.
- If IP source guard is enabled on private VLAN ports, then DHCP snooping must be enabled on the primary VLAN.
- Do not configure private VLAN ports on interfaces configured for voice VLAN.
- If static MAC addresses are added for the host port, the same static MAC address entry must be added to the associated primary VLAN. This does not need to be replicated for dynamic MAC addresses.
- A private VLAN cannot be enabled on a management VLAN.
- A private VLAN cannot be enabled on the default VLAN.
- VLAN routing can be enabled on private VLANs. It is not very useful to enable routing on secondary VLANs, as the access to them is restricted. However, primary VLANs can be enabled for routing.

- It is recommended that the private VLAN IDs be removed from the trunk ports connected to devices that do not participate in the private VLAN traffic.

Private VLAN Configuration Example

See "Configuring a Private VLAN" on page 602.

Additional VLAN Features

The PowerConnect 8000-series and 8100-series switches also support the following VLANs and VLAN-related features:

- VLAN routing interfaces — See "Configuring Routing Interfaces" on page 843.
- Guest VLAN — See "Configuring Port and System Security" on page 457.

Default VLAN Behavior

One VLAN exists on the PowerConnect 8000-series and 8100-series switches by default. The VLAN ID is 1, and all ports are included in the VLAN as access ports, which are untagged. This means when a device connects to any port on the switch, the port forwards the packets without inserting a VLAN tag. If a device sends a tagged frame to a port, the frame is dropped. Since all ports are members of this VLAN, all ports are in the same broadcast domain and receive all broadcast and multicast traffic received on any port.

When you add a new VLAN to the VLAN database, no ports are members. The configurable VLAN range is 2–4093. VLANs 4094 and 4095 are reserved.

Ports in trunk and access mode have the default behavior shown in Table 22-2 and cannot be configured with different tagging or ingress filtering values. When you add a VLAN to a port in general mode, the VLAN has the behavior shown in Table 22-6.

Table 22-6. General mode Default Settings

Feature	Default Value
Frames accepted	Untagged
	Incoming untagged frames are classified into the VLAN whose VLAN ID is the currently configured PVID.
Frames sent	Untagged
Ingress Filtering	On
PVID	1

Table 22-7 shows the default values or maximum values for VLAN features.

Table 22-7. Additional VLAN Default and Maximum Values

Feature	Value
Default VLAN	VLAN 1
VLAN Name	No VLAN name is configured except for VLAN 1, whose name "default" cannot be changed.
VLAN Range	2–4093
Switchport mode	Access
Double-VLAN tagging	Disabled If double-VLAN tagging is enabled, the default EtherType value is 802.1Q
Maximum number of configurable MAC-to-VLAN bindings	128
Maximum number of configurable IP Subnet-to-VLAN bindings	64
GVRP	Disabled If GVRP is enabled, the default port parameters are: • GVRP State: Disabled • Dynamic VLAN Creation: Disabled • GVRP Registration: Disabled
Number of dynamic VLANs that can be assigned through GVRP	1024
Voice VLAN	Disabled
Voice VLAN DSCP value	46
Voice VLAN authentication mode	Enabled

Configuring VLANs (Web)

This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring VLANs on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

VLAN Membership

Use the **VLAN Membership** page to create VLANs and define VLAN groups stored in the VLAN membership table.

To display the **VLAN Membership** page, click **Switching** → **VLAN** → **VLAN Membership** in the navigation panel.

The **VLAN Membership** tables display which Ports and LAGs are members of the VLAN, and whether they're tagged (T), untagged (U), or forbidden (F). The tables have two rows: **Static** and **Current**. Only the **Static** row is configurable. The **Current** row is updated either dynamically through GVRP or when the **Static** row is changed and **Apply** is clicked.

There are two tables on the page:

- **Ports** — Displays and assigns VLAN membership to ports. To assign membership, click in **Static** for a specific port. Each click toggles between U, T, and blank. See Table 22-8 for definitions.
- **LAGs** — Displays and assigns VLAN membership to LAGs. To assign membership, click in **Static** for a specific LAG. Each click toggles between U, T, and blank. See Table 22-8 for definitions.

Table 22-8. VLAN Port Membership Definitions

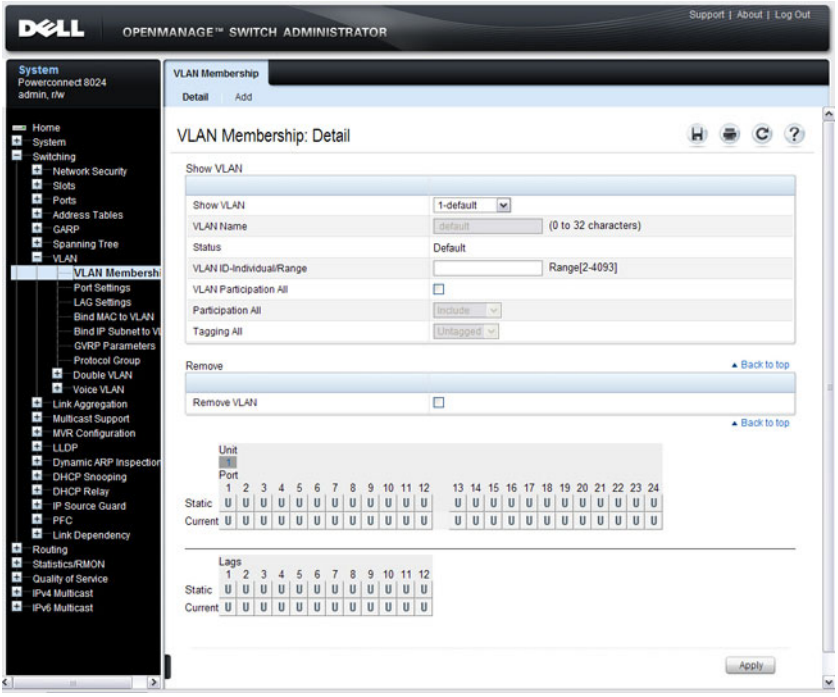
Port Control	Definition
T	Tagged: the interface is a member of a VLAN. All packets forwarded by the interface in this VLAN are tagged. The packets contain VLAN information.
U	Untagged: the interface is a VLAN member. Packets forwarded by the interface in this VLAN are untagged.
F	Forbidden: indicates that the interface is forbidden from becoming a member of the VLAN. This setting is primarily for GVRP, which enables dynamic VLAN assignment.

Table 22-8. VLAN Port Membership Definitions

Port Control	Definition
Blank	Blank: the interface is not a VLAN member. Packets in this VLAN are not forwarded on this interface.

To perform additional port configuration, such as making the port a trunk port, use the **Port Settings** page.

Figure 22-4. VLAN Membership



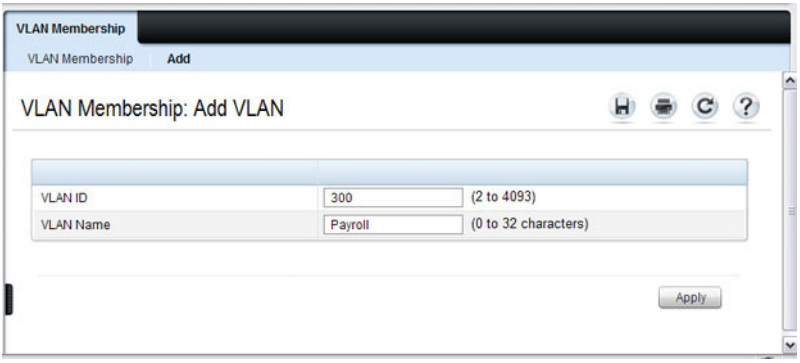
Adding a VLAN

To create a VLAN:

1. Open the **VLAN Membership** page.
2. Click **Add** to display the **Add VLAN** page.

3 Specify a VLAN ID and a VLAN name.

Figure 22-5. Add VLAN



4 Click **Apply**.

Configuring Ports as VLAN Members

To add member ports to a VLAN:

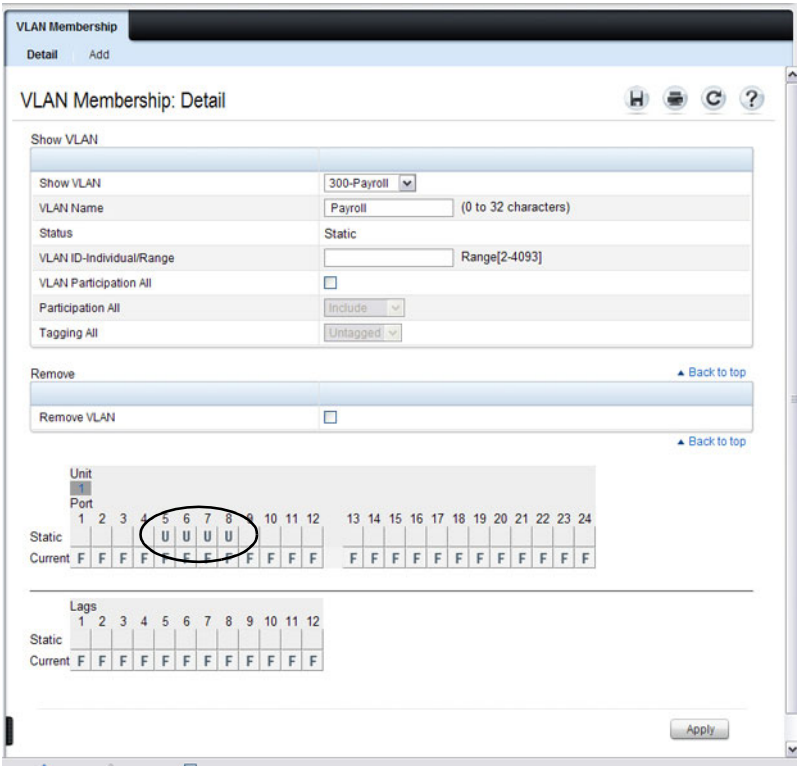
1 Open the **VLAN Membership** page.

2 From the **Show VLAN** menu, select the VLAN to which you want to assign ports.

3 In the **Static** row of the **VLAN Membership** table, click the blank field to assign the port as an untagged member.

Figure 22-6 shows 10-Gigabit Ethernet ports 5–8 being added to VLAN 300.

Figure 22-6. Add Ports to VLAN

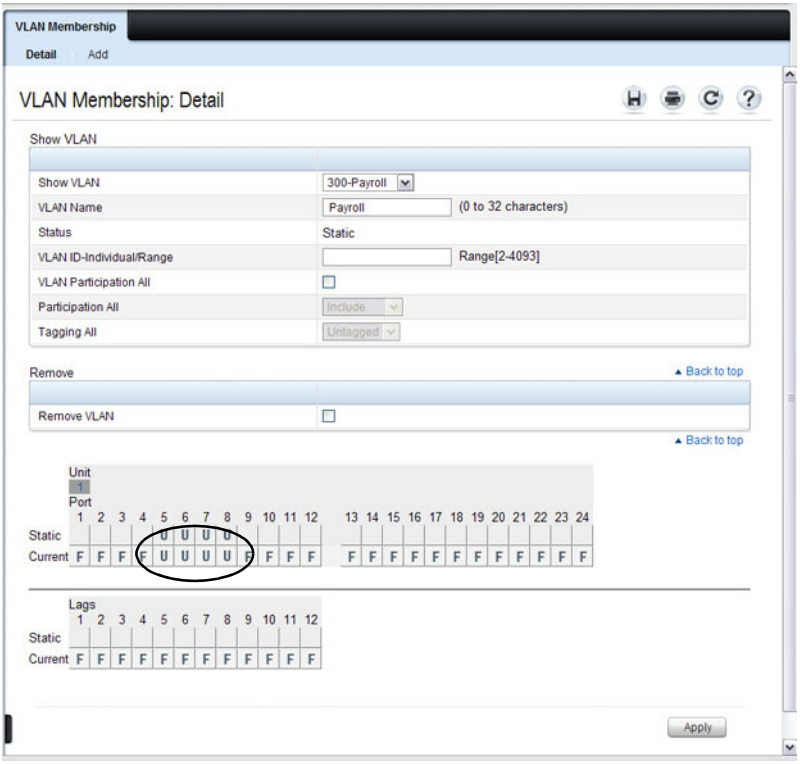


4 Click **Apply**.

5 Verify that the ports have been added to the VLAN.

In Figure 22-7, the presence of the letter **U** in the **Current** row indicates that the port is an untagged member of the VLAN.

Figure 22-7. Add Ports to VLAN



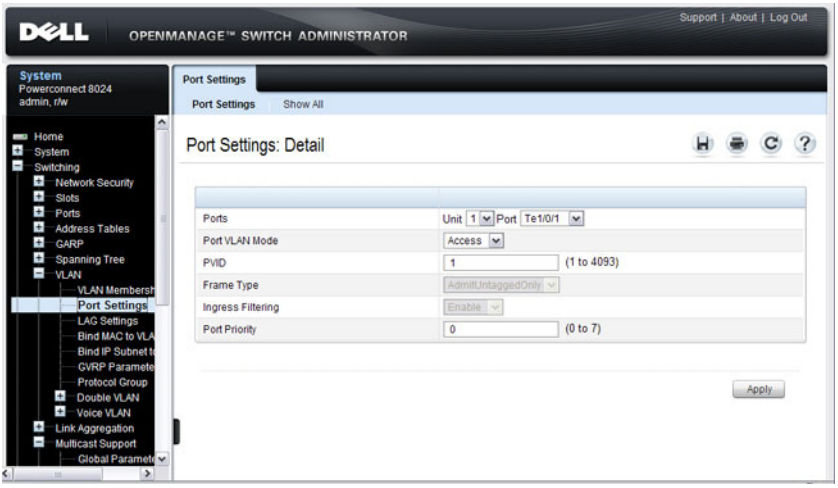
VLAN Port Settings

Use the **VLAN Port Settings** page to add ports to an existing VLAN and to configure settings for the port. If you select Trunk or Access as the **Port VLAN Mode**, some of the fields are not configurable because of the requirements for that mode.

NOTE: You can add ports to a VLAN through the table on the **VLAN Membership** page or through the **PVID** field on the **Port Settings** page. The PVID is the VLAN that untagged received packets are assigned to. To include a general-mode port in multiple VLANs, use the **VLAN Membership** page.

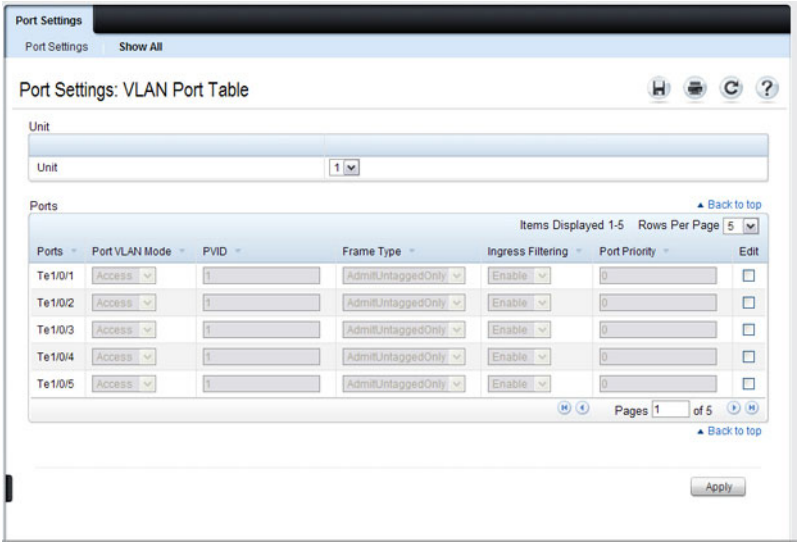
To display the **Port Settings** page, click **Switching** → **VLAN** → **Port Settings** in the navigation panel.

Figure 22-8. VLAN Port Settings



From the **Port Settings** page, click **Show All** to see the current VLAN settings for all ports. You can change the settings for one or more ports by clicking the **Edit** option for a port and selecting or entering new values.

Figure 22-9. VLAN Settings for All Ports

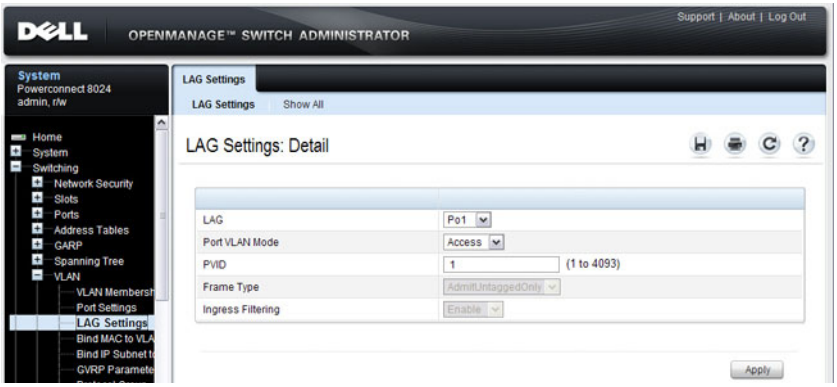


VLAN LAG Settings

Use the **VLAN LAG Settings** page to map a LAG to a VLAN and to configure specific VLAN settings for the LAG.

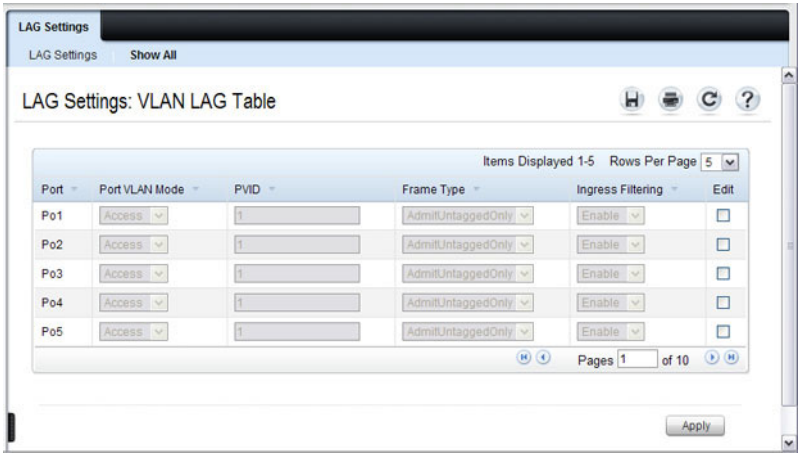
To display the **LAG Settings** page, click **Switching → VLAN → LAG Settings** in the navigation panel.

Figure 22-10. VLAN LAG Settings



From the **LAG Settings** page, click **Show All** to see the current VLAN settings for all LAGs. You can change the settings for one or more LAGs by clicking the **Edit** option for a port and selecting or entering new values.

Figure 22-11. VLAN LAG Table

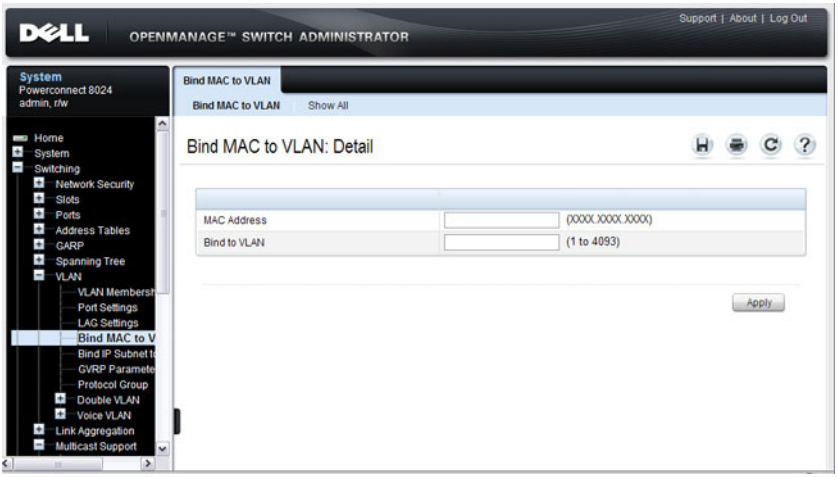


Bind MAC to VLAN

Use the **Bind MAC to VLAN** page to map a MAC address to a VLAN. After the source MAC address and the VLAN ID are specified, the MAC to VLAN configurations are shared across all ports of the switch. The MAC to VLAN table supports up to 128 entries.

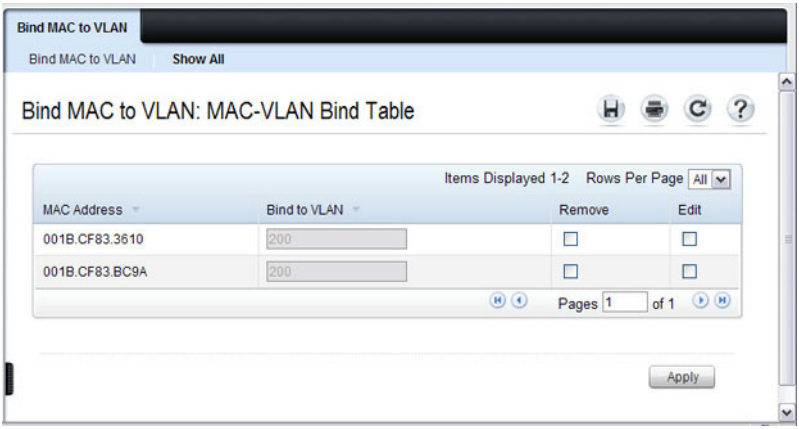
To display the **Bind MAC to VLAN** page, click **Switching → VLAN → Bind MAC to VLAN** in the navigation panel.

Figure 22-12. Bind MAC to VLAN



From the **Bind MAC to VLAN** page, click **Show All** to see the MAC addresses that are mapped to VLANs. From this page, you can change the settings for one or more entries or remove an entry.

Figure 22-13. MAC-VLAN Bind Table

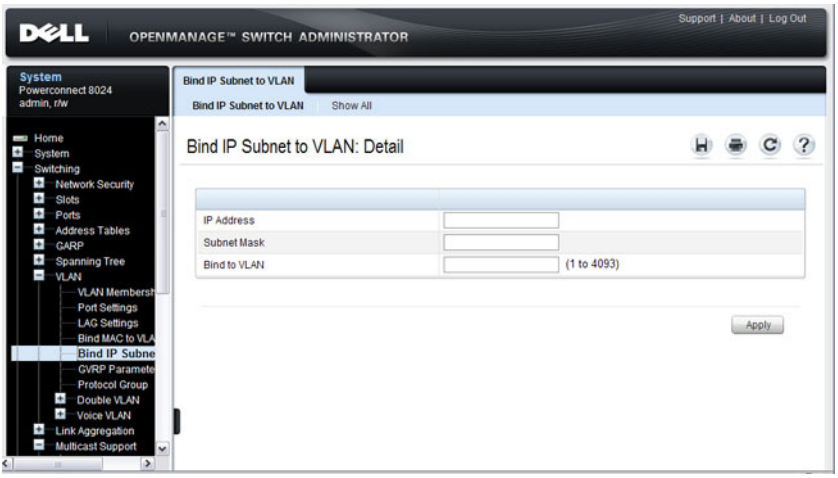


Bind IP Subnet to VLAN

Use the **Bind IP Subnet to VLAN** page to assign an IP Subnet to a VLAN. The IP Subnet to VLAN configurations are shared across all ports of the switch. There can be up to 64 entries configured in this table.

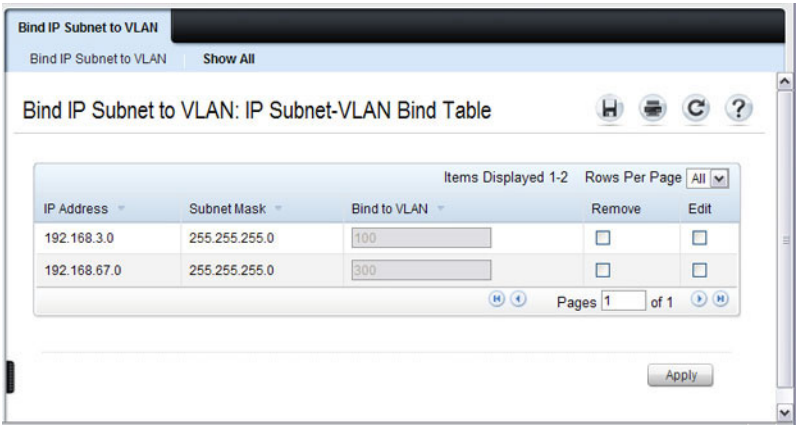
To display the **Bind IP Subnet to VLAN** page, click **Switching** → **VLAN** → **Bind IP Subnet to VLAN** in the navigation panel.

Figure 22-14. Bind IP Subnet to VLAN



From the **Bind IP Subnet to VLAN** page, click **Show All** to see the IP subnets that are mapped to VLANs. From this page, you can change the settings for one or more entries or remove an entry.

Figure 22-15. Subnet-VLAN Bind Table

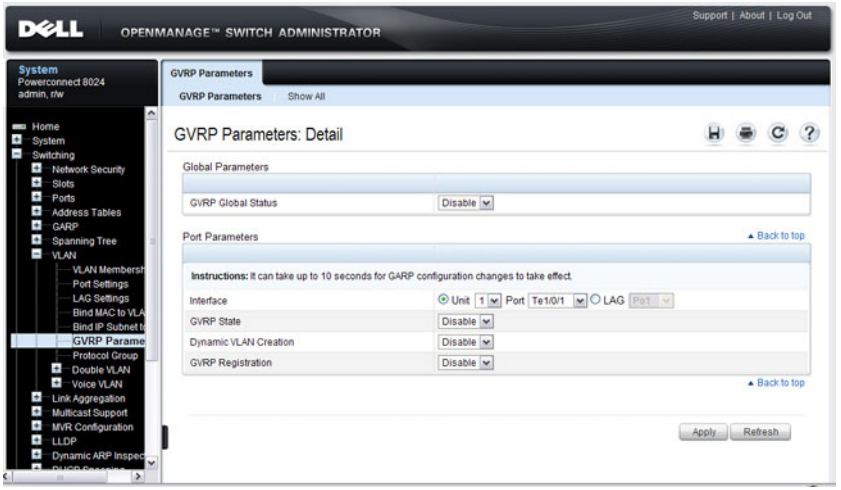


GVRP Parameters

Use the **GVRP Parameters** page to enable GVRP globally and configure the port settings.

To display the **GVRP Parameters** page, click **Switching → VLAN → GVRP Parameters** in the navigation panel.

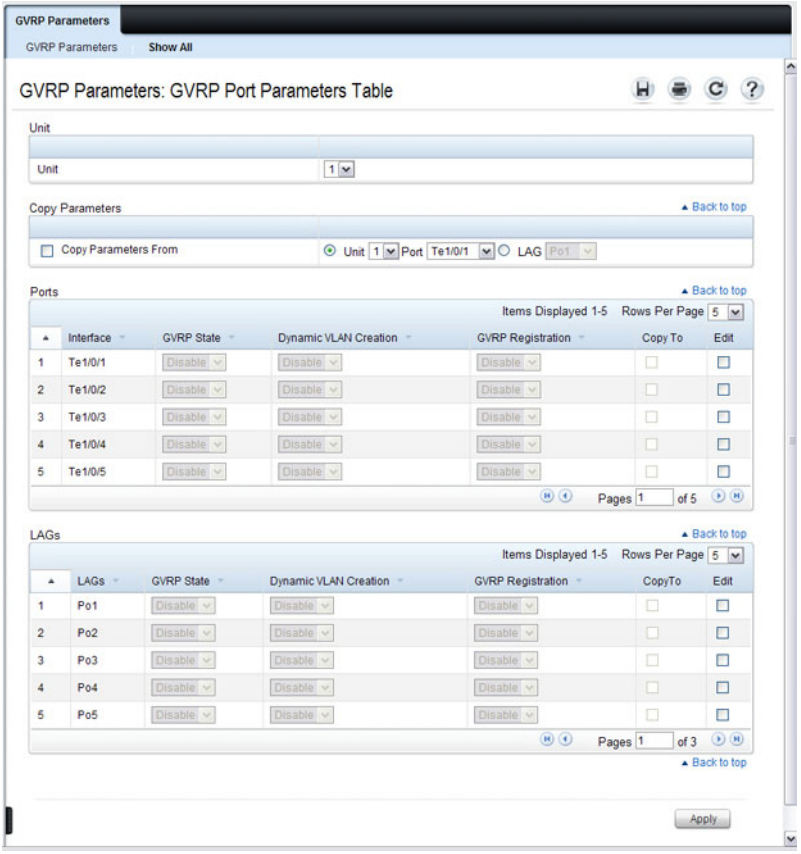
Figure 22-16. GVRP Parameters



From the **GVRP Parameters** page, click **Show All** to see the GVRP configuration for all ports. From this page, you can change the settings for one or more entries.

NOTE: Per-port and per-LAG GVRP Statistics are available from the Statistics/RMON page. For more information, see "Monitoring Switch Traffic" on page 347.

Figure 22-17. GVRP Port Parameters Table

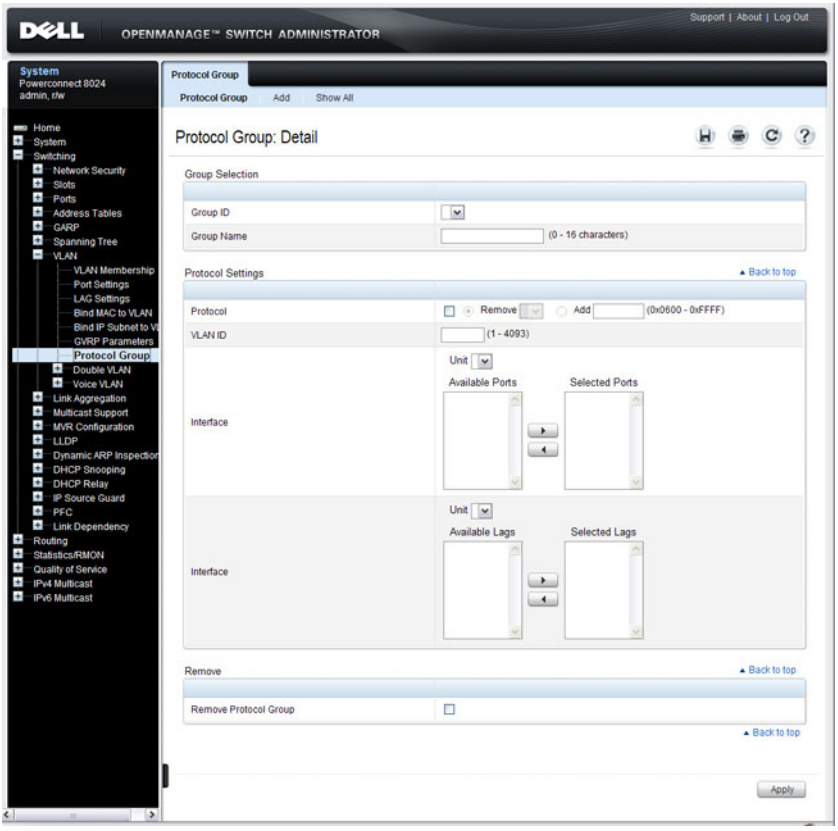


Protocol Group

Use the **Protocol Group** page to configure which EtherTypes go to which VLANs, and then enable certain ports to use these settings. Protocol-based VLANs are most often used in situations where network segments contain hosts running multiple protocols.

To display the **Protocol Group** page, click **Switching** → **VLAN** → **Protocol Group** in the navigation panel.

Figure 22-18. Protocol Group

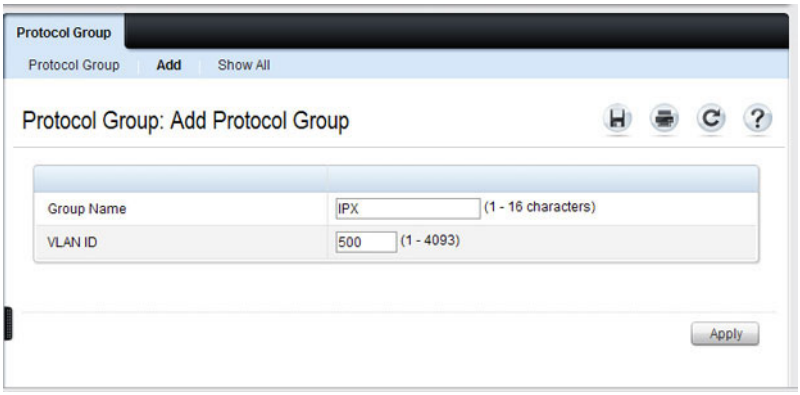


Adding a Protocol Group

To add a protocol group:

1. Open the **Protocol Group** page.
2. Click **Add** to display the **Add Protocol Group** page.
3. Create a name for the group and associate a VLAN with the group.

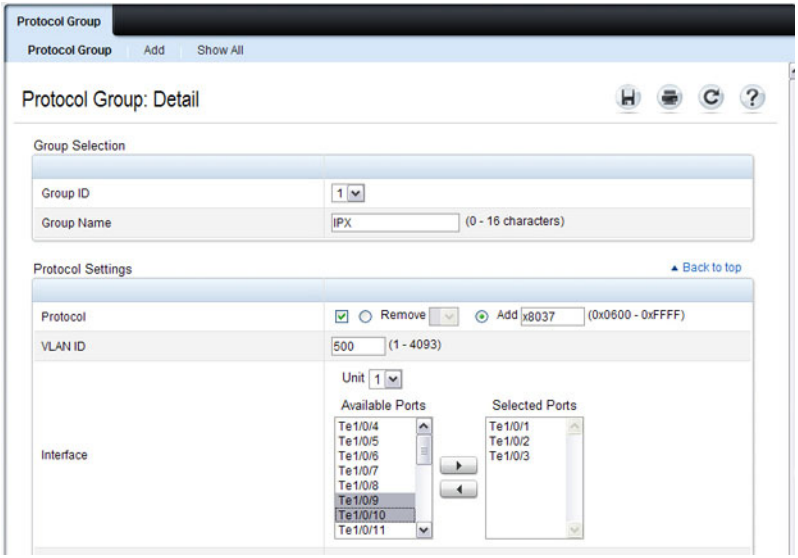
Figure 22-19. Add Protocol Group



4. Click **Apply**.
5. Click **Protocol Group** to return to the main **Protocol Group** page.
6. From the **Group ID** field, select the group to configure.
7. In the **Protocol Settings** table, select the protocol and interfaces to associate with the protocol-based VLAN.

   In Figure 22-20, the Protocol Group 1 (named IPX) is associated with the IPX protocol and ports 14–16. Ports 20-22 are selected in **Available Ports** list. After clicking the right arrow, they will be added to the **Selected Ports** list.

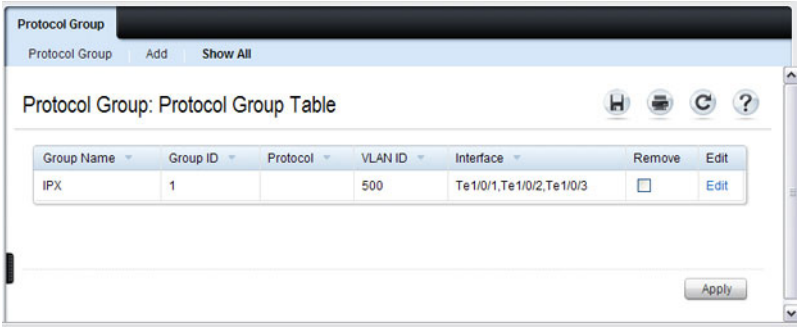
Figure 22-20. Configure Protocol Group



8 Click **Apply**.

9 Click **Show All** to see the protocol-based VLANs and their members.

Figure 22-21. Protocol Group Table

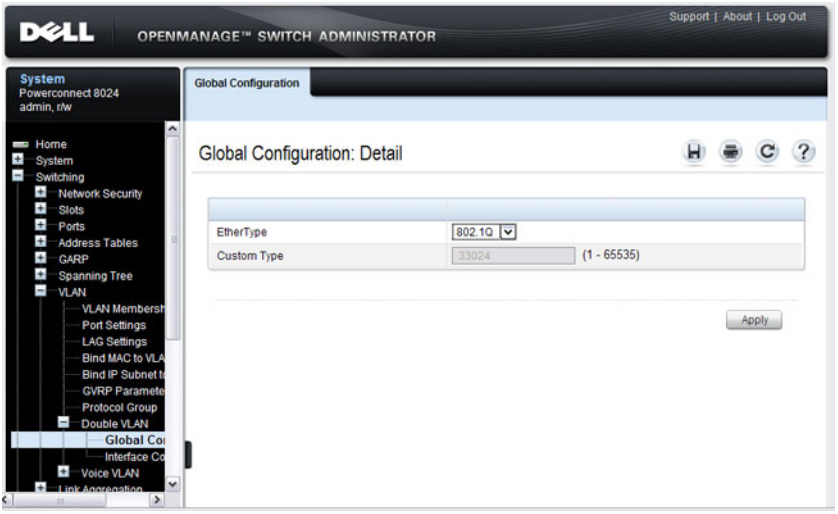


Double VLAN Global Configuration

Use the **Double VLAN Global Configuration** page to specify the value of the **EtherType** field in the first EtherType/tag pair of the double-tagged frame.

To display the **Double VLAN Global Configuration** page, click **Switching** → **VLAN** → **Double VLAN** → **Global Configuration** in the navigation panel.

Figure 22-22. Double VLAN Global Configuration

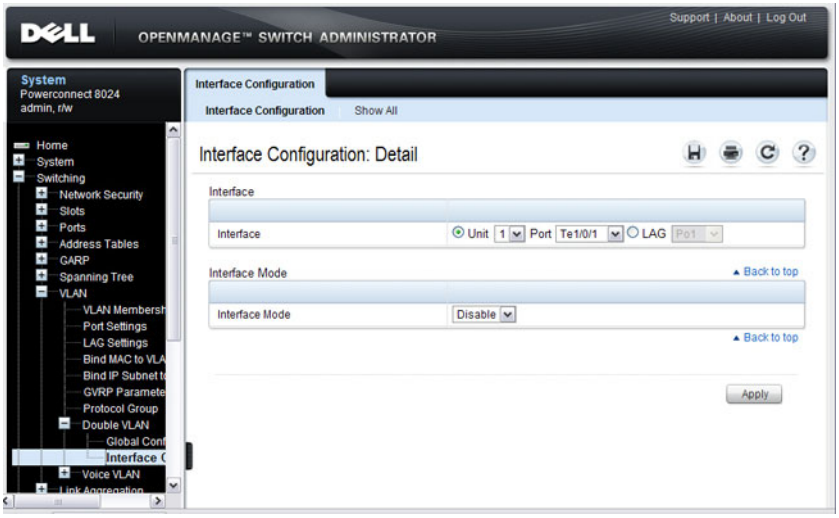


Double VLAN Interface Configuration

Use the **Double VLAN Interface Configuration** page to specify the value of the **EtherType** field in the first EtherType/tag pair of the double-tagged frame.

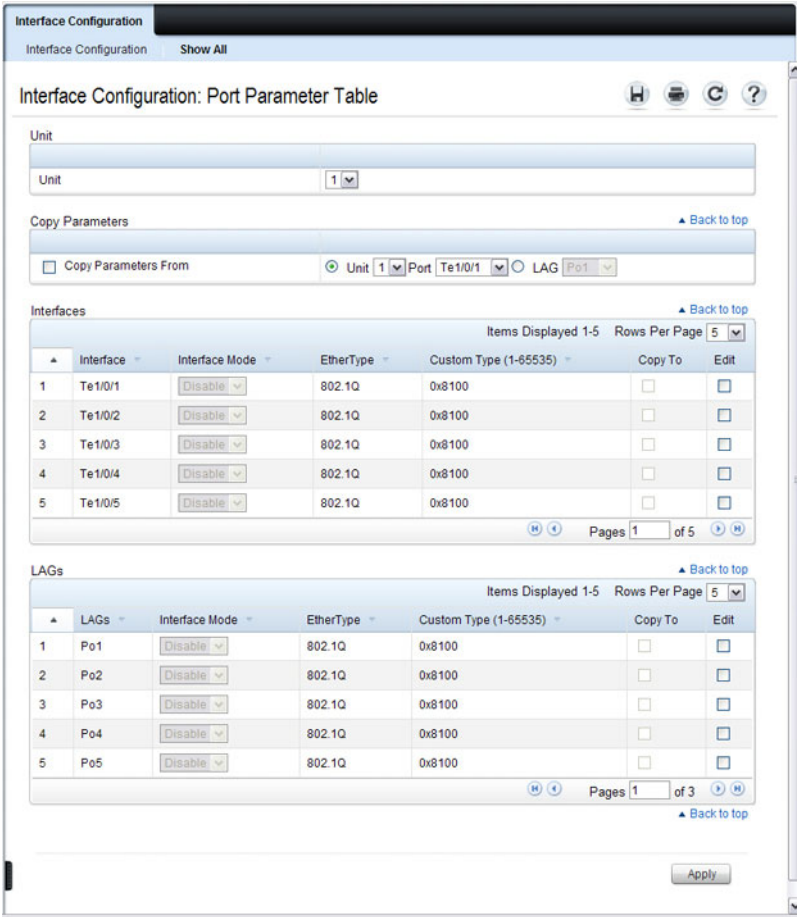
To display the **Double VLAN Interface Configuration** page, click **Switching** → **VLAN** → **Double VLAN** → **Interface Configuration** in the navigation panel.

Figure 22-23. Double VLAN Interface Configuration



To view a summary of the double VLAN configuration for all interfaces and to edit settings for one or more interfaces, click **Show All**.

Figure 22-24. Double VLAN Port Parameter Table

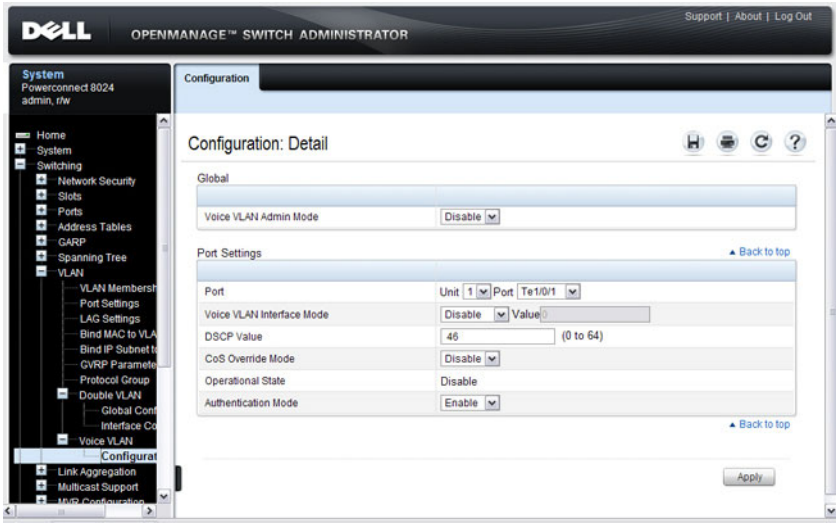


Voice VLAN

Use the **Voice VLAN Configuration** page to configure and view voice VLAN settings that apply to the entire system and to specific interfaces.

To display the page, click **Switching → VLAN → Voice VLAN → Configuration** in the navigation panel.

Figure 22-25. Voice VLAN Configuration



NOTE: IEEE 802.1X must be enabled on the switch before you disable voice VLAN authentication. Voice VLAN authentication can be disabled in order to allow VoIP phones that do not support authentication to send and receive unauthenticated traffic on the Voice VLAN.

Configuring VLANs (CLI)

This section provides information about the commands you use to create and configure VLANs. For more information about the commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Creating a VLAN

Beginning in Privileged EXEC mode, use the following commands to configure a VLAN and associate a name with the VLAN.

Command	Purpose
configure	Enter global configuration mode.
vlan {*vlan-id* | *vlan-range*}	Create a new VLAN or a range of VLANs and enter the interface configuration mode for the specified VLAN or VLAN range. • *vlan-id*—A valid VLAN IDs (Range: 2–4093). • *vlan-range* — A list of valid VLAN IDs to be added. List separate, non-consecutive VLAN IDs separated by commas (without spaces); use a hyphen to designate a range of IDs. (Range: 2–4093) **NOTE:** You can also create with this command in VLAN Database mode . To enter VLAN Database mode, use the **vlan database** command in Global Configuration mode.
name *string*	Add a name to the specified VLAN. *string* — Comment or description to help identify a specific VLAN (Range: 1–32 characters).
CTRL + Z	Exit to Privileged EXEC mode.
show vlan [**id** *vlan-id* | **name** *vlan-name*]	Display VLAN information. • *vlan-id* — A valid VLAN ID. (Range: 1–4093) • *vlan-name* — A valid VLAN name string. (Range: 1–32 characters)

Configuring a Port in Access Mode

Beginning in Privileged EXEC mode, use the following commands to configure an untagged layer 2 VLAN interface and assign the interface to a VLAN. When a port is in access mode, it can only be a member of one untagged VLAN. When you configure the interface as a VLAN member, the interface is automatically removed from its previous VLAN membership. You can configure each interface separately, or you can configure a range of interfaces with the same settings.

Command	Purpose
configure	Enter global configuration mode.
interface *interface*	Enter interface configuration mode for the specified interface. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**.
	You can also specify a range of interfaces with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12.
switchport mode access	Configure the interface as an untagged layer 2 VLAN interface.
switchport access vlan *vlan-id*	Configure the interface as a member of the specified VLAN.
	vlan-id — A valid VLAN ID of the VLAN to which the port is configured. (Range: 1–4093)
CTRL + Z	Exit to Privileged EXEC mode.
show interfaces switchport *interface*	Display information about the VLAN settings configured for the specified interface.

Configuring a Port in General Mode

Beginning in Privileged EXEC mode, use the following commands to configure an interface with full 802.1q support and configure the VLAN membership information for the interface.

Command	Purpose
configure	Enter global configuration mode.
interface *interface*	Enter interface configuration mode for the specified interface. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**. You can also specify a range of interfaces with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12.
switchport mode general	Configure the interface as an untagged layer 2 VLAN interface.
switchport general allowed vlan [add|remove] *vlan-list* **{tagged|untagged}**	Configure the VLAN membership for the port. You can also use this command to change the egress tagging for packets without changing the VLAN assignment. • **add** *vlan-list* — List of VLAN IDs to add. Separate nonconsecutive VLAN IDs with a comma and no spaces. Use a hyphen to designate a range of IDs. (Range: 1–4093) • **remove** *vlan-list* — List of VLAN IDs to remove. Separate nonconsecutive VLAN IDs with a comma and no spaces. Use a hyphen to designate a range of IDs. • **tagged** — Sets the port to transmit tagged packets for the VLANs. If the port is added to a VLAN without specifying tagged or untagged, the default is untagged. • **untagged** — Sets the port to transmit untagged packets for the VLANs.
switchport general pvid *vlan-id*	(Optional) Set the port VLAN ID. Untagged traffic that enters the switch through this port is tagged with the PVID. *vlan-id* — PVID. The VLAN ID may mot belong to a VLAN that does not exist on the switch. (Range: 1–4093)

Command	Purpose
switchport general acceptable-frame-type tagged-only	(Optional) Specifies that the port will only accept tagged frames. Untagged frames are dropped at ingress.
switchport general ingress-filtering disable	(Optional) Turn off ingress filtering so that all received tagged frames are forwarded whether or not the port is a member of the VLAN in the tag.
CTRL + Z	Exit to Privileged EXEC mode.
show interfaces switchport *interface*	Display information about the VLAN settings configured for the specified interface. The *interface* variable includes the interface type and number.

Configuring a Port in Trunk Mode

Beginning in Privileged EXEC mode, use the following commands to configure an interface as a layer 2 trunking interface, which connects two switches.

Command	Purpose
configure	Enter global configuration mode.
interface *interface*	Enter interface configuration mode for the specified interface. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**. You can also specify a range of interfaces with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12.
switchport mode trunk	Configure the interface as an untagged layer 2 VLAN interface.

Command	Purpose
switchport trunk {allowed vlan *vlan-list*|**native vlan** *vlan-id*}	Set the list of allowed VLANs that can receive and send traffic on this interface in tagged format when in trunking mode. • **allowed** *vlan-list* — Set the list of allowed VLANs that can receive and send traffic on this interface in tagged format when in trunking mode. Separate non-consecutive VLAN IDs with a comma and no spaces. Use a hyphen to designate a range of IDs. The *vlan-list* format is **all** | [**add** | **remove** | **except**] *vlan-atom* [*vlan-atom*...] where: • **all**—Specifies all VLANs from 1 to 4093. This keyword is not allowed on commands that do not permit all VLANs in the list to be set at the same time. • **add**—Adds the defined list of VLANs to those currently set instead of replacing the list. • **remove**—Removes the defined list of VLANs from those currently set instead of replacing the list. Valid IDs are from 1 to 1005; extended-range VLAN IDs are valid in. • **except**—Lists the VLANs that should be calculated by inverting the defined list of VLANs. (VLANs are added except the ones specified.) . • *vlan-atom* —Either a single VLAN number from 1 to 4093 or a continuous range of VLANs described by two VLAN numbers, the lesser one first, separated by a hyphen. • **native** *vlan-id*— The untagged VLAN. Untagged packets received on this interface are switched in this VLAN. Transmitted packets in this VLAN are sent untagged.
CTRL + Z	Exit to Privileged EXEC mode.
show interfaces switchport *interface*	Display information about the VLAN settings configured for the specified interface. The *interface* variable includes the interface type and number.

Configuring VLAN Settings for a LAG

The VLAN mode and memberships settings you configure for a port are also valid for a LAG (port channel). Beginning in Privileged EXEC mode, use the following commands to configure the VLAN mode for a LAG. Once you specify the switchport mode settings for a LAG, you can configure other VLAN memberships settings that are valid that the switchport mode.

Command	Purpose
configure	Enter global configuration mode.
interface port-channel *channel-id*	Enter interface configuration mode for the specified interface. *channel-id* — Specific port channel. (Range 1–48). You can also specify a range of LAGs with the **interface range port-channel** command, for example, **interface range port-channel 4-**8.
switchport mode [access|general|trunk]	Configure the interface as an untagged layer 2 VLAN interface.
CTRL + Z	Exit to Privileged EXEC mode.
show interfaces switchport port-channel *channel-id*	Display information about the VLAN settings configured for the specified LAG.

Configuring Double VLAN Tagging

Beginning in Privileged EXEC mode, use the following commands to configure an interface to send and accept frames with double VLAN tagging.

Command	Purpose
configure	Enter global configuration mode.
interface *interface*	Enter interface configuration mode for the specified interface. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**. You can also specify a range of interfaces with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12.
mode dvlan-tunnel	Enable Double VLAN Tunneling on the specified interface.
exit	Exit to global configuration mode
dvlan-tunnel ethertype {**802.1Q** | **vman** | **custom** <*0-65535*>} [primary-tpid]	Configure the EtherType to use for interfaces with double VLAN tunneling enabled. • 802.1Q — Configures the EtherType as 0x8100. • vman — Configures the EtherType as 0x88A8. • custom — Custom configures the EtherType for the DVLAN tunnel. The value must be 0-65535. • primary-tpid — Configure the primary (outer) TPID. If this parameter is not present, the inner TPID is configured.
CTRL + Z	Exit to Privileged EXEC mode.
show dvlan-tunnel	Display all interfaces enabled for Double VLAN Tunneling
show dvlan-tunnel interface {*interface* | **all**}	Display detailed information about Double VLAN Tunneling for the specified interface or all interfaces.

Configuring MAC-Based VLANs

Beginning in Privileged EXEC mode, use the following commands to associate a MAC address with a configured VLAN. The VLAN does not need to be configured on the system to associate a MAC address with it. You can create up to 256 VLAN to MAC address associations.

Command	Purpose
configure	Enter global configuration mode.
vlan database	Enter VLAN database mode.
vlan association mac *mac-address vlan-id*	Associate a MAC address with a VLAN. • *mac-address* — MAC address to associate. (Range: Any MAC address in the format xxxx.xxxx.xxxx or xx:xx:xx:xx:xx:xx) • *vlanid* — VLAN to associate with subnet. (Range: 1-4093)
CTRL + Z	Exit to Privileged EXEC mode.
show vlan association mac [*mac-address*]	Display the VLAN associated with a specific configured MAC address. If no MAC address is specified, the VLAN associations of all the configured MAC addresses are displayed.

Configuring IP-Based VLANs

Beginning in Privileged EXEC mode, use the following commands to associate an IP subnet with a configured VLAN. The VLAN does not need to be configured on the system to associate an IP subnet with it. You can create up to 256 VLAN to MAC address associations.

Command	Purpose
configure	Enter global configuration mode.
vlan database	Enter VLAN database mode.
vlan association subnet *ip-address subnet-mask vlanid*	Associate an IP subnet with a VLAN. • *ip-address* — Source IP address. (Range: Any valid IP address) • *subnet-mask* — Subnet mask. (Range: Any valid subnet mask) • *vlanid* — VLAN to associated with subnet. (Range: 1-4093)
CTRL + Z	Exit to Privileged EXEC mode.
show vlan association subnet [*ip-address ip-mask]*	Display the VLAN associated with a specific configured IP-Address and netmask. If no IP Address and net mask are specified, the VLAN associations of all the configured IP-subnets are displayed.

Configuring a Protocol-Based VLAN

Beginning in Privileged EXEC mode, use the following commands to create and name a protocol group, and associate VLANs with the protocol group. When you create a protocol group, the switch automatically assigns it a unique group ID number. The group ID is used for both configuration and script generation to identify the group in subsequent commands.

A protocol group may have more than one interface associated with it, but each interface and protocol combination can be associated with one group only. If adding an interface to a group causes any conflicts with protocols currently associated with the group, adding the interface(s) to the group fails and no interfaces are added to the group. Ensure that the referenced VLAN is created prior to the creation of the protocol-based group except when GVRP is expected to create the VLAN.

Command	Purpose
configure	Enter global configuration mode.
vlan protocol group *name*	Create a new protocol group.
exit	Exit to Privileged EXEC mode.
show port protocol all	Obtain the group ID for the newly configured group.
configure	Enter global configuration mode.
vlan protocol group add protocol *groupid* **ethertype** *value*	Add any EtherType protocol to the protocol-based VLAN groups identified by *groupid*. A group may have more than one protocol associated with it. Each interface and protocol combination can be associated with one group only. If adding a protocol to a group causes any conflicts with interfaces currently associated with the group, this command fails and the protocol is not added to the group. • *groupid* — The protocol-based VLAN group ID. • *protocol* — The protocol you want to add. The ethertype can be any valid number in the range 0x0600-0xffff.
protocol vlan group all *groupid*	(Optional) Add all physical interfaces to the protocol-based group identified by *groupid*. You can add individual interfaces to the protocol-based group as shown in the next two commands. *groupid* — The protocol-based VLAN group ID.
interface *interface*	Enter interface configuration mode for the specified interface. *interface* — Specific interface type and number, such as te1/0/8.
protocol vlan group *groupid*	Add the physical unit/port interface to the protocol-based group identified by groupid. *groupid* — The protocol-based VLAN group ID.
exit	Exit to global configuration mode.
vlan database	Enter VLAN database mode.

Command	Purpose
protocol group *groupid vlanid*	Attach a VLAN ID to the protocol-based group identified by groupid. A group may only be associated with one VLAN at a time. However, the VLAN association can be changed. • *groupid* — The protocol-based VLAN group ID, which is automatically generated when you create a protocol-based VLAN group with the **vlan protocol group** command. To see the group ID associated with the name of a protocol group, use the **show port protocol all** command. • *vlanid* — A valid VLAN ID.
CTRL + Z	Exit to Privileged EXEC mode.
show port protocol [**all**| *groupid*]	Display the Protocol-Based VLAN information for either the entire system or for the indicated group.

Configuring GVRP

Beginning in Privileged EXEC mode, use the following commands to enable GVRP on the switch and on an interface, and to configure various GVRP settings.

Command	Purpose
configure	Enter global configuration mode.
gvrp enable	Enable GVRP on the switch.
interface *interface*	Enter interface configuration mode for the specified port or LAG. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3** or **port-channel 3**. You can also specify a range of interfaces with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12.
gvrp enable	Enable GVRP on the interface.

Command	Purpose
switchport forbidden vlan {add *vlan-list*|**remove** *vlan-list*}	(Optional) Forbids adding the specified VLANs to a port. To revert to allowing the addition of specific VLANs to the port, use the remove parameter of this command. **add** *vlan-list* — List of valid VLAN IDs to add to the forbidden list. Separate nonconsecutive VLAN IDs with a comma and no spaces. Use a hyphen to designate a range of IDs. **remove** *vlan-list* — List of valid VLAN IDs to remove from the forbidden list. Separate nonconsecutive VLAN IDs with a comma and no spaces. Use a hyphen to designate a range of IDs.
gvrp registration-forbid	(Optional) Deregister all VLANs on a port and prevent any dynamic registration on the port.
gvrp vlan-creation-forbid	(Optional) Disable dynamic VLAN creation.
exit	Exit to global configuration mode.
vlan database	Enter VLAN database mode.
vlan makestatic *vlan-id*	(Optional) Change a dynamically created VLAN (one that is created by GVRP registration) to a static VLAN (one that is permanently configured and defined). *vlan-id* — Valid vlan ID. Range is 2-4093.
CTRL + Z	Exit to Privileged EXEC mode.
show gvrp configuration	Display GVRP configuration information. Timer values are displayed. Other data shows whether GVRP is enabled and which ports are running GVRP.

Configuring Voice VLANs

Beginning in Privileged EXEC mode, use the following commands to enable the Voice VLAN feature on the switch and on an interface.

Command	Purpose
configure	Enter global configuration mode.
voice vlan	Enable the voice vlan capability on the switch.
interface *interface*	Enter interface configuration mode for the specified interface. *interface* — Specific interface, such as te1/0/8. You can also specify a range of interfaces with the **interface range** command, for example, **interface range te1/0/8-12** enters Interface Configuration mode for ports 8–12.
voice vlan {*vlanid* | **dot1p** *priority* | **none** | **untagged** | **data** *priority* {**trust** | **untrust**} | **auth** {**enable** | **disable**} | **dscp** *value*}	Enable the voice vlan capability on the interface. • *vlanid*—The voice VLAN ID. • *priority*—The Dot1p priority for the voice VLAN on the port. • **trust**—Trust the dot1p priority or DSCP values contained in packets arriving on the voice vlan port. • **untrust**—Do not trust the dot1p priority or DSCP values contained in packets arriving on the voice vlan port. • **auth {enable | disable}** — Use **enable** to allow voice traffic on unauthorized voice vlan port. Use **disable** to prevent voice traffic on an Unauthorized voice vlan port • **dscp** *value*—The DSCP value (Range: 0–64).
CTRL + Z	Exit to Privileged EXEC mode.
show voice vlan [**interface** {*interface* |**all**}]	Display voice VLAN configuration information for the switch, for the specified interface, or for all interfaces.

VLAN Configuration Examples

This section contains the following examples:

- Configuring VLANs Using Dell OpenManage Administrator
- Configuring VLANs Using the CLI
- Configuring a Voice VLAN

NOTE: For an example that shows how to use a RADIUS server to provide VLAN information, see "Controlling Authentication-Based VLAN Assignment" on page 486. For an example that shows how to allow the switch to dynamically create RADIUS-assigned VLANS, see "Allowing Dynamic VLAN Creation of RADIUS-Assigned VLANs" on page 490.

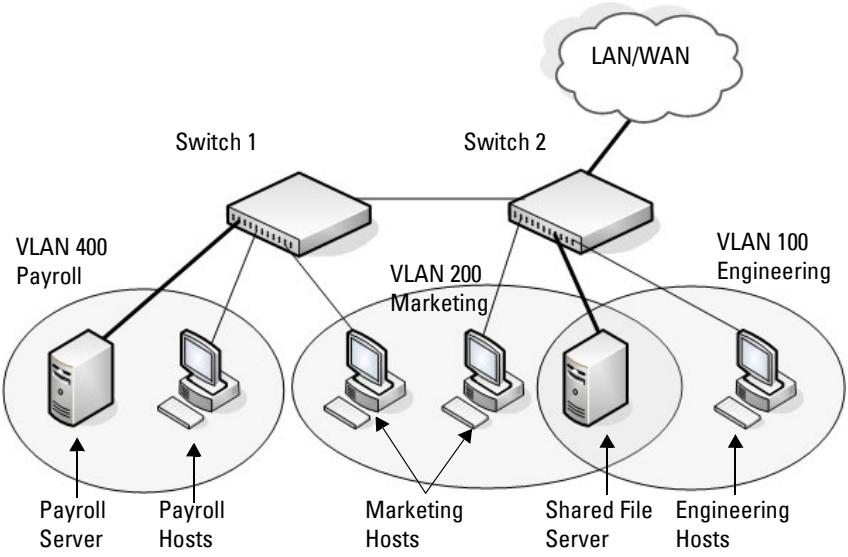
This example assumes that network administrator wants to create the VLANs in Table 22-9:

Table 22-9. Example VLANs

VLAN ID	VLAN Name	VLAN Type	Purpose
100	Engineering	Port-based	All employees in the Engineering department use this VLAN. Confining this department's traffic to a single VLAN helps reduce the amount of traffic in the broadcast domain, which increases bandwidth.
200	Marketing	Port-based	All employees in the Marketing department use this VLAN.
300	Sales	MAC-based	The sales staff works remotely but occasionally comes to the office. Since these employees do not have assigned work areas, they typically plug their laptops into a network port in an available cubicle, office, or conference room.
400	Payroll	Port-based	The payroll department has sensitive traffic and needs its own VLAN to help keep that traffic private.

Figure 22-26 shows the network topology for this example. As the figure shows, there are two switches, two file servers, and many hosts. One switch has an uplink port that connects it to a layer 3 device and the rest of the corporate network.

Figure 22-26. Network Topology for Port-Based VLAN Configuration



The network in Figure 22-26 has the following characteristics:

- Each connection to a host represents multiple ports and hosts.
- The Payroll and File servers are connected to the switches through a LAG.
- Some of the Marketing hosts connect to Switch 1, and some connect to Switch 2.
- The Engineering and Marketing departments share the same file server.
- Because security is a concern for the Payroll VLAN, the ports and LAG that are members of this VLAN will accept and transmit only traffic tagged with VLAN 400.
- The Sales staff might connect to a port on Switch 1 or Switch 2.

Table 22-10 shows the port assignments on the switches.

Table 22-10. Switch Port Connections

Port/LAG	Function
Switch 1	
1	Connects to Switch 2
2–15	Host ports for Payroll
16–20	Host ports for Marketing
LAG1 (ports 21–24)	Connects to Payroll server
Switch 2	
1	Connects to Switch 1
2–10	Host ports for Marketing
11–30	Host ports for Engineering
LAG1 (ports 35–39)	Connects to file server
LAG2 (ports 40–44)	Uplink to router.

Configuring VLANs Using Dell OpenManage Administrator

This example shows how to perform the configuration by using the web-based interface.

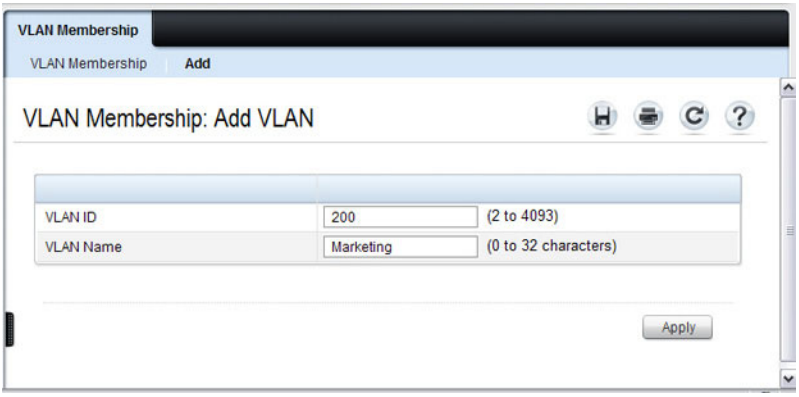
Configure the VLANs and Ports on Switch 1

Use the following steps to configure the VLANs and ports on Switch 1. None of the hosts that connect to Switch 1 use the Engineering VLAN (VLAN 100), so it is not necessary to create it on that switch.

To configure Switch 1:

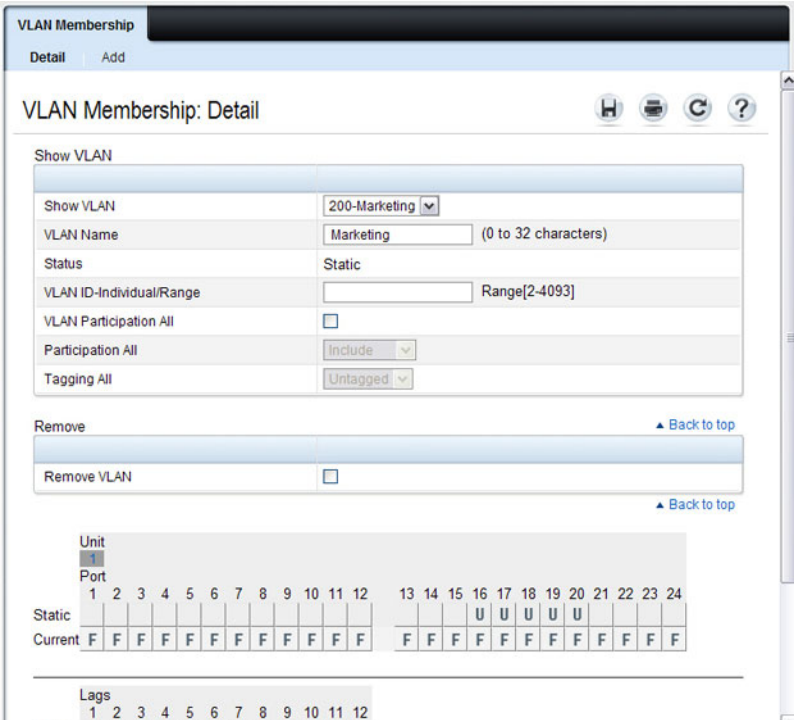
1. Create the Marketing, Sales, and Payroll VLANs.
   a. From the **Switching → VLAN → VLAN Membership** page, click **Add**.
   b. In the **VLAN ID** field, enter 200.
   c. In the **VLAN Name** field, enter Marketing.
   d. Click **Apply**.

Figure 22-27. Add VLANs



   e. Repeat steps b–d to create VLANs 300 (Sales) and 400 (Payroll).
2. Assign ports 16–20 to the Marketing VLAN.
   a. From the **Switching → VLAN → VLAN Membership** page, select 200-Marketing from the **Show VLAN** field.
   b. In the **Static** row, click the space for ports 16–20 so the U (untagged) displays for each port.

Figure 22-28. VLAN Membership - VLAN 200



3 Click **Apply**.

4 Assign ports 2–15 and LAG1 to the Payroll VLAN.

   a From the **Switching → VLAN → VLAN Membership** page, select 400-Payroll from the Show VLAN field.

   b In the Static row, click the space for ports 2–15 and LAG 1 so the U (untagged) displays for each port, and then click **Apply**.

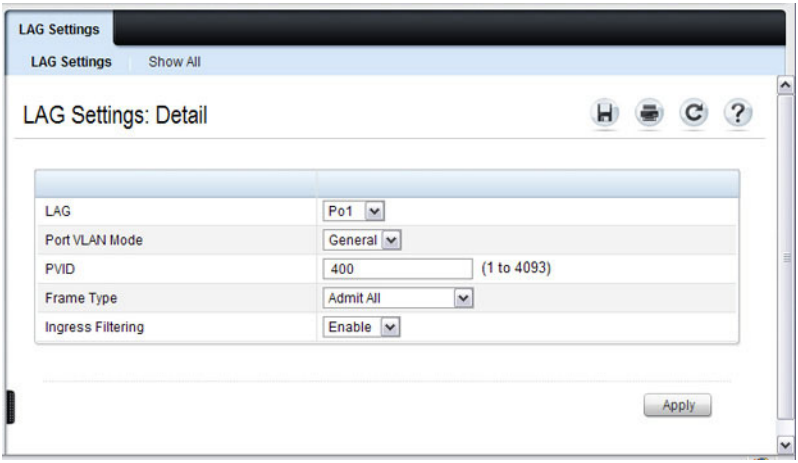
5. Configure LAG 1 to be in general mode and specify that the LAG will accept tagged or untagged frames, but that untagged frames will be transmitted tagged with PVID 400.

   a. From the **Switching → VLAN → LAG Settings** page, make sure Po1 is selected.

b. Configure the following settings:
   - Port VLAN Mode — General
   - PVID — 400
   - Frame Type — AdmitAll

c. Click **Apply**.

Figure 22-29. LAG Settings



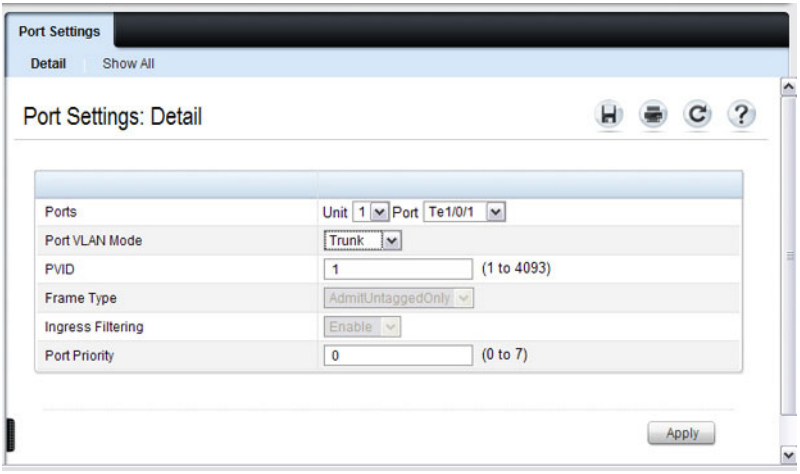
6 Configure port 1 as a trunk port.

a From the **Switching → VLAN → Port Settings** page, make sure port Te1/0/1 is selected.

b From the **Port VLAN Mode** field, select Trunk.

c Click **Apply**.

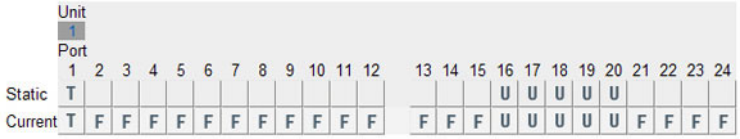
Figure 22-30. Trunk Port Configuration



7 From the **Switching → VLAN → VLAN Membership** page, verify that port 1 is marked as a tagged member (T) for each VLAN.

Figure 22-31 shows VLAN 200, in which port 1 is a tagged member, and ports 16–20 are untagged members.

Figure 22-31. Trunk Port Configuration



8 Configure the MAC-based VLAN information.

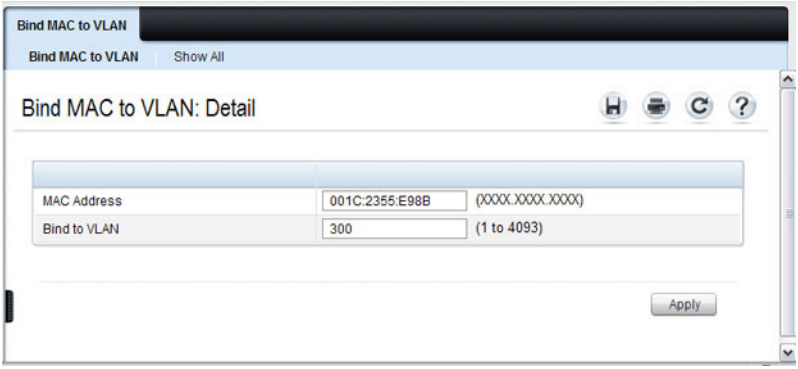
a Go to the **Switching → VLAN → Bind MAC to VLAN** page.

b In the **MAC Address** field, enter a valid MAC address, for example 00:1C:23:55:E9:8B.

c In the **Bind to VLAN** field, enter 300, which is the Sales VLAN ID.

d Click **Apply**.

Figure 22-32. Trunk Port Configuration



e Repeat steps b–d to add additional MAC address-to-VLAN information for the Sales department.

9 To save the configuration so that it persists across a system reset, use the following steps:

a Go to the **System → File Management→ Copy Files** page

b Select Copy Configuration and ensure that Running Config is the source and Startup Config is the destination.

c Click **Apply**.

Configure the VLANs and Ports on Switch 2

Use the following steps to configure the VLANs and ports on Switch 2. Many of the procedures in this section are the same as procedures used to configure Switch 1. For more information about specific procedures, see the details and figures in the previous section.

To configure Switch 2:

1. Create the Engineering, Marketing, Sales, and Payroll VLANs.

   Although the Payroll hosts do not connect to this switch, traffic from the Payroll department must use Switch 2 to reach the rest of the network and Internet through the uplink port. For that reason, Switch 2 must be aware of VLAN 400 so that traffic is not rejected by the trunk port.

2. Configure LAG 1 as a general port so that it can be a member of multiple VLANs.
   a. From the **Switching → VLAN → LAG Settings** page, make sure Po1 is selected.
   b. From the **Port VLAN Mode** field, select General.
   c. Click **Apply**.
3. Configure port 1 as a trunk port.
4. Configure LAG2 as a trunk port.
5. Assign ports 1–10 to VLAN 200 as untagged (U) members.
6. Assign ports 11–30 to VLAN 100 as untagged (U) members.
7. Assign LAG1 to VLAN 100 and 200 as a tagged (T) member.
8. Assign port 1 and LAG2 to VLAN 100, VLAN 200, VLAN 300, and VLAN 400 as a tagged (T) member.
9. Configure the MAC-based VLAN information.
10. If desired, copy the running configuration to the startup configuration.

Configuring VLANs Using the CLI

This example shows how to perform the same configuration by using CLI commands.

Configure the VLANs and Ports on Switch 1

Use the following steps to configure the VLANs and ports on Switch 1. None of the hosts that connect to Switch 1 use the Engineering VLAN (VLAN 100), so it is not necessary to create it on that switch.

To configure Switch 1:

1. Create VLANs 200 (Marketing), 300 (Sales), and 400 (Payroll), and associate the VLAN ID with the appropriate name.

```
console#configure
console(config-vlan)#vlan 200,300,400
console(config)#vlan 200
console(config-vlan200)#name Marketing
console(config-vlan200)#exit
console(config)#vlan 300
```

```
console(config-vlan300)#name Sales
console(config-vlan300)#exit
console(config)#vlan 400
console(config-vlan400)#name Payroll
console(config-vlan400)#exit
```

2. Assign ports 16–20 to the Marketing VLAN.

```
console(config)#interface range tengigabitEthernet
1/0/16-20
console(config-if)#switchport mode access
console(config-if)#switchport access vlan 200
console(config-if)#exit
```

3. Assign ports 2–15 to the Payroll VLAN

```
console(config)#interface range tengigabitEthernet
1/0/2-15
console(config-if)#switchport mode access
console(config-if)#switchport access vlan 400
console(config-if)#exit
```

4. Assign LAG1 to the Payroll VLAN and specify that frames will always be transmitted tagged with a PVID of 400.

```
console(config)#interface port-channel 1
console(config-if-ch1)#switchport mode general
console(config-if-ch1)#switchport general allowed
vlan add 400 tagged
console(config-if-ch1)#switchport general pvid 400
console(config-if-ch1)#exit
```

5. Configure port 1 as a trunk port and add VLAN 200, VLAN 300, and VLAN 400 as members.

```
console(config)#interface tengigabitEthernet 1/0/1
console(config-if-Te1/0/1)#switchport mode trunk
console(config-if-Te1/0/1)#switchport trunk
allowed vlan add 200,300,400
console(config-if-Te1/0/1)#exit
```

6. Configure the MAC-based VLAN information.

   The following commands show how to associate a system with a MAC address of 00:1C:23:55:E9:8B with VLAN 300. Repeat the **vlan association mac** command to associate additional MAC addresses with VLAN 300.

```
console(config)#vlan database
console(config-vlan)#vlan association mac
00:1C:23:55:E9:8B 300
console(config-vlan)#exit
console(config)#exit
```

7. To save the configuration so that it persists across a system reset, use the following command:

```
console#copy running-config startup-config
```

8. View the VLAN settings.

```
console#show vlan

VLAN     Name        Ports          Type        Authorization
-----    ---------   ------------   ---------   -------------
1        Default     Po1-12,        Default     Required
                     te1/0/2-15,
                     te1/0/21-24

200      Marketing   te1/0/1,       Static      Required
                     te1/0/16-20
300      Sales       te1/0/1        Static      Required
400      Payroll     te1/0/1-15     Static      Required
```

9. View the VLAN membership information for a port.

```
console#show interfaces switchport te1/0/1

Port: Te1/0/1
VLAN Membership mode:Trunk Mode

Operating parameters:
PVID: 1
Ingress Filtering: Enabled
Acceptable Frame Type: VLAN Only
Default Priority: 0
GVRP status:Disabled
```

```
Protected:Disabled

Port Te1/0/1 is member in:

VLAN    Name                Egress rule   Type
----    ----------------- -----------   --------
200     Marketing           Tagged        Static
300     Sales               Tagged        Static
400     Payroll             Tagged        Static
```

Configure the VLANs and Ports on Switch 2

Use the following steps to configure the VLANs and ports on Switch 2. Many of the procedures in this section are the same as procedures used to configure Switch 1. For more information about specific procedures, see the details and figures in the previous section.

To configure Switch 2:

1. Create the Engineering, Marketing, Sales, and Payroll VLANs.

   Although the Payroll hosts do not connect to this switch, traffic from the Payroll department must use Switch 2 to reach the rest of the network and Internet through the uplink port. For that reason, Switch 2 must be aware of VLAN 400 so that traffic is not rejected by the trunk port.

2. Configure ports 2-10 as access ports and add VLAN 200 to the ports.
3. Configure ports 11–30 as access ports and add VLAN 100 to the ports.
4. Configure LAG 1 as a general port so that it can be a member of multiple untagged VLANs and add VLAN 100 and VLAN 200 to the LAG.
5. Configure port 1 and LAG 2 trunk ports and add VLAN 100, VLAN 200, VLAN 300, and VLAN 400 to the port and LAG.
6. Configure the MAC-based VLAN information.
7. If desired, copy the running configuration to the startup configuration.
8. View VLAN information for the switch and ports.

Configuring a Voice VLAN

The commands in this example create a VLAN for voice traffic with a VLAN ID of 25. Port 10 is set to an 802.1Q VLAN. In in this example, there are multiple devices connected to port 10, so the port must be in general mode in order to enable MAC-based 802.1X authentication. Next, Voice VLAN is enabled on the port with the Voice VLAN ID set to 25. Finally, Voice VLAN authentication is disabled on port 10 because the phone connected to that port does not support 802.1X authentication. All other devices are required to use 802.1X authentication for network access. For more information about 802.1X authentication, see "Configuring Port and System Security" on page 457.



NOTE: In an environment where the IP phone uses LLDP-MED to obtain configuration information, an additional step to enable LLDP-MED on the interface would be required by issuing the **lldp med** command in Interface Configuration mode.

To configure the switch:

1. Create the voice VLAN

```
console#configure
console(config)#vlan database
console(config-vlan)#vlan 25
console(config-vlan)#exit
```

2. Enable the Voice VLAN feature on the switch.

```
console(config)#voice vlan
```

3. Configure port 10 to be in general mode.

```
console(config)#interface te1/0/10
console(config-if-Te1/0/10)#switchport mode
general
```

4. Enable port-based 802.1X authentication on the port. This step is required only if there are multiple devices that use port-based authentication connected to the port.

```
console(config-if-Te1/0/10)#dot1x port-control
mac-based
```

5 Enable the voice VLAN feature on the interface

```
console(config-if-Te1/0/10)#voice vlan 25
```

6 Disable authentication for the voice VLAN on the port. This step is required only if the voice phone does not support port-based authentication.

```
console(config-if-Te1/0/10)#voice vlan auth
disable
```

7 Exit to Privileged Exec mode.

```
console(config-if-Te1/0/10)#<CTRL+Z>
```

8 View the voice VLAN settings for port 10.

```
console#show voice vlan interface te1/0/10

Interface.............................. Te1/0/10
Voice VLAN Interface Mode.............. Enabled
Voice VLAN ID.......................... 25
Voice VLAN COS Override................ False
Voice VLAN DSCP Value.................. 46
Voice VLAN Port Status................. Disabled
Voice VLAN Authentication.............. Disabled
```

Configuring a Private VLAN

1 Configure the VLANs and their roles:

This example configures VLAN 100 as the primary VLAN, secondary VLAN 101 as the community VLAN and secondary VLANs 102 and 103 as the isolated VLANs:

```
switch# configure
switch(config)# vlan 100
switch(config-vlan-100)# private-vlan primary
switch(config-vlan-100)# exit
switch(config)# vlan 101
switch(config-vlan-101)# private-vlan community
switch(config-vlan-101)# exit
switch(config)# vlan 102
switch(config-vlan-102)# private-vlan isolated
switch(config-vlan-102)# exit
switch(config)# vlan 103
switch(config-vlan-103)# private-vlan isolated
```

```
switch(config-vlan-103)# exit
```

2 Associate the community and isolated VLANs with the primary VLAN.

```
switch(config)# vlan 100
switch(config-vlan-100)# private-vlan association 101-102
switch(config-vlan-100)# exit
```

This completes the configuration of the private VLAN. The only remaining step is to assign the ports to the private VLAN.

3 Assign the router connected port to the primary VLAN:

```
console(config)#interface te1/1/1
console(config-if-Te1/1/1)#switchport mode private-vlan
promiscuous
console(config-if-Te1/1/1)#switchport private-vlan mapping 100
101-102
console(config-if-Te1/1/1)#exit
```

4 Assign the community VLAN ports:

```
console(config)#interface gi1/0/11
console(config-if-Gi1/0/11)#switchport mode private-vlan host
console(config-if-Gi1/0/11)#switchport private-vlan host-
association 100 101
console(config-if-Gi1/0/11)#interface gi1/0/12
console(config-if-Gi1/0/12)#switchport mode private-vlan host
console(config-if-Gi1/0/12)#switchport private-vlan host-
association 100 101
```

5 Assign the isolated VLAN ports:

```
console(config)#interface gi1/0/10
console(config-if-Gi1/0/10)#switchport mode private-vlan host
console(config-if-Gi1/0/10)#switchport private-vlan host-
association 100 102
console(config-if-Gi1/0/10)#interface gi2/0/10
console(config-if-Gi2/0/10)#switchport mode private-vlan host
console(config-if-Gi2/0/10)#switchport private-vlan host-
association 100 102
console(config-if-Gi2/0/10)#interface gi2/0/11
console(config-if-Gi2/0/11)#switchport mode private-vlan host
console(config-if-Gi2/0/11)#switchport private-vlan host-
association 100 102
```

6 Show the configuration:

```
console(config)#show vlan private-vlan type

VLAN Type
```

```
---- -----------------------
100  primary
101  community
102  isolated
103  isolated

console#show vlan private-vlan

Primary VLAN Secondary VLAN Community
------------ -------------- -------------------
100          102            101

console(config)#show vlan

VLAN   Name         Ports          Type
-----  ----------- -------------  -------------
1      default     Po1-128,       Default
                   Te1/1/1,
                   Gi1/0/1-10,
                   Gi1/0/13-24
100    VLAN0100    Te1/1/1,       Static
                   Gi1/0/11-12
101    VLAN0101    Gi1/0/11       Static
102    VLAN0102    Gi1/0/12       Static
```

23

Configuring the Spanning Tree Protocol

This chapter describes how to configure the Spanning Tree Protocol (STP) settings on the switch.

The topics covered in this chapter include:

- STP Overview
- Default STP Values
- Configuring Spanning Tree (Web)
- Configuring Spanning Tree (CLI)
- STP Configuration Examples

STP Overview

STP is a layer 2 protocol that provides a tree topology for switches on a bridged LAN. STP allows a network to have redundant paths without the risk of network loops. STP uses the spanning-tree algorithm to provide a single path between end stations on a network.

PowerConnect 8000-series and 8100-series switches support Classic STP, Multiple STP, and Rapid STP.

What Are Classic STP, Multiple STP, and Rapid STP?

Classic STP provides a single path between end stations, avoiding and eliminating loops.

Multiple Spanning Tree Protocol (MSTP) supports multiple instances of Spanning Tree to efficiently channel VLAN traffic over different interfaces. Each instance of the Spanning Tree behaves in the manner specified in IEEE 802.1w, Rapid Spanning Tree (RSTP), with slight modifications in the working but not the end effect (chief among the effects, is the rapid transitioning of the port to Forwarding). The difference between the RSTP and the traditional STP (IEEE 802.1d) is the ability to configure and

recognize full-duplex connectivity and ports which are connected to end stations, resulting in rapid transitioning of the port to the Forwarding state and the suppression of Topology Change Notifications.

MSTP is compatible to both RSTP and STP. It behaves appropriately to STP and RSTP bridges. A MSTP bridge can be configured to behave entirely as a RSTP bridge or a STP bridge.

How Does STP Work?

The switches (bridges) that participate in the spanning tree elect a switch to be the root bridge for the spanning tree. The root bridge is the switch with the lowest bridge ID, which is computed from the unique identifier of the bridge and its configurable priority number. When two switches have an equal bridge ID value, the switch with the lowest MAC address is the root bridge.

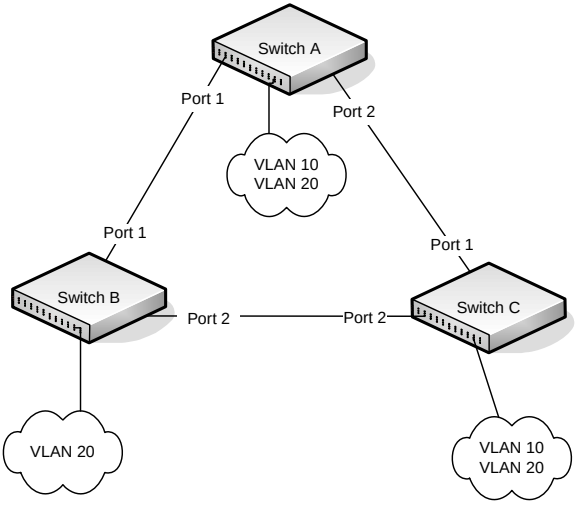
After the root bridge is elected, each switch finds the lowest-cost path to the root bridge. The port that connects the switch to the lowest-cost path is the root port on the switch. The switches in the spanning tree also determine which ports have the lowest-path cost for each segment. These ports are the designated ports. Only the root ports and designated ports are placed in a forwarding state to send and receive traffic. All other ports are put into a blocked state to prevent redundant paths that might cause loops.

To determine the root path costs and maintain topology information, switches that participate in the spanning tree use Bridge Protocol Data Units (BPDUs) to exchange information.

How Does MSTP Operate in the Network?

In the following diagram of a small 802.1d bridged network, STP is necessary to create an environment with full connectivity and without loops.

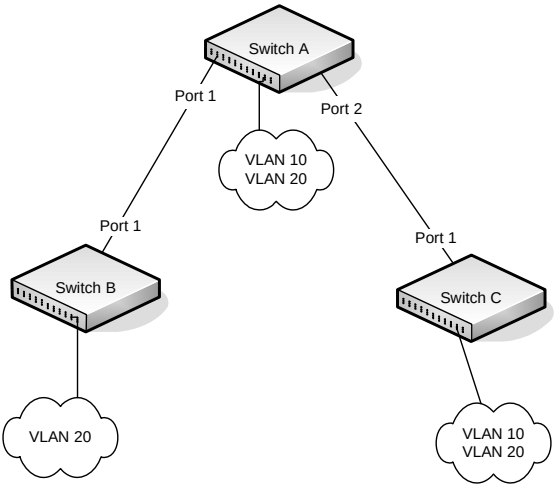
Figure 23-1. Small Bridged Network



Assume that Switch A is elected to be the Root Bridge, and Port 1 on Switch B and Switch C are calculated to be the root ports for those bridges, Port 2 on Switch B and Switch C would be placed into the Blocking state. This creates a loop-free topology. End stations in VLAN 10 can talk to other devices in VLAN 10, and end stations in VLAN 20 have a single path to communicate with other VLAN 20 devices.

Figure 23-2 shows the logical single STP network topology.

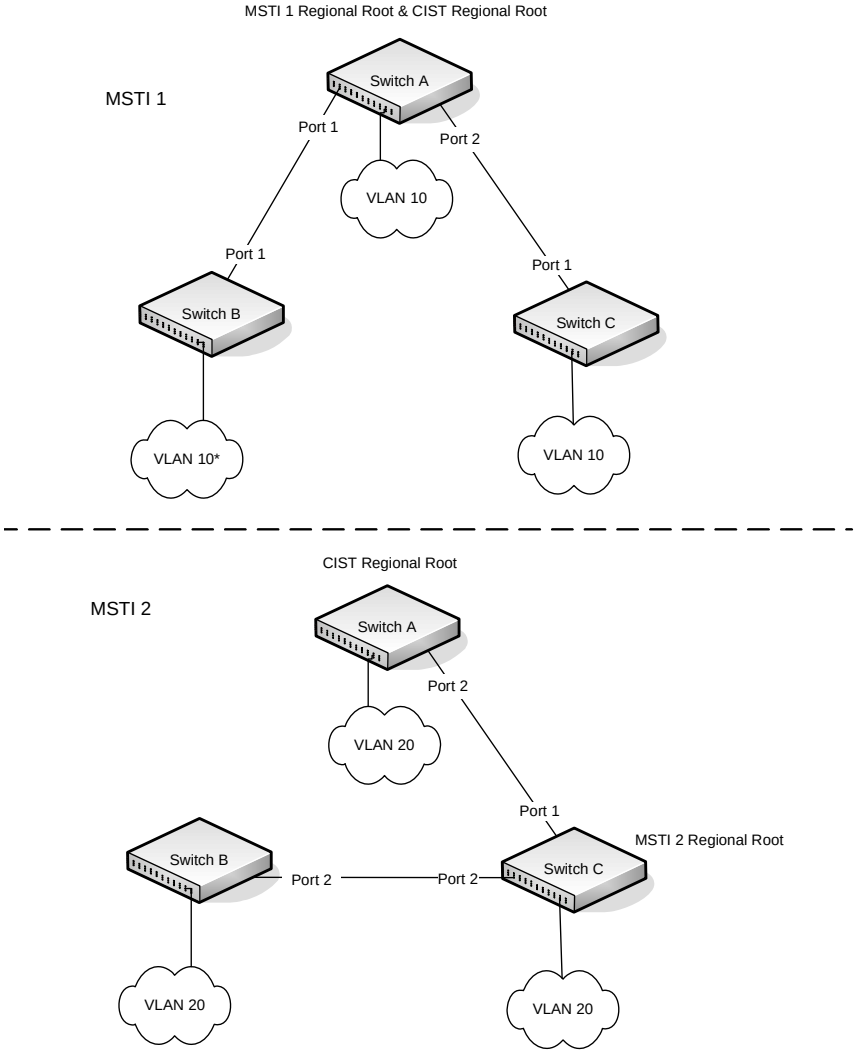
Figure 23-2. Single STP Topology



For VLAN 10 this single STP topology is fine and presents no limitations or inefficiencies. On the other hand, VLAN 20's traffic pattern is inefficient. All frames from Switch B will have to traverse a path through Switch A before arriving at Switch C. If the Port 2 on Switch B and Switch C could be used, these inefficiencies could be eliminated. MSTP does just that, by allowing the configuration of MSTIs based upon a VLAN or groups of VLANs. In this simple case, VLAN 10 could be associated with Multiple Spanning Tree Instance (MSTI)1 with an active topology similar to Figure 23-2 and VLAN 20 could be associated with MSTI 2 where Port 1 on both Switch A and Switch B begin discarding and all others forwarding. This simple modification creates an active topology with a better distribution of network traffic and an increase in available bandwidth.

The logical representation of the MSTP environment for these three switches is shown in Figure 23-3.

Figure 23-3. Logical MSTP Environment



In order for MSTP to correctly establish the different MSTIs as above, some additional changes are required. For example, the configuration would have to be the same on each and every bridge. That means that Switch B would have to add VLAN 10 to its list of supported VLANs (shown in Figure 23-3 with a *). This is necessary with MSTP to allow the formation of Regions made up of all switches that exchange the same MST Configuration Identifier. It is within only these MST Regions that multiple instances can exist. It will also allow the election of Regional Root Bridges for each instance. One common and internal spanning tree (CIST) Regional Root for the CIST and an MSTI Regional Root Bridge per instance will enable the possibility of alternate paths through each Region. Above Switch A is elected as both the MSTI 1 Regional Root and the CIST Regional Root Bridge, and after adjusting the Bridge Priority on Switch C in MSTI 2, it would be elected as the MSTI 2 Regional Root.

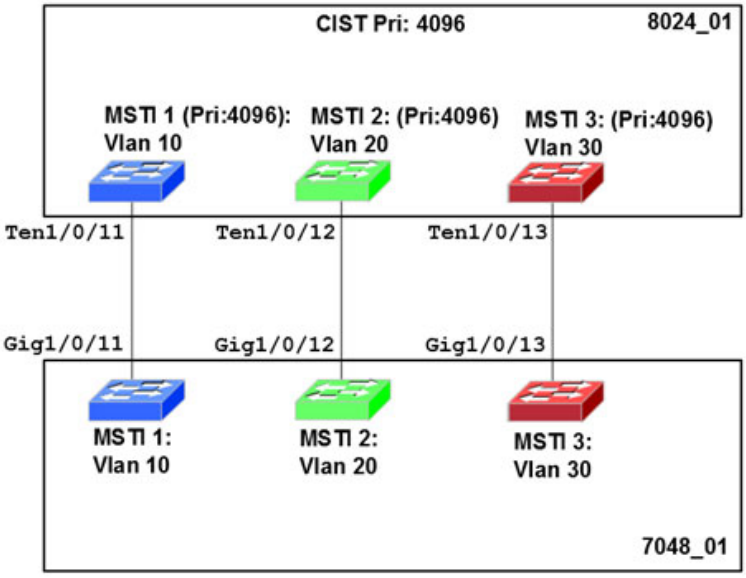
To further illustrate the full connectivity in an MSTP active topology, the following rules apply:

1. Each Bridge or LAN is in only one Region.
2. Every frame is associated with only one VID.
3. Frames are allocated either to the IST or MSTI within any given Region.
4. The internal spanning tree (IST) and each MSTI provides full and simple connectivity between all LANs and Bridges in a Region.
5. All Bridges within a Region reach a consistent agreement as to which ports interconnect that Region to a different Region and label those as Boundary Ports.
6. At the Boundary Ports, frames allocated to the CIST or MSTIs are forwarded or not forwarded alike.
7. The CIST provides full and simple connectivity between all LANs and Bridges in the network.

MSTP with Multiple Forwarding Paths

Consider the physical topology shown in Figure 23-4. It might be assumed that MSTI 2 and MSTI 3 would follow the most direct path for VLANs 20 and 30. However, using the default path costs, this is not the case. MSTI operates without considering the VLAN membership of the ports. This results in unexpected behavior if the active topology of an MSTI depends on a port that is not a member of the VLAN assigned to the MSTI and the port is selected as root port. In this configuration, port TE 1/0/11 is selected as the root port and ports TE1/0/12 and TE1/0/13 are blocked. To resolve the issue, set the port path cost of the directly connected links to allow the MSTIs to connect directly.

Figure 23-4. MSTP with Multiple Forwarding Paths



What are the Optional STP Features?

The PowerConnect 8000-series and 8100-series switches support the following optional STP features:

- BPDU flooding
- PortFast
- BPDU filtering
- Root guard
- Loop guard
- BPDU protection

BPDU Flooding

The BPDU flooding feature determines the behavior of the switch when it receives a BPDU on a port that is disabled for spanning tree. If BPDU flooding is configured, the switch will flood the received BPDU to all the ports on the switch which are similarly disabled for spanning tree.

Port Fast

The PortFast feature reduces the STP convergence time by allowing edge ports that are connected to end devices (such as a desktop computer, printer, or file server) to transition to the forwarding state without going through the listening and learning states.

BPDU Filtering

Ports that have the PortFast feature enabled continue to transmit BPDUs. The BPDU filtering feature prevents PortFast-enabled ports from sending BPDUs.

If BPDU filtering is configured globally on the switch, the feature is automatically enabled on all operational PortFast-enabled ports. These ports are typically connected to hosts that drop BPDUs. However, if an operational edge port receives a BPDU, the BPDU filtering feature disables PortFast and allows the port to participate in the spanning-tree calculation.

Enabling BPDU filtering on a specific port prevents the port from sending BPDUs and allows the port to drop any BPDUs it receives.

Root Guard

Enabling root guard on a port ensures that the port does not become a root port or a blocked port. When a switch is elected as the root bridge, all ports are designated ports unless two or more ports of the root bridge are connected together. If the switch receives superior STP BPDUs on a root-guard enabled port, the root guard feature moves this port to a root-inconsistent STP state, which is effectively equal to a listening state. No traffic is forwarded across this port. In this way, the root guard feature enforces the position of the root bridge.

When the STP mode is MSTP, the port may be a designated port in one MSTI and an alternate port in the CIST, etc. Root guard is a per port (not a per port per instance command) configuration, so all the MSTP instances this port participates in should not be in a root role.

Loop Guard

Loop guard protects a network from forwarding loops induced by BPDU packet loss. The reasons for failing to receive packets are numerous, including heavy traffic, software problems, incorrect configuration, and unidirectional link failure. When a non-designated port no longer receives BPDUs, the spanning-tree algorithm considers that this link is loop free and begins transitioning the link from blocking to forwarding. Once in forwarding state, the link may create a loop in the network.

Enabling loop guard prevents such accidental loops. When a port is no longer receiving BPDUs and the max age timer expires, the port is moved to a *loop-inconsistent blocking state*. In the loop-inconsistent blocking state, traffic is not forwarded so the port behaves as if it is in the blocking state. The port will remain in this state until it receives a BPDU. It will then transition through the normal spanning tree states based on the information in the received BPDU.



NOTE: Loop Guard should be configured only on non-designated ports. These include ports in alternate or backup roles. Root ports and designated ports should not have loop guard enabled so that they can forward traffic.

BPDU Protection

When the switch is used as an access layer device, most ports function as edge ports that connect to a device such as a desktop computer or file server. The port has a single, direct connection and is configured as an edge port to implement the fast transition to a forwarding state. When the port receives a BPDU packet, the system sets it to non-edge port and recalculates the spanning tree, which causes network topology flapping. In normal cases, these ports do not receive any BPDU packets. However, someone may forge BPDU to maliciously attack the switch and cause network flapping.

BPDU protection can be enabled in RSTP to prevent such attacks. When BPDU protection is enabled, the switch disables an edge port that has received BPDU and notifies the network manager about it.

Default STP Values

Spanning tree is globally enabled on the switch and on all ports and LAGs. Table 23-1 summarizes the default values for STP.

Table 23-1. STP Defaults

Parameter	Default Value
Enable state	Enabled (globally and on all ports)
Spanning-tree mode	RSTP (Classic STP and MSTP are disabled)
Switch priority	32768
BPDU flooding	Disabled
PortFast mode	Disabled
PortFast BPDU filter	Disabled
Loop guard	Disabled
BPDU protection	Disabled
Spanning tree port priority	128
Maximum-aging time	20 seconds
Forward-delay time	15 seconds
Maximum hops	20
Spanning-tree transmit hold count	6
MSTP region name	MAC address of switch
MSTP included VLANs	1

Configuring Spanning Tree (Web)

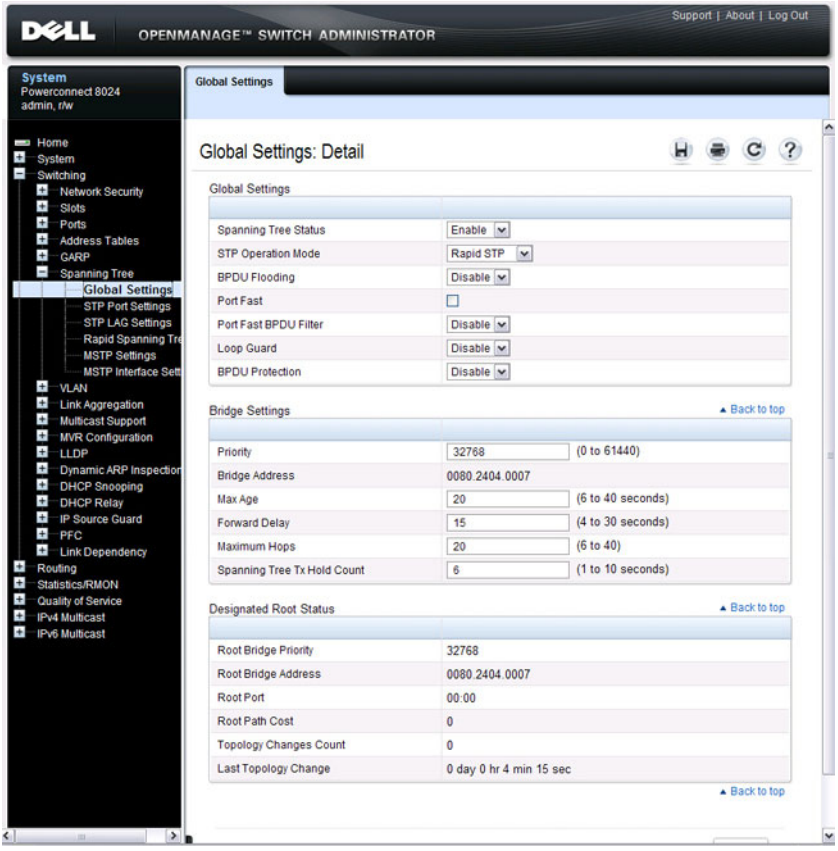
This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring STP settings on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click  at the top of the page.

STP Global Settings

The **STP Global Settings** page contains fields for enabling STP on the switch.

To display the **STP Global Settings** page, click **Switching** → **Spanning Tree** → **Global Settings** in the navigation panel.

Figure 23-5. Spanning Tree Global Settings

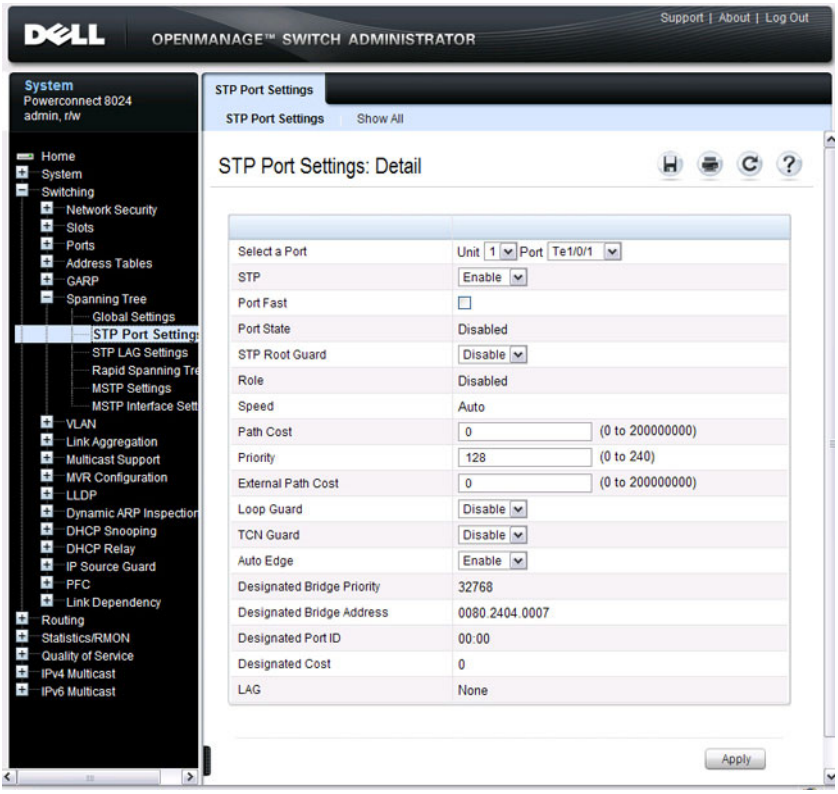


STP Port Settings

Use the **STP Port Settings** page to assign STP properties to individual ports.

To display the **STP Port Settings** page, click **Switching** → **Spanning Tree** → **STP Port Settings** in the navigation panel.

Figure 23-6. STP Port Settings

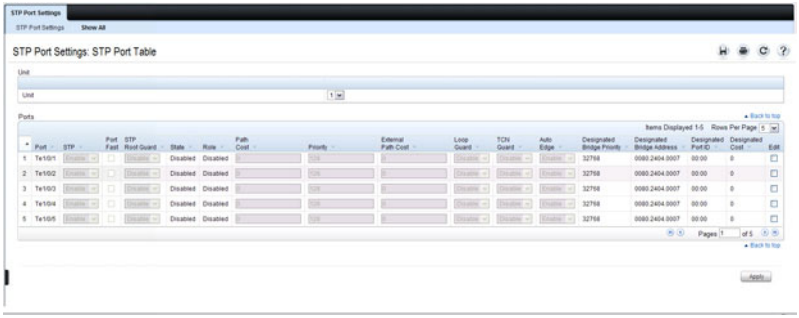


Configuring STP Settings for Multiple Ports

To configure STP settings for multiple ports:

1. Open the **STP Port Settings** page.
2. Click **Show All** to display the **STP Port Table**.

Figure 23-7. Configure STP Port Settings



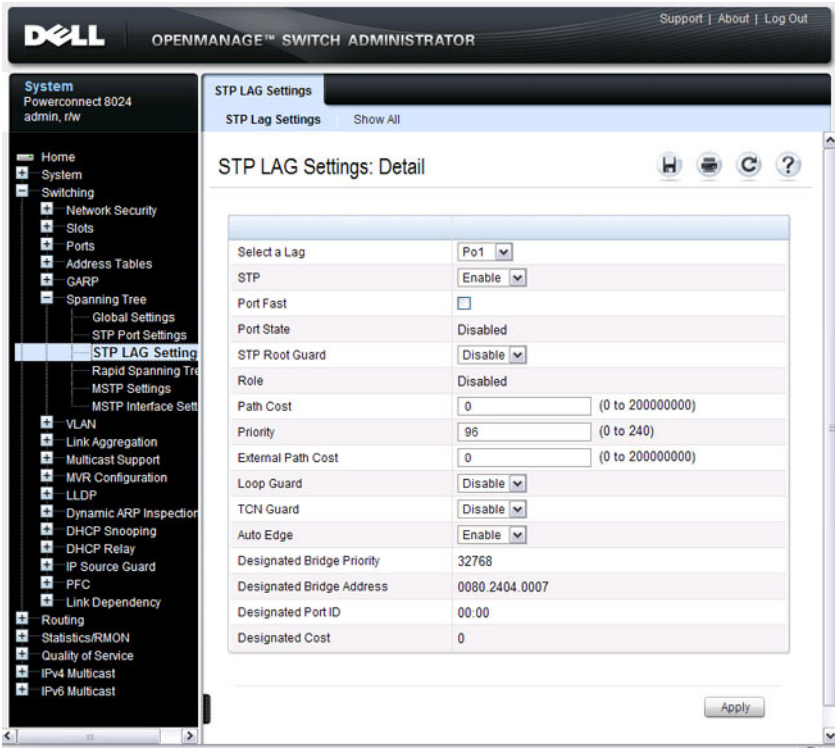
3. For each port to configure, select the check box in the **Edit** column in the row associated with the port.
4. Select the desired settings.
5. Click **Apply**.

STP LAG Settings

Use the **STP LAG Settings** page to assign STP aggregating ports parameters.

To display the **STP LAG Settings** page, click **Switching** → **Spanning Tree** → **STP LAG Settings** in the navigation panel.

Figure 23-8. STP LAG Settings

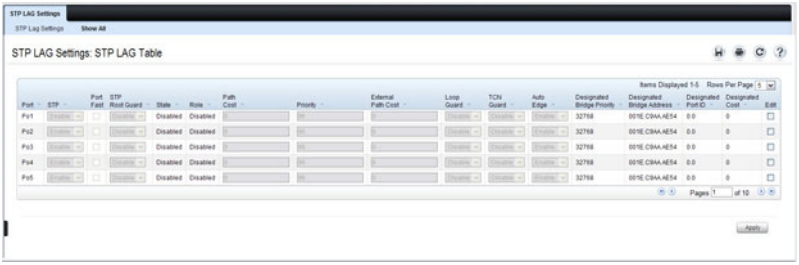


Configuring STP Settings for Multiple LAGs

To configure STP settings on multiple LAGS:

1. Open the **STP LAG Settings** page.
2. Click **Show All** to display the **STP LAG Table**.

Figure 23-9. Configure STP LAG Settings



3 For each LAG to configure, select the check box in the **Edit** column in the row associated with the LAG.

4 Select the desired settings.

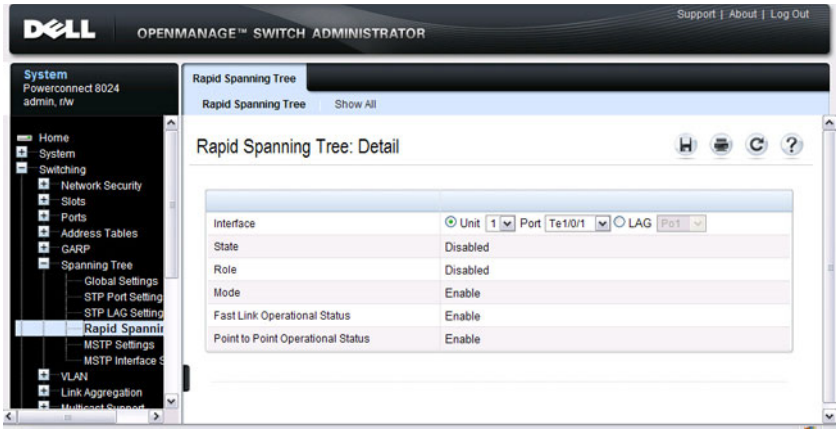
5 Click **Apply**.

Rapid Spanning Tree

Rapid Spanning Tree Protocol (RSTP) detects and uses network topologies that allow a faster convergence of the spanning tree without creating forwarding loops.

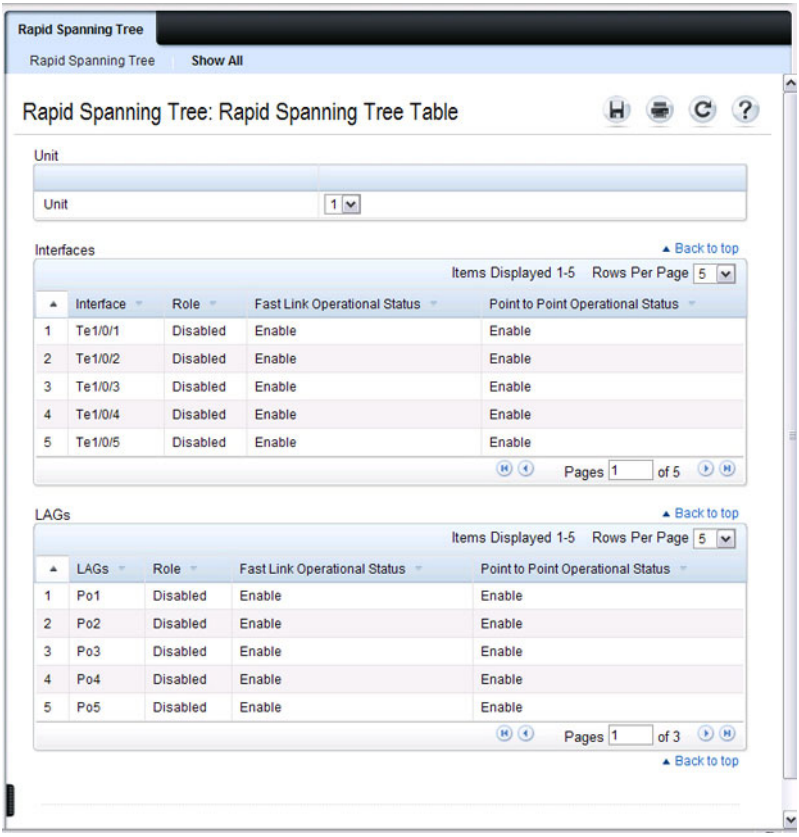
To display the **Rapid Spanning Tree** page, click **Switching** → **Spanning Tree** → **Rapid Spanning Tree** in the navigation panel.

Figure 23-10. Rapid Spanning Tree



To view RSTP Settings for all interfaces, click the **Show All** link. The **Rapid Spanning Tree Table** displays.

Figure 23-11. RSTP LAG Settings

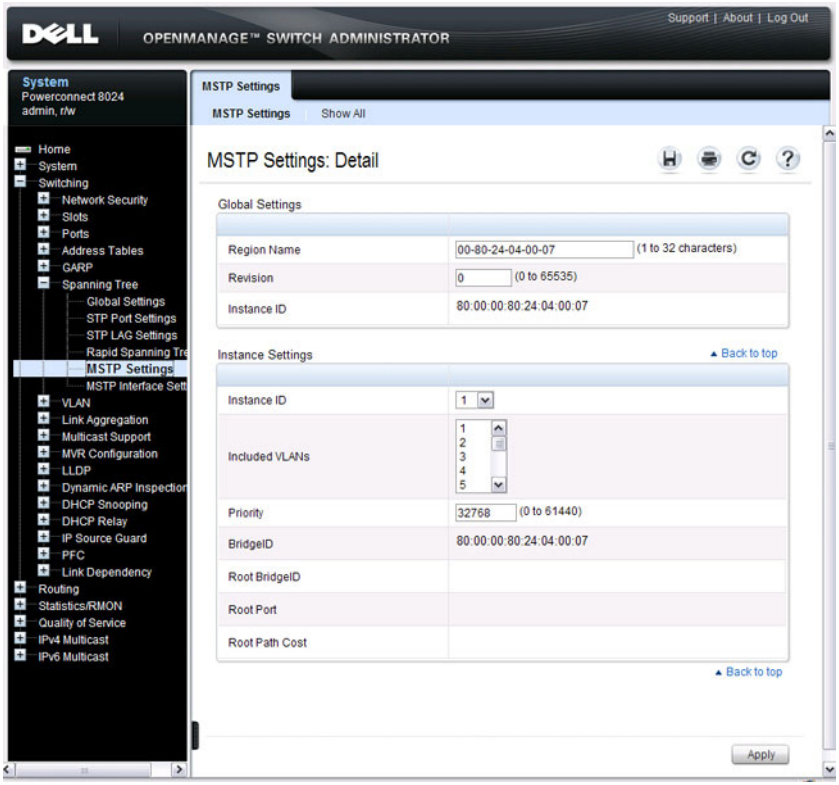


MSTP Settings

The Multiple Spanning Tree Protocol (MSTP) supports multiple instances of Spanning Tree to efficiently channel VLAN traffic over different interfaces. MSTP is compatible with both RSTP and STP; a MSTP bridge can be configured to behave entirely as a RSTP bridge or a STP bridge.

To display the **MSTP Settings** page, click **Switching** → **Spanning Tree** → **MSTP Settings** in the navigation panel.

Figure 23-12. MSTP Settings

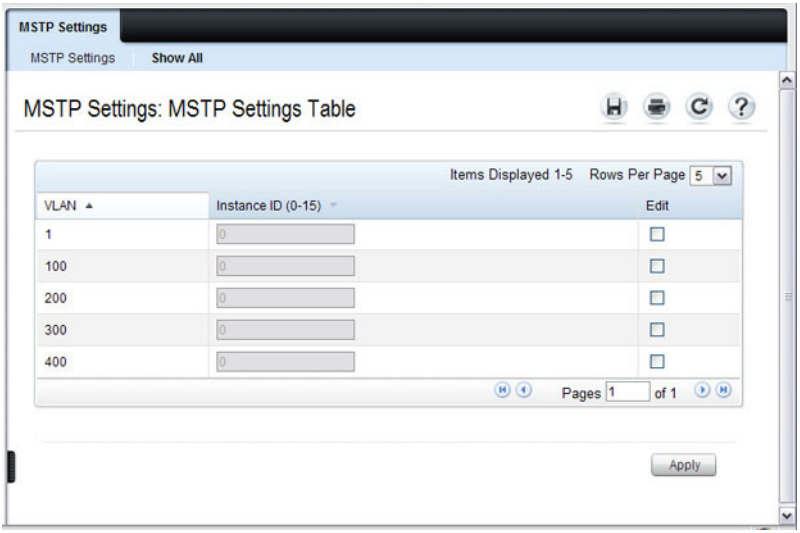


Viewing and Modifying the Instance ID for Multiple VLANs

To configure MSTP settings for multiple VLANS:

1. Open the **MSTP Settings** page.
2. Click **Show All** to display the **MSTP Settings Table**.

Figure 23-13. Configure MSTP Settings



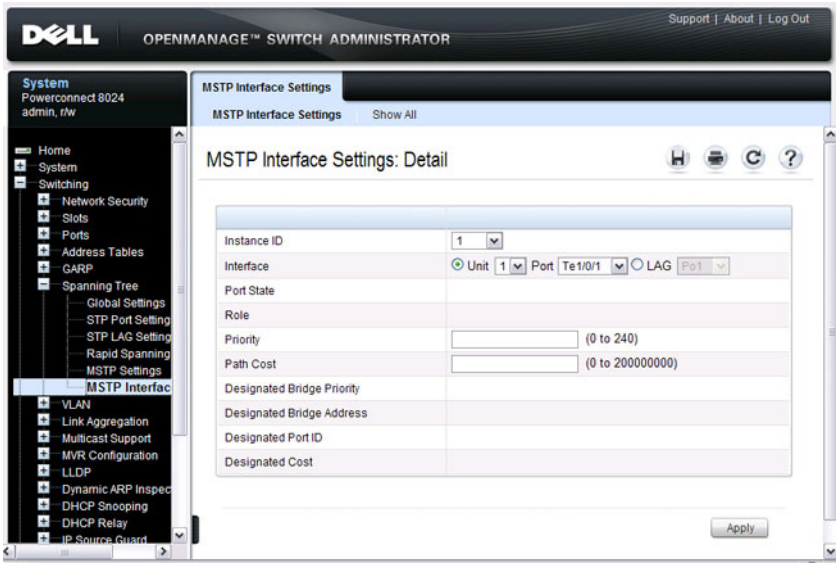
3. For each Instance ID to modify, select the check box in the **Edit** column in the row associated with the VLAN.
4. Update the **Instance ID** settings for the selected VLANs.
5. Click **Apply**.

MSTP Interface Settings

Use the **MSTP Interface Settings** page to assign MSTP settings to specific interfaces.

To display the **MSTP Interface Settings** page, click **Switching** → **Spanning Tree** → **MSTP Interface Settings** in the navigation panel.

Figure 23-14. MSTP Interface Settings

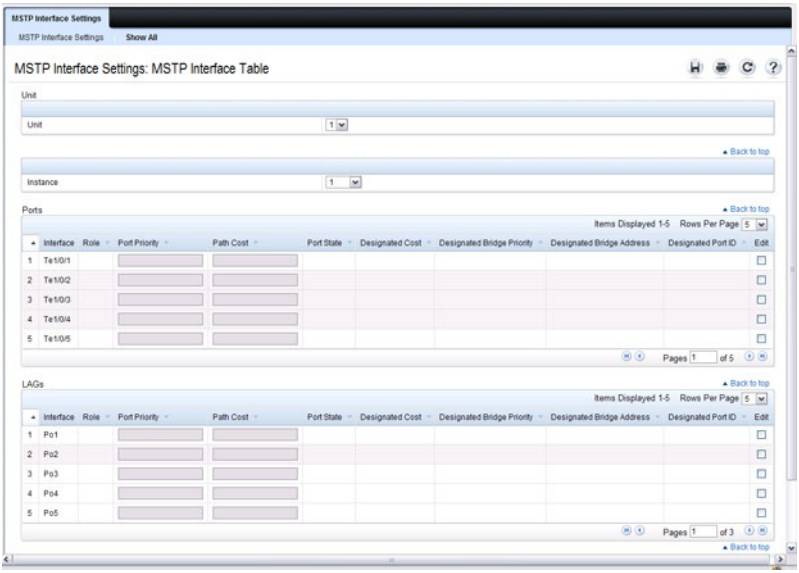


Configuring MSTP Settings for Multiple Interfaces

To configure MSTP settings for multiple interfaces:

1. Open the **MSTP Interface Settings** page.
2. Click **Show All** to display the **MSTP Interface Table**.

Figure 23-15. Configure MSTP Interface Settings



3 For each interface to configure, select the check box in the **Edit** column in the row associated with the interface.

4 Update the desired settings.

5 Click **Apply**.

Configuring Spanning Tree (CLI)

This section provides information about the commands you use to configure STP settings on the switch. For more information about the commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Configuring Global STP Bridge Settings

Beginning in Privileged EXEC mode, use the following commands to configure the global STP settings for the switch, such as the priority and timers.

Command	Purpose
configure	Enter global configuration mode.
spanning-tree	Enable spanning tree on the switch.
spanning tree mode {stp | rstp |mst}	Specify which spanning tree mode to use on the switch.
spanning-tree priority *priority*	Specify the priority of the bridge. (Range: 0–61440). The switch with the lowest priority value is elected as the root switch.
spanning-tree max-age *seconds*	Specify the switch maximum age time, which indicates the amount of time in seconds a bridge waits before implementing a topological change. Valid values are from (6 to 40) seconds.
spanning-tree forward-time *seconds*	Specify the switch forward delay time, which indicates the amount of time in seconds a bridge remains in a listening and learning state before forwarding packets. Valid values are from (4 to 30) seconds.
spanning-tree max-hops *hops*	Configure the maximum number of hops for the Spanning tree. Valid values are from (6 to 40).
spanning-tree transmit hold-count [*value*]	Set the maximum number of BPDUs that a bridge is allowed to send within a hello time window (2 seconds). The range for *value* is 1–10.
CTRL + Z	Exit to Privileged EXEC mode.

Command	Purpose
show spanning-tree [detail] [active | blockedports]	View information about spanning tree and the spanning tree configuration on the switch.

Configuring Optional STP Features

Beginning in Privileged EXEC mode, use the following commands to configure the optional STP features on the switch or on specific interfaces.

Command	Purpose
configure	Enter global configuration mode.
spanning-tree bpdu flooding	Allow the flooding of BPDUs received on non-spanning-tree ports to all other non-spanning-tree ports.
spanning-tree portfast	Enable PortFast on all switch ports.
spanning-tree portfast bpdufilter default	Prevent ports configured in PortFast mode from sending BPDUs.
spanning-tree loopguard default	Enable loop guard on all ports.
spanning-tree bpdu-protection	Enable BPDU protection on the switch.
interface *interface*	Enter interface configuration mode for the specified interface. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3** or **port-channel 4**. You can also specify a range of interfaces with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12. The **range** keyword is also valid for LAGs (port channels).
spanning-tree auto-portfast	Set the port to auto portfast mode. This enables the port to become a portfast port if it does not see any BPDUs for 3 seconds.
spanning-tree guard {root | loop | none}	Enable loop guard or root guard (or disable both) on the interface.

Command	Purpose
spanning-tree tcnguard	Prevent the port from propagating topology change notifications.
CTRL + Z	Exit to Privileged EXEC mode.
show spanning-tree summary	View various spanning tree settings and parameters for the switch.

Configuring STP Interface Settings

Beginning in Privileged EXEC mode, use the following commands to configure the STP settings for a specific interface.

Command	Purpose
configure	Enter global configuration mode.
interface *interface*	Enter interface configuration mode for the specified interface. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3** or **port-channel 4**.
	You can also specify a range of interfaces with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12. The **range** keyword is also valid for LAGs (port channels).
spanning-tree disable	Disable spanning-tree on the port.
spanning-tree port-priority *priority*	Specify the priority of the port. (Range: 0–240).
	The priority value is used to determine which ports are put in the forwarding state and which ports are put in the blocking state. A port with a lower priority value is more likely to be put into a forwarding state.
spanning-tree cost *cost*	Specify the spanning-tree path cost for the port. (Range: 0–200,000,000). The default cost is 0, which signifies that the cost is automatically calculated based on port speed.
CTRL + Z	Exit to Privileged EXEC mode.
show spanning-tree *interface*	View spanning tree configuration information for the specified port or LAG (port channel).

Configuring MSTP Switch Settings

Beginning in Privileged EXEC mode, use the following commands to configure MSTP settings for the switch.

Command	Purpose
configure	Enter global configuration mode.
spanning-tree mst configuration	Enable configuring an MST region by entering the multiple spanning-tree (MST) mode.
name *string*	Define the MST configuration name
revision *version*	Identify the MST configuration revision number.
instance *instance-id* {**add** | **remove**} **vlan** *vlan-range*	Map VLANs to an MST instance. • *instance-ID* — ID of the MST instance. (Range: 1-4094) • *vlan-range* — VLANs to be added to the existing MST instance. To specify a range of VLANs, use a hyphen. To specify a series of VLANs, use a comma. (Range: 1-4093)
exit	Return to global configuration mode.
spanning-tree mst *instance-id* **priority** *priority*	Set the switch priority for the specified spanning-tree instance. • *instance-id* — ID of the spanning-tree instance. (Range: 1-4094) • *priority* — Sets the switch priority for the specified spanning-tree instance. This setting affects the likelihood that the switch is selected as the root switch. A lower value increases the probability that the switch is selected as the root switch. (Range: 0-61440)
CTRL + Z	Exit to Privileged EXEC mode.
show spanning-tree mst-configuration	View multiple spanning tree configuration information.
show spanning-tree instance *instance-id*	View information about the specified MSTI.

Configuring MSTP Interface Settings

Beginning in Privileged EXEC mode, use the following commands to configure MSTP settings for the switch.

Command	Purpose
configure	Enter global configuration mode.
interface *interface*	Enter interface configuration mode for the specified interface. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3** or **port-channel 4**. You can also specify a range of interfaces with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12. The **range** keyword is also valid for LAGs (port channels).
spanning-tree mst 0 external-cost *cost*	Set the external cost for the common spanning tree. (Range: 0–200000000)
spanning-tree mst *instance-id* **cost** *cost*	Configure the path cost for MST calculations. If a loop occurs, the spanning tree considers path cost when selecting an interface to put in the forwarding state. • *instance-ID* — ID of the spanning -tree instance. (Range: 1-4094) • *cost* — The port path cost. (Range: 0–200,000,000)
spanning-tree mst *instance-id* **port-priority** *priority*	Specify the priority of the port. The priority value is used to determine which ports are put in the forwarding state and which ports are put in the blocking state. A port with a lower priority value is more likely to be put into a forwarding state. • *instance-ID* — ID of the spanning-tree instance. (Range: 1-4094) • *priority* — The port priority. (Range: 0–240 in multiples of 16)
CTRL + Z	Exit to Privileged EXEC mode.
show spanning-tree *interface* **instance** *instance-id*	View MST configuration information for the specified port or LAG (port channel) and instance.

STP Configuration Examples

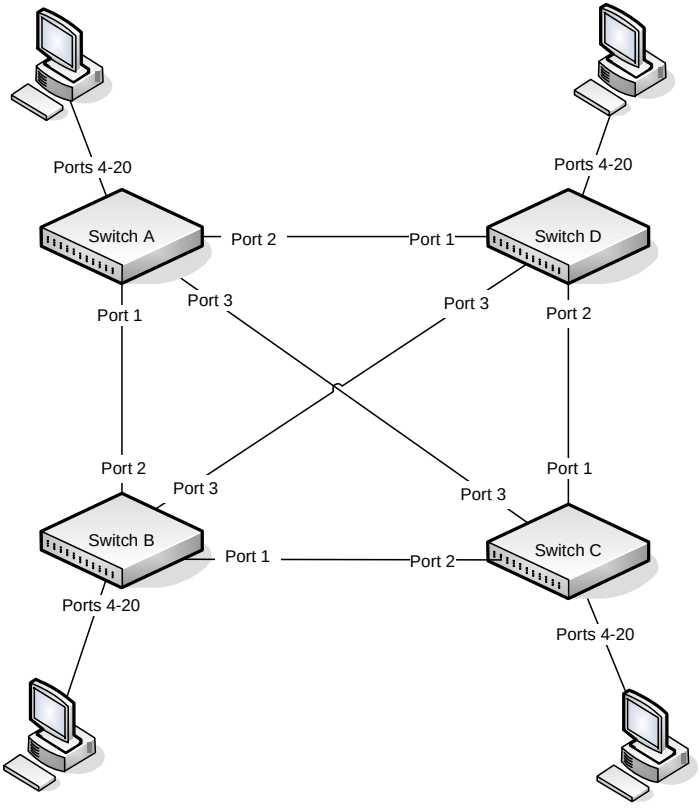
This section contains the following examples:

- Configuring STP
- Configuring MSTP

Configuring STP

This example shows a LAN with four switches. On each switch, ports 1, 2, and 3 connect to other switches, and ports 4–20 connect to hosts (in Figure 23-16, each PC represents 17 host systems).

Figure 23-16. STP Example Network Diagram



Of the four switches in Figure 23-16, the administrator decides that Switch A is the most centrally located in the network and is the least likely to be moved or redeployed. For these reasons, the administrator selects it as the root bridge for the spanning tree. The administrator configures Switch A with the highest priority and uses the default priority values for Switch B, Switch C, and Switch D.

For all switches, the administrator also configures ports 4–17 in Port Fast mode because these ports are connected to hosts and can transition directly to the Forwarding state to speed up the connection time between the hosts and the network.

The administrator also configures Port Fast BPDU filtering and Loop Guard to extend STP's capability to prevent network loops. For all other STP settings, the administrator uses the default STP values.

To configure the switch:

1. Connect to Switch A and configure the priority to be higher (a lower value) than the other switches, which use the default value of 32768.

   ```
   console#config
   console(config)#spanning-tree priority 8192
   ```

2. Configure ports 4–20 to be in Port Fast mode.

   ```
   console(config)#interface range te1/0/4-20
   console(config-if)#spanning-tree portfast
   console(config-if)#exit
   ```

3. Enable Loop Guard on ports 1–3 to help prevent network loops that might be caused if a port quits receiving BPDUs.

   ```
   console(config)#interface range te1/0/1-3
   console(config-if)#spanning-tree guard loop
   console(config-if)#exit
   ```

4. Enable Port Fast BPDU Filter. This feature is configured globally, but it affects only Port Fast-enabled access ports.

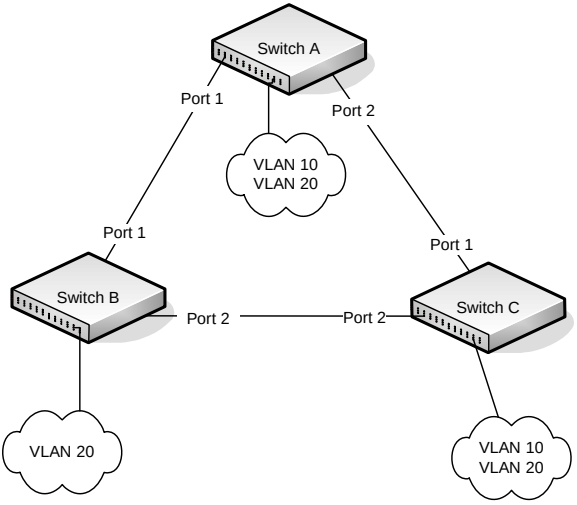
   ```
   console(config)#spanning-tree portfast bpdufilter
   default
   ```

5. Repeat step 2 through step 4 on Switch B, Switch C, and Switch D to complete the configuration.

Configuring MSTP

This example shows how to configure IEEE 802.1s Multiple Spanning Tree (MST) protocol on the switches shown in Figure 23-17.

Figure 23-17. MSTP Configuration Example



To make multiple switches be part of the same MSTP region, make sure the STP operational mode for all switches is MSTP. Also, make sure the MST region name and revision level are the same for all switches in the region.

To configure the switches:

1 Create VLAN 10 (Switch A and Switch B) and VLAN 20 (all switches).

NOTE: Even Switch B does not have any ports that are members of VLAN 10, this VLAN must be created to allow the formation of MST regions made up of all bridges that exchange the same MST Configuration Identifier. It is only within these MST Regions that multiple instances can exist.

```
console#configure
console(config)#vlan database
console(config-vlan)#vlan 10
console(config-vlan)#vlan 20
console(config-vlan)#exit
```

2 Set the STP operational mode to MSTP.

```
console(config)#spanning-tree mode mst
```

3 Create MST instance 10 and associate it to VLAN 10.

```
console(config)#spanning-tree mst configuration
console(config-mst)#instance 10 add vlan 10
```

4 Create MST instances 20 and associate it to VLAN 20.

```
console(config-mst)#instance 20 add vlan 20
```

5 Change the region name so that all the bridges that want to be part of the same region can form the region.

```
console(config-mst)#name dell
console(config-mst)#exit
```

6 (Switch A only) Configure Switch A to be the root bridge of the spanning tree (CIST Regional Root) by configuring a higher root bridge priority.

```
console(config)#spanning-tree priority 8192
```

7 (Switch A only) Make Switch A the Regional Root for MSTI 1 by configuring a higher priority for MST ID 10.

```
console(config)#spanning-tree mst 10 priority
12288
```

8 (Switch A only) Change the priority of MST ID 20 to ensure Switch C is the Regional Root bridge for this MSTI.

```
console(config)#spanning-tree mst 20 priority
61440
console(config)#spanning-tree priority 8192
```

9 (Switch C only) Change the priority of port 1 to force it to be the root port for MST 20.

```
console(config)#interface te1/0/1
console(config-if-Te1/0/1)#spanning-tree mst 20
port-priority 64
console(config-if-Te1/0/1)#exit
```

24

Discovering Network Devices

This chapter describes the Industry Standard Discovery Protocol (ISDP) feature and the Link Layer Discovery Protocol (LLDP) feature, including LLDP for Media Endpoint Devices (LLDP-MED).

The topics covered in this chapter include:

- Device Discovery Overview
- Default IDSP and LLDP Values
- Configuring ISDP and LLDP (Web)
- Configuring ISDP and LLDP (CLI)
- Device Discovery Configuration Examples

Device Discovery Overview

The switch software includes two different device discovery protocols: IDSP and LLDP. These protocols allow the switch to broadcast information about itself and to learn information about neighboring devices.

What Is ISDP?

The Industry Standard Discovery Protocol (ISDP) is a proprietary Layer 2 network protocol that inter-operates with Cisco devices running the Cisco Discovery Protocol (CDP). ISDP is used to share information between neighboring devices. The switch software participates in the CDP protocol and is able to both discover and be discovered by other CDP-supporting devices.

What is LLDP?

LLDP is a standardized discovery protocol defined by IEEE 802.1AB. It allows stations residing on an 802 LAN to advertise major capabilities physical descriptions, and management information to physically adjacent devices allowing a network management system (NMS) to access and display this information.

LLDP is a one-way protocol; there are no request/response sequences. Information is advertised by stations implementing the transmit function, and is received and processed by stations implementing the receive function. The transmit and receive functions can be enabled/disabled separately on each switch port.

What is LLDP-MED?

LLDP-MED is an extension of the LLDP standard. LLDP-MED uses LLDP's organizationally-specific Type- Length-Value (TLV) extensions and defines new TLVs that make it easier for a VoIP deployment in a wired or wireless LAN/MAN environment. It also makes mandatory a few optional TLVs from LLDP and recommends not transmitting some TLVs.

The TLVs only communicate information; these TLVs do not automatically translate into configuration. An external application may query the MED MIB and take management actions in configuring functionality.

Why are Device Discovery Protocols Needed?

The device discovery protocols are used primarily in conjunction with network management tools to provide information about network topology and configuration, and to help troubleshoot problems that occur on the network. The discovery protocols can also facilitate inventory management within a company.

LLDP and the LLDP-MED extension are vendor-neutral discovery protocols that can discover devices made by numerous vendors. LLDP-MED is intended to be used on ports that connect to VoIP phones. Additional applications for LLDP-MED include device location (including for Emergency Call Service/E911) and Power over Ethernet management.

ISDP interoperates with the Cisco-proprietary CDP protocol and is most effective in an environment that contains many Cisco devices.

Default IDSP and LLDP Values

ISDP and LLDP are globally enabled on the switch and enabled on all ports by default. By default, the switch transmits and receives LLDP information on all ports. LLDP-MED is disabled on all ports.

Table 24-1 summarizes the default values for ISDP.

Table 24-1. ISDP Defaults

Parameter	Default Value
ISDP Mode	Enabled (globally and on all ports)
ISDPv2 Mode	Enabled (globally and on all ports)
Message Interval	30 seconds
Hold Time Interval	180 seconds
Device ID	none
Device ID Format Capability	Serial Number, Host Name
Device ID Format	Serial Number

Table 24-2 summarizes the default values for LLDP.

Table 24-2. LLDP Defaults

Parameter	Default Value
Transmit Mode	Enabled on all ports
Receive Mode	Enabled on all ports
Transmit Interval	30 seconds
Hold Multiplier	4
Reinitialization Delay	2 seconds
Notification Interval	5 seconds
Transmit Management Information	Disabled
Notification Mode	Disabled
Included TLVs	None

Table 24-3 summarizes the default values for LLDP-MED.

Table 24-3. LLDP-MED Defaults

Parameter	Default Value
LLDP-MED Mode	Disabled on all ports
Config Notification Mode	Disabled on all ports
Transmit TVLs	MED Capabilities Network Policy

Configuring ISDP and LLDP (Web)

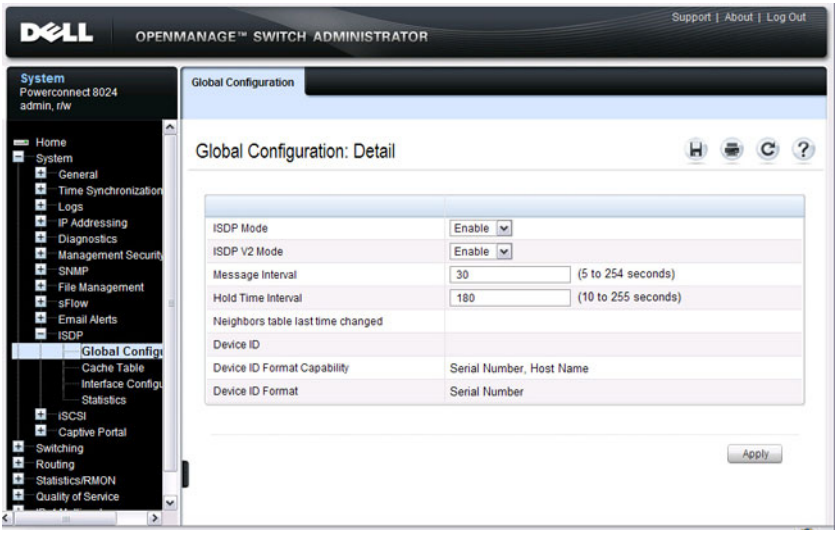
This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring IDSP and LLDP/LLDP-MED on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

ISDP Global Configuration

From the **ISDP Global Configuration** page, you can configure the ISDP settings for the switch, such as the administrative mode.

To access the **ISDP Global Configuration** page, click **System** → **ISDP** → **Global Configuration** in the navigation panel.

Figure 24-1. ISDP Global Configuration

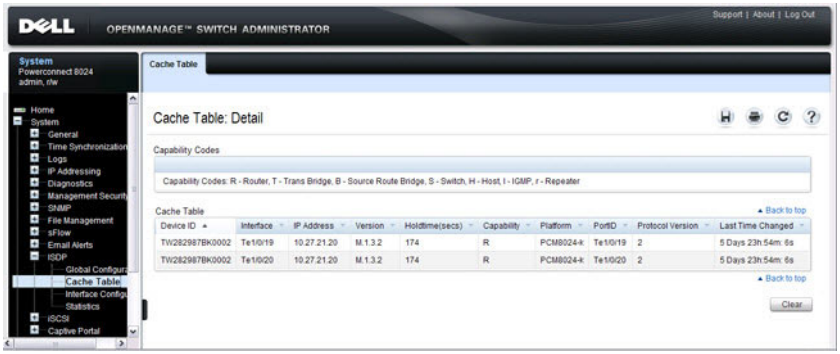


ISDP Cache Table

From the **ISDP Cache Table** page, you can view information about other devices the switch has discovered through the ISDP.

To access the **ISDP Cache Table** page, click **System → ISDP → Cache Table** in the navigation panel.

Figure 24-2. ISDP Cache Table



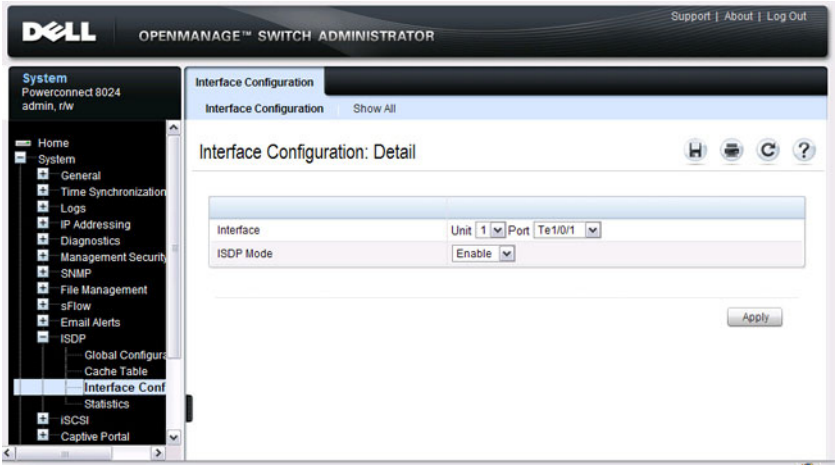
ISDP Interface Configuration

From the **ISDP Interface Configuration** page, you can configure the ISDP settings for each interface.

If ISDP is enabled on an interface, it must also be enabled globally in order for the interface to transmit ISDP packets. If the ISDP mode on the ISDP Global Configuration page is disabled, the interface will not transmit ISDP packets, regardless of the mode configured on the interface.

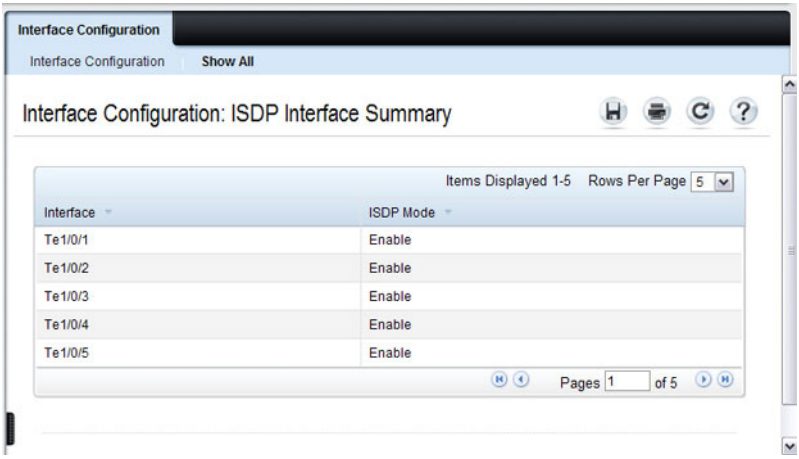
To access the **ISDP Interface Configuration** page, click **System** → **ISDP** → **Interface Configuration** in the navigation panel.

Figure 24-3. ISDP Interface Configuration



To view view the ISDP mode for multiple interfaces, click **Show All**.

Figure 24-4. ISDP Interface Summary

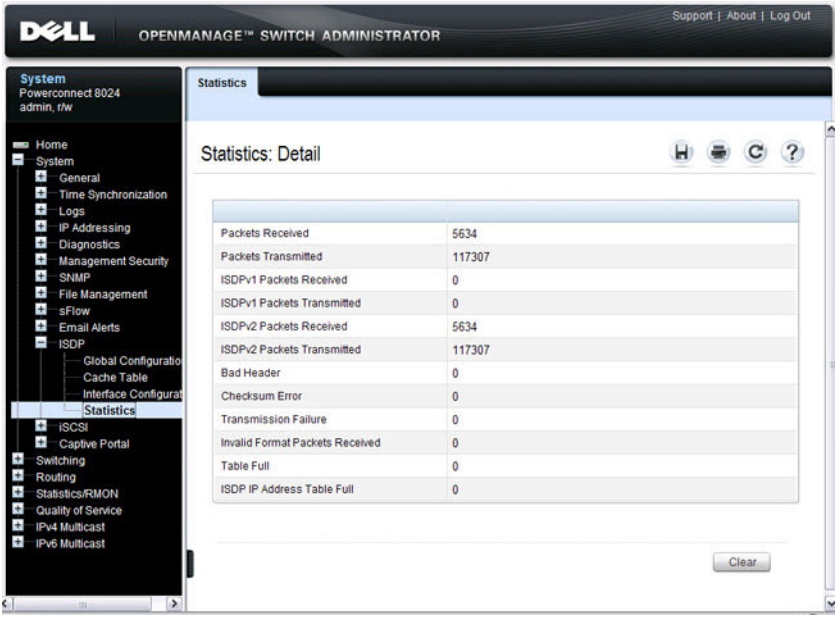


ISDP Statistics

From the **ISDP Statistics** page, you can view information about the ISDP packets sent and received by the switch.

To access the **ISDP Statistics** page, click **System** → **ISDP** → **Statistics** in the navigation panel.

Figure 24-5. ISDP Statistics

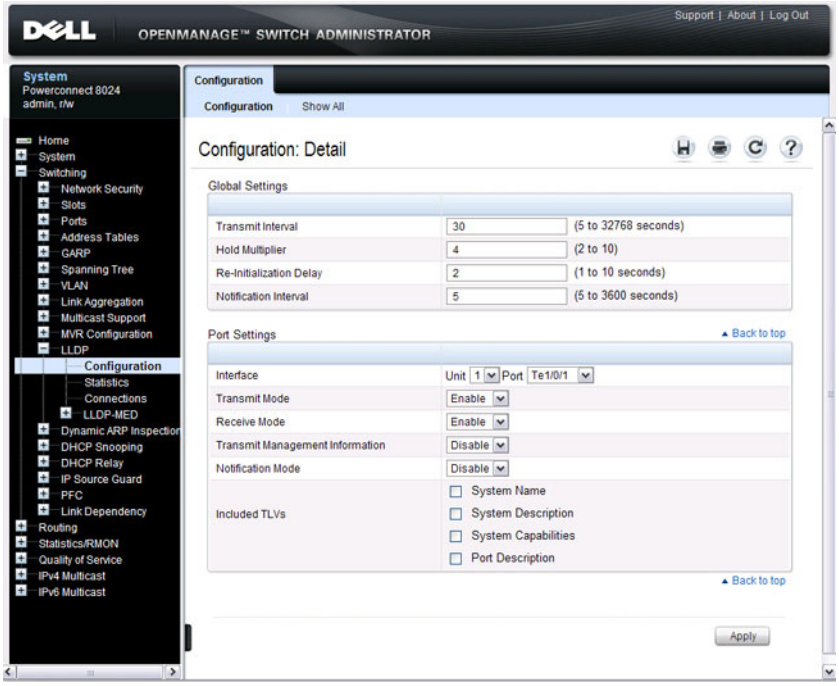


LLDP Configuration

Use the **LLDP Configuration** page to specify LLDP parameters. Parameters that affect the entire system as well as those for a specific interface can be specified here.

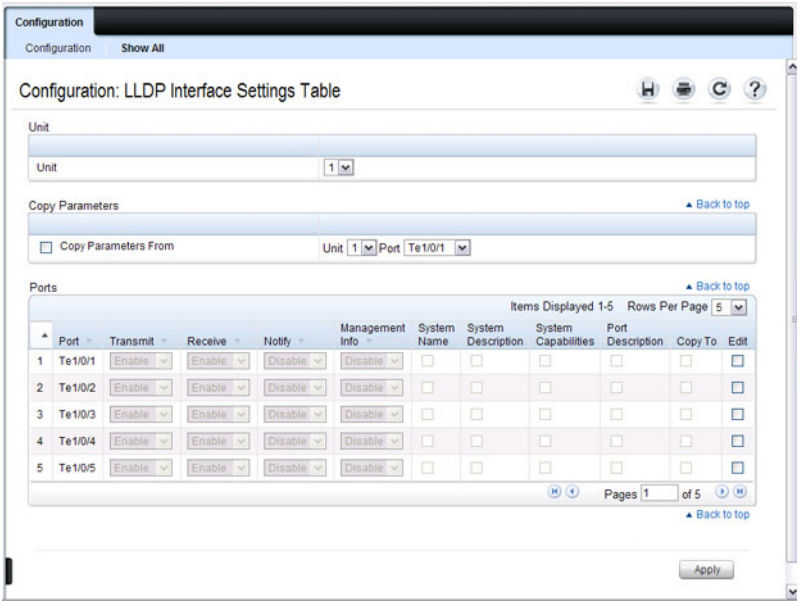
To display the **LLDP Configuration** page, click **Switching** → **LLDP** → **Configuration** in the navigation panel.

Figure 24-6. LLDP Configuration



To view the **LLDP Interface Settings Table**, click **Show All**. From the LLDP Interface Settings Table page, you can view and edit information about the LLDP settings for multiple interfaces.

Figure 24-7. LLDP Interface Settings Table

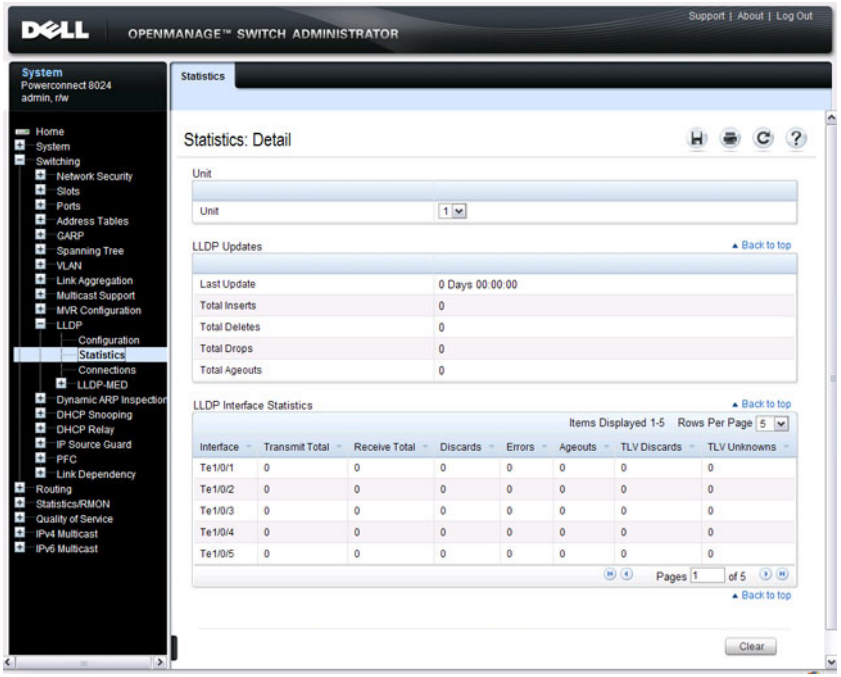


LLDP Statistics

Use the LLDP Statistics page to view LLPD-related statistics.

To display the **LLDP Statistics** page, click **Switching** → **LLDP** → **Statistics** in the navigation panel.

Figure 24-8. LLDP Statistics

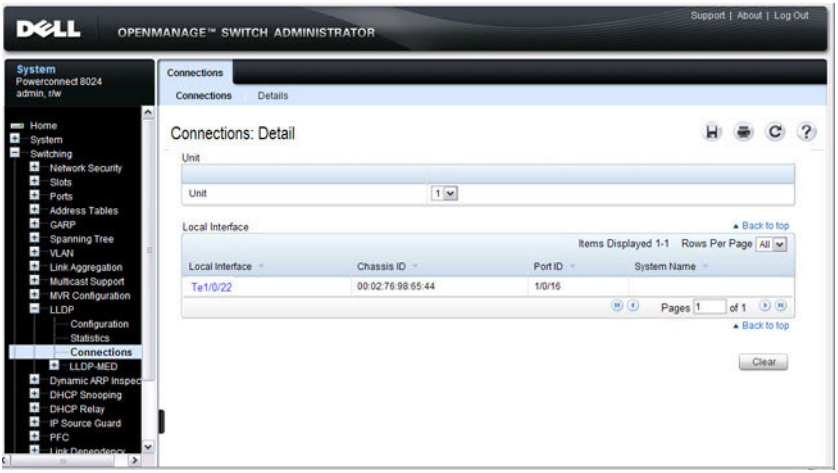


LLDP Connections

Use the **LLDP Connections** page to view the list of ports with LLDP enabled. Basic connection details are displayed.

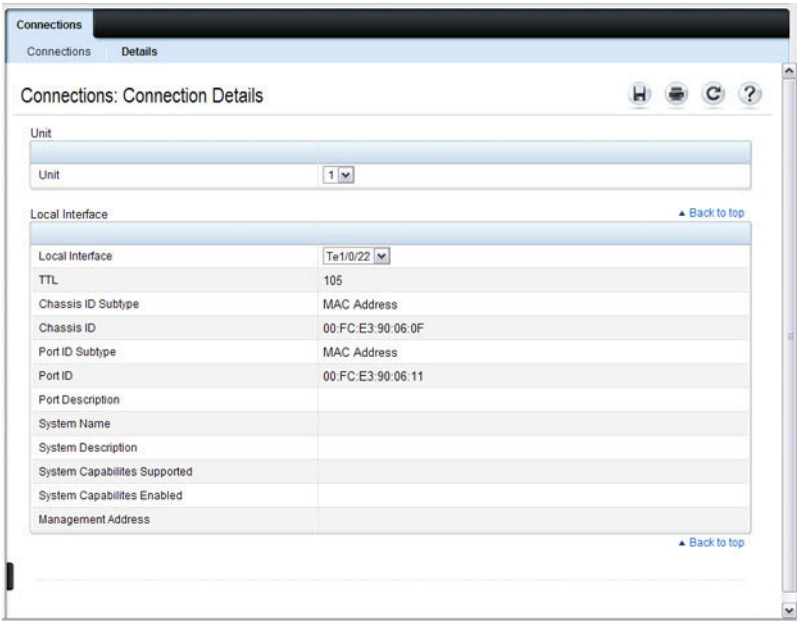
To display the **LLDP Connections** page, click **Switching** → **LLDP** → **Connections** in the navigation panel.

Figure 24-9. LLDP Connections



To view additional information about a device connected to a port that has been discovered through LLDP, click the port number in the Local Interface table (it is a hyperlink), or click **Details** and select the port with the connected device.

Figure 24-10. LLDP Connection Detail

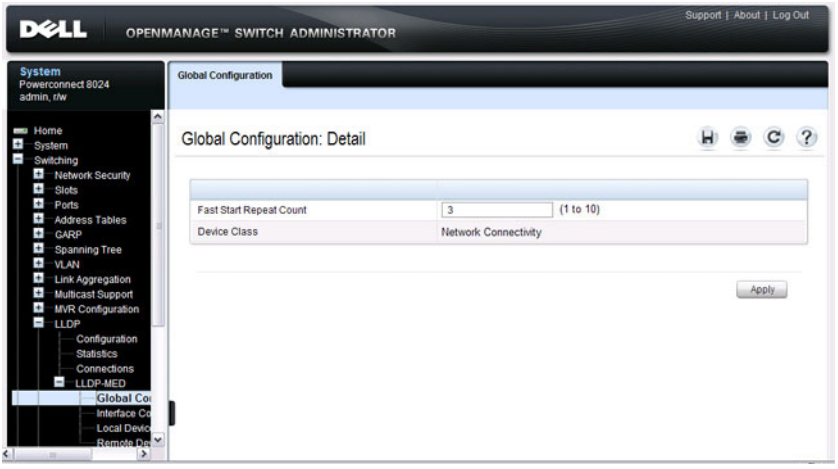


LLDP-MED Global Configuration

Use the **LLDP-MED Global Configuration** page to change or view the LLDP-MED parameters that affect the entire system.

To display the **LLDP-MED Global Configuration** page, click **Switching→ LLDP → LLDP-MED → Global Configuration** in the navigation panel.

Figure 24-11. LLDP-MED Global Configuration

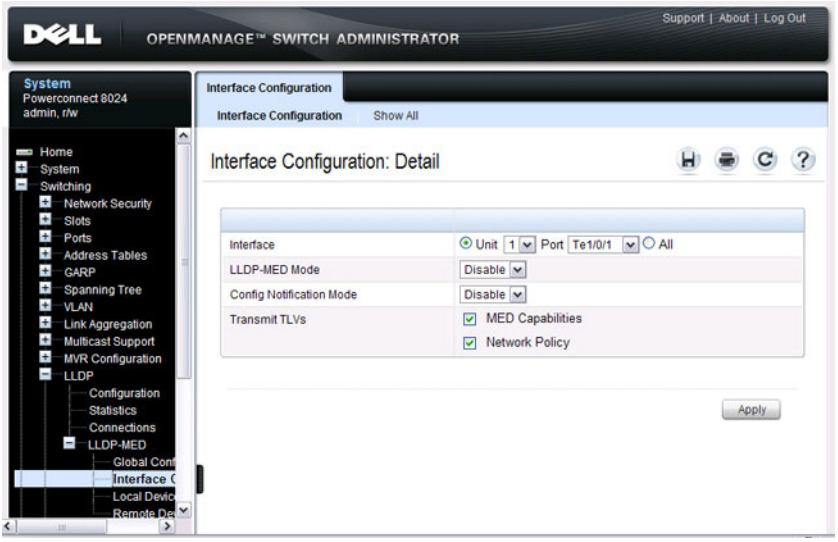


LLDP-MED Interface Configuration

Use the **LLDP-MED Interface Configuration** page to specify LLDP-MED parameters that affect a specific interface.

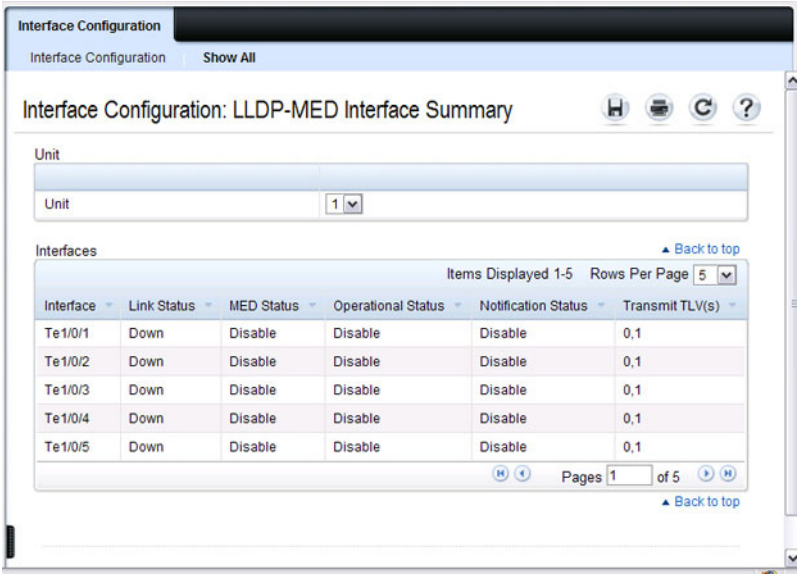
To display the **LLDP-MED Interface Configuration** page, click **Switching** → **LLDP** → **LLDP-MED** → **Interface Configuration** in the navigation panel.

Figure 24-12. LLDP-MED Interface Configuration



To view the **LLDP-MED Interface Summary** table, click **Show All**.

Figure 24-13. LLDP-MED Interface Summary

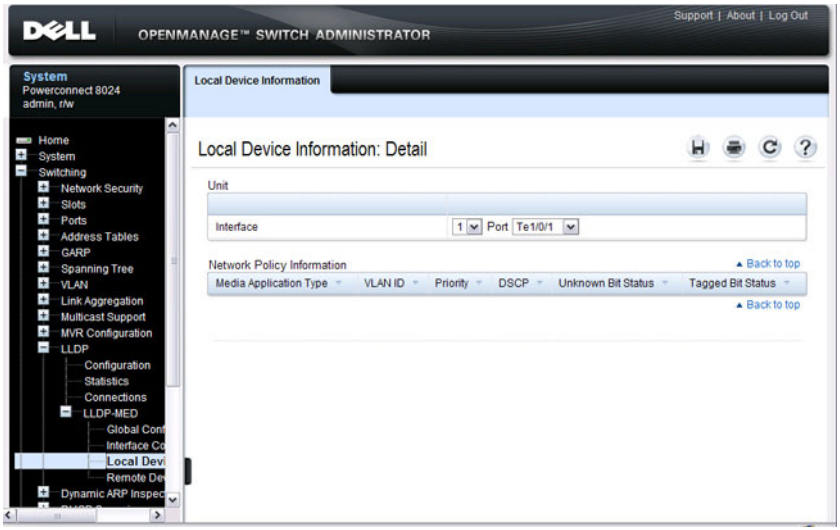


LLDP-MED Local Device Information

Use the **LLDP-MED Local Device Information** page to view the advertised LLDP local data for each port.

To display the **LLDP-MED Local Device Information** page, click **Switching→ LLDP→ LLDP-MED→ Local Device Information** in the navigation panel.

Figure 24-14. LLDP-MED Local Device Information

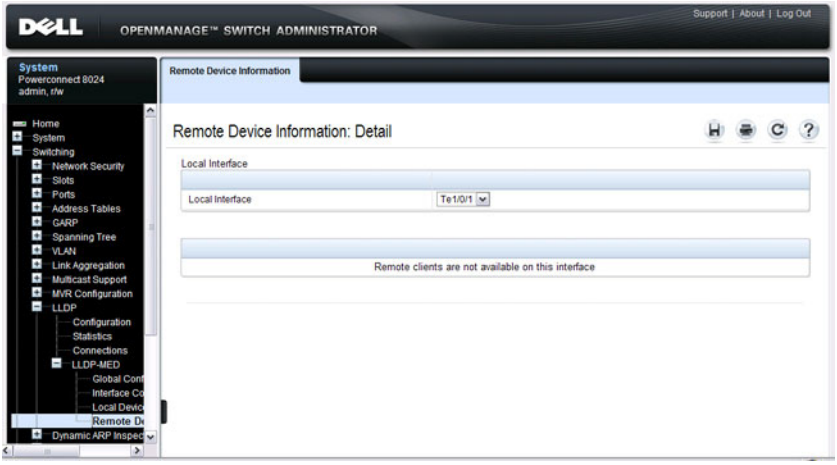


LLDP-MED Remote Device Information

Use the **LLDP-MED Remote Device Information** page to view the advertised LLDP data advertised by remote devices.

To display the **LLDP-MED Remote Device Information** page, click **Switching→ LLDP→ LLDP-MED→ Remote Device Information** in the navigation panel.

Figure 24-15. LLDP-MED Remote Device Information



Configuring ISDP and LLDP (CLI)

This section provides information about the commands you use to manage and view the device discovery protocol features on the switch. For more information about these commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Configuring Global ISDP Settings

Beginning in Privileged EXEC mode, use the following commands to configure ISDP settings that affect the entire switch.

Command	Purpose
configure	Enter Global Configuration mode.
isdp enable	Administratively enable ISDP on the switch.
isdp advertise-v2	Allow the switch to send ISDPv2 packets.
isdp holdtime *time*	Specify the number of seconds the device that receives ISDP packets from the switch should store information sent in the ISDP packet before discarding it.
isdp timer *time*	Specify the number of seconds to wait between sending new ISDP packets.
exit	Exit to Privileged EXEC mode.
show isdp	View global ISDP settings.

Enabling ISDP on a Port

Beginning in Privileged EXEC mode, use the following commands to enable ISDP on a port.

Command	Purpose
configure	Enter Global Configuration mode.
interface *interface*	Enter interface configuration mode for the specified interface.
isdp enable	Administratively enable ISDP on the switch.
exit	Exit to Global Config mode.
exit	Exit to Privileged Exec mode.
show isdp interface all	View the ISDP mode on all interfaces.

Viewing and Clearing ISDP Information

Beginning in Privileged EXEC mode, use the following commands to view and clear the contents of the ISDP table and to view and clear ISDP statistics.

Command	Purpose
show isdp entry {**all** | *deviceid*}	View information about all entries or a specific entry in the ISDP table.
show isdp neighbors	View the neighboring devices discovered through ISDP.
clear isdp table	Clear all entries, including discovered neighbors, from the ISDP table.
show isdp traffic	View ISDP statistics.
clear isdp counters	Reset all ISDP statistics to zero.

Configuring Global LLDP Settings

Beginning in Privileged EXEC mode, use the following commands to configure LLDP settings that affect the entire switch.

Command	Purpose
configure	Enter Global Configuration mode.
lldp notification-interval *interval*	Specify how often, in seconds, the switch should send remote data change notifications.
lldp timers [interval *transmit-interval*] [**hold** *hold-value*] [**reinit** *reinit-delay*]	Configure the timing for local data transmission on ports enabled for LLDP. • *transmit-interval* — The interval in seconds at which to transmit local data LLDP PDUs. (Range: 5–32768 seconds) • *hold-value* — Multiplier on the transmit interval used to set the TTL in local data LLDP PDUs. (Range: 2–10) • *reinit-delay* — The delay in seconds before re-initialization. (Range: 1–10 seconds)
exit	Exit to Privileged EXEC mode.
show lldp	View global LLDP settings.

Configuring Port-based LLDP Settings

Beginning in Privileged EXEC mode, use the following commands to configure per-port LLDP settings.

Command	Purpose
configure	Enter Global Configuration mode.
interface *interface*	Enter interface configuration mode for the specified Ethernet interface.
lldp transmit	Enable the LLDP advertise (transmit) capability.
lldp receive	Enable the LLDP receive capability so that the switch can receive LLDP Protocol Data Units (LLDP PDUs) from other devices.
lldp transmit-mgmt	Include the transmission of local system management address information in the LLDP PDUs.

Command	Purpose
lldp notification	Enable remote data change notifications on the interface.
lldp transmit-tlv [**sys-desc**] [**sys-name**] [**sys-cap**] [**port-desc**]	Specify which optional type-length-value settings (TLVs) in the 802.1AB basic management set will be transmitted in the LLDP PDUs. • sys-name — Transmits the system name TLV • sys-desc — Transmits the system description TLV • sys-cap — Transmits the system capabilities TLV • port desc — Transmits the port description TLV
exit	Exit to Global Config mode.
exit	Exit to Privileged EXEC mode.
show lldp interface all	View LLDP settings for all interfaces.

Viewing and Clearing LLDP Information

Beginning in Privileged EXEC mode, use the following commands to view transmitted and received LLDP information and to view and clear LLDP statistics.

Command	Purpose
show lldp local-device {**all** | *interface* | **detail** *interface*}	View LLDP information advertised by all ports or the specified port. Include the keyword detail to see additional information.
show lldp remote-device {**all** | *interface* | **detail** *interface*}	View LLDP information received by all ports or by the specified port. Include the keyword detail to see additional information.
clear lldp remote-data	Delete all LLDP information from the remote data table.
show lldp statistics	View LLDP traffic statistics.
clear lldp statistics	Reset the LLDP statistics counters to zero.

Configuring LLDP-MED Settings

Beginning in Privileged EXEC mode, use the following commands to configure LLDP-MED settings that affect the entire switch.

Command	Purpose
configure	Enter Global Configuration mode.
lldp med faststartrepeatcount *count*	Specifies the number of LLDP PDUs that will be transmitted when the protocol is enabled.
interface *interface*	Enter interface configuration mode for the specified Ethernet interface.
lldp med	Enable LLDP-MED on the interface.
lldp med confignotification	Allow the port to send topology change notifications.
lldp med transmit-tlv [**capabilities**] [**network-policy**] [**location**] [**inventory**]	Specify which optional TLVs in the LLDP MED set are transmitted in the LLDP PDUs.
exit	Exit to Global Config mode.
exit	Exit to Privileged EXEC mode.
show lldp med	View global LLDP-MED settings.
show lldp med interface {**all** | *interface*}	View LLDP-MED settings for all ports or for the specified port.

Viewing LLDP-MED Information

Beginning in Privileged EXEC mode, use the following commands to view information about the LLDP-MED Protocol Data Units (PDUs) that are sent and have been received.

Command	Purpose
show lldp med local-device detail *interface*	View LLDP information advertised by the specified port.
show lldp remote-device {**all** | *interface* | **detail** *interface*}	View LLDP-MED information received by all ports or by the specified port. Include the keyword detail to see additional information.

Device Discovery Configuration Examples

This section contains the following examples:

- Configuring ISDP
- Configuring LLDP

Configuring ISDP

This example shows how to configure ISDP settings on the switch.

To configure the switch:

1. Specify the number of seconds that a remote device should keep the ISDP information sent by the switch before discarding it.

   ```
   console#configure
   console(config)#isdp holdtime 60
   ```

2. Specify how often, in seconds, the ISDP-enabled ports should transmit information.

   ```
   console(config)#isdp timer 45
   ```

3. Enable ISDP on interface 1/0/3.

   ```
   console(config)#interface tengigabitEthernet1/0/3
   console(config-if-Te1/0/3)#isdp enable
   ```

4 Exit to Privileged EXEC mode and view the LLDP settings for the switch and for interface 1/0/3.

```
console(config-if-Te1/0/3)# <CTRL + Z>
console#show isdp
Timer.....................................45
Hold Time.................................60
Version 2 Advertisements.................Enabled
Neighbors table time since last change...00 days
                                          00:00:00
Device ID.................................none
Device ID format capability........ Serial Number,
                                    Host Name
Device ID format.................. Serial Number

console#show isdp interface te1/0/3

Interface        Mode
---------------  ----------
Te1/0/3          Enabled
```

Configuring LLDP

This example shows how to configure LLDP settings for the switch and to allow 10-Gigabit Ethernet port 1/0/3 to transmit all LLDP information available.

To configure the switch:

1 Configure the transmission interval, hold multiplier, and reinitialization delay for LLDP PDUs sent from the switch.

```
console#configure
console(config)#lldp timers interval 60 hold 5
reinit 3
```

2 Enable port 1/0/3 to transmit and receive LLDP PDUs.

```
console(config)#interface TengigabitEthernet1/0/3
console(config-if-Te1/0/3)#lldp transmit
console(config-if-Te1/0/3)#lldp receive
```

3 Enable port 1/0/3 to transmit management address information in the LLDP PDUs and to send topology change notifications if a device is added or removed from the port.

```
console(config-if-Te1/0/3)#lldp transmit-mgmt
console(config-if-Te1/0/3)#lldp notification
```

4 Specify the TLV information to be included in the LLDP PDUs transmitted from port 1/0/3.

```
console(config-if-Te1/0/3)#lldp transmit-tlv sys-
name sys-desc sys-cap port-desc
```

5 Set the port description to be transmitted in LLDP PDUs.

```
console(config-if-Te1/0/3)#description "Test Lab
Port"
```

6 Exit to Privileged EXEC mode.

```
console(config-if-Te1/0/3)# <CTRL + Z>
```

7 View global LLDP settings on the switch.

```
console#show lldp

LLDP Global Configuration

Transmit Interval..................... 60 seconds
Transmit Hold Multiplier.............. 5
Reinit Delay.......................... 3 seconds
Notification Interval................. 5 seconds
```

8 View summary information about the LLDP configuration on port 1/0/3.

```
console#show lldp interface te1/0/3

LLDP Interface Configuration

Interface  Link    Transmit  Receive   Notify    TLVs     Mgmt
---------  ------  --------  --------  --------  -------  ----
Te1/0/3    Down    Enabled   Enabled   Enabled   0,1,2,3  Y

TLV Codes: 0- Port Description,   1- System Name
           2- System Description, 3- System Capabilities
```

9 View detailed information about the LLDP configuration on port 1/0/3.

```
console#show lldp local-device detail te1/0/3

LLDP Local Device Detail

Interface: Te1/0/3

Chassis ID Subtype: MAC Address
Chassis ID: 00:1E:C9:AA:AA:07
Port ID Subtype: Interface Name
Port ID: te 1/0/3
System Name: console
System Description: PowerConnect 8024 3.16.22.30,
VxWorks 6.5
Port Description: Test Lab Port
System Capabilities Supported: bridge, router
System Capabilities Enabled: bridge
Management Address:
    Type: IPv4
    Address: 192.168.2.1
```

25

Configuring Port-Based Traffic Control

This chapter describes how to configure features that provide traffic control through filtering the type of traffic or limiting the speed or amount of traffic on a per-port basis. The features this section describes includes flow control, storm control, protected ports, and Link Local Protocol Filtering (LLPF), which is also known as Cisco Protocol Filtering.

The topics covered in this chapter include:

- Port-Based Traffic Control Overview
- Default Port-Based Traffic Control Values
- Configuring Port-Based Traffic Control (Web)
- Configuring Port-Based Traffic Control (CLI)
- Port-Based Traffic Control Configuration Example

Port-Based Traffic Control Overview

Table 25-1 provides a summary of the features this chapter describes.

Table 25-1. Port-Based Traffic Control Features

Feature	Description
Flow control	Allows traffic transmission between a switch port and another Ethernet device to be paused for a specified period of time when congestion occurs.
Storm control	Limits the amount of broadcast, unknown unicast, and multicast frames accepted and forwarded by the switch.
Protected ports	Prevents traffic from flowing between members of the same protected port group.
LLPF	Filters proprietary protocols that should not normally be relayed by a bridge.

For information about Priority Flow Control (PFC), which provides a way to distinguish which traffic on a physical link is paused when congestion occurs based on the priority of the traffic, see "Configuring Data Center Bridging Features" on page 799

What is Flow Control?

IEEE 802.3 Annex 31B flow control allows nodes that transmit at slower speeds to communicate with higher speed switches by requesting that the higher speed switch refrain from sending packets. Transmissions are temporarily halted to prevent buffer overflows. Enabling the flow control feature allows PowerConnect 8000-series and 8100-series switches to process pause frames received from connected devices. PowerConnect switches do not transmit pause frames.

Flow control is supported only on ports that are configured for full-duplex mode of operation. Since ports set to auto negotiate may not be added as LAG members, LAG member ports cannot have flow control configured to auto.

What is Storm Control?

A LAN storm is the result of an excessive number of broadcast, multicast, or unknown unicast messages simultaneously transmitted across a network by a single port. Forwarded message responses can overload network resources and cause network congestion.

The storm control feature allows the switch to measure the incoming broadcast, multicast, and/or unknown unicast packet rate per port and discard packets when the rate exceeds the defined threshold. Storm control is enabled per interface, by defining the packet type and the rate at which the packets are transmitted. For each type of traffic (broadcast, multicast, or unknown unicast) you can configure a threshold level, which is expressed as a percentage of the total available bandwidth on the port. If the ingress rate of that type of packet is greater than the configured threshold level the port drops the excess traffic until the ingress rate for the packet type falls below the threshold.

The actual rate of ingress traffic required to activate storm-control is based on the size of incoming packets and the hard-coded average packet size of 512 bytes - used to calculate a packet-per-second (pps) rate - as the forwarding-plane requires PPS versus an absolute rate Kbps. For example, if the

configured limit is 10%, this is converted to ~25000 PPS, and this PPS limit is set in the hardware. You get the approximate desired output when 512 bytes packets are used.

What are Protected Ports?

The switch supports up to three separate groups of protected ports. Traffic can flow between protected ports belonging to different groups, but not within the same group.

A port can belong to only one protected port group. You must remove an interface from one group before adding it to another group.

Port protection occurs within a single switch. Protected port configuration does not affect traffic between ports on two different switches. No traffic forwarding is possible between two protected ports.

What is Link Local Protocol Filtering?

The Link Local Protocol Filtering (LLPF) feature can help troubleshoot network problems that occur when a network includes proprietary protocols running on standards-based switches. LLPF allows a PowerConnect 8000-series and 8100-series switches to filter out various Cisco proprietary protocol data units (PDUs) and/or ISDP if problems occur with these protocols running on standards-based switches. If certain protocol PDUs cause unexpected results, LLPF can be enabled to prevent those protocol PDUs from being processed by the switch.

The LLPF feature can be configured per-port to block any combination (or all) of the following PDUs:

- Industry Standard Discovery Protocol (ISDP)
- VLAN Trunking Protocol (VTP)
- Dynamic Trunking Protocol (DTP)
- UniDirectional Link Detection (UDLD)
- Port Aggregation Protocol (PAgP)
- Shared Spanning Tree Protocol (SSTP)

Access Control Lists (ACLs) and LLPF can exist on the same interface. However, the ACL rules override the LLPF rules when there is a conflict. Similarly, DiffServ and LLPF can both be enabled on an interface, but DiffServ rules override LLPF rules when there is a conflict.

If Industry Standard Discovery Protocol (ISDP) is enabled on an interface, and the LLPF feature on an interface is enabled and configured to drop ISDP PDUs, the ISDP configuration overrides the LLPF configuration, and the ISDP PDUs are allowed on the interface.

Default Port-Based Traffic Control Values

Table 25-2 lists the default values for the port-based traffic control features that this chapter describes.

Table 25-2. Default Port-Based Traffic Control Values

Feature	Default
Flow control	Enabled
Storm control	Disabled
Protected ports	None
LLPF	No protocols are blocked

Configuring Port-Based Traffic Control (Web)

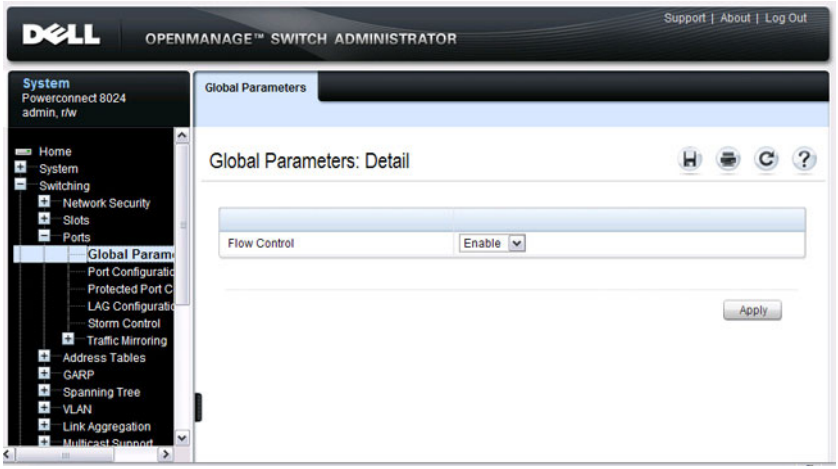
This section provides information about the OpenManage Switch Administrator pages to use to control port-based traffic on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

Flow Control (Global Port Parameters)

Use the **Global Parameters** page for ports to enable or disable flow control support on the switch.

To display the **Global Parameters** page, click **Switching** → **Ports** → **Global Parameters** in the navigation menu.

Figure 25-1. Global Port Parameters

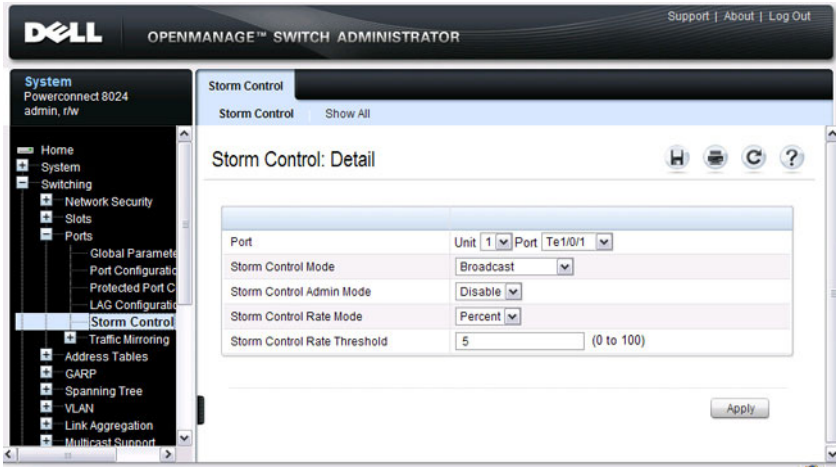


Storm Control

Use the **Storm Control** page to enable and configure the storm control feature.

To display the **Storm Control** interface, click **Switching** → **Ports** → **Storm Control** in the navigation menu.

Figure 25-2. Storm Control

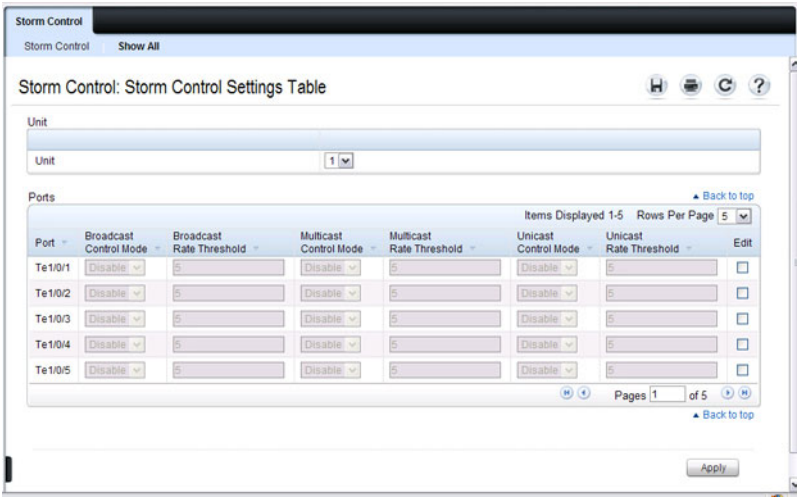


Configuring Storm Control Settings on Multiple Ports

To configure storm control on multiple ports:

1. Open the **Storm Control** page.
2. Click **Show All** to display the **Storm Control Settings Table**.
3. In the Ports list, select the check box in the **Edit** column for the port to configure.
4. Select the desired storm control settings.

Figure 25-3. Storm Control



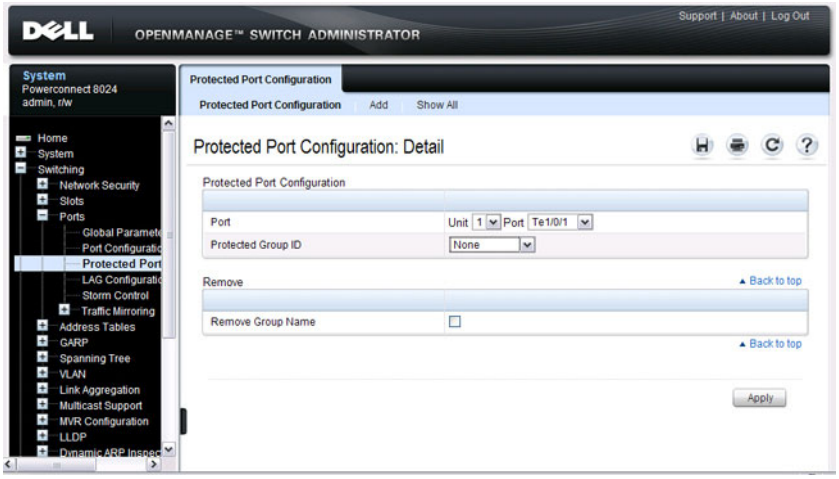
5 Click **Apply**.

Protected Port Configuration

Use the **Protected Port Configuration** page to prevent ports in the same protected ports group from being able to see each other's traffic.

To display the **Protected Port Configuration** page, click **Switching** → **Ports** → **Protected Port Configuration** in the navigation menu.

Figure 25-4. Protected Port Configuration

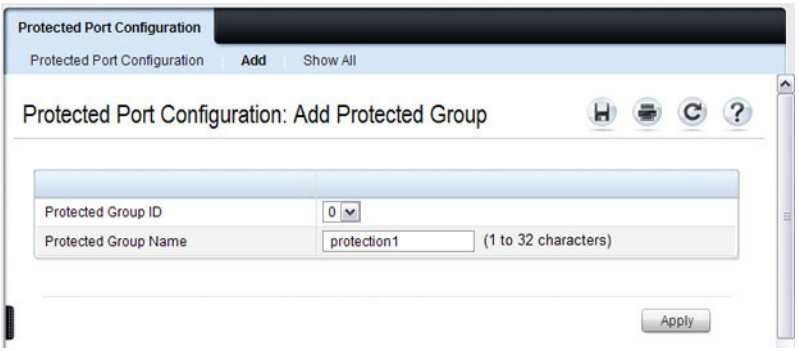


Configuring Protected Ports

To configure protected ports:

1. Open the **Protected Ports** page.
2. Click **Add** to display the **Add Protected Group** page.
3. Select a group (0–2).
4. Specify a name for the group.

Figure 25-5. Add Protected Ports Group



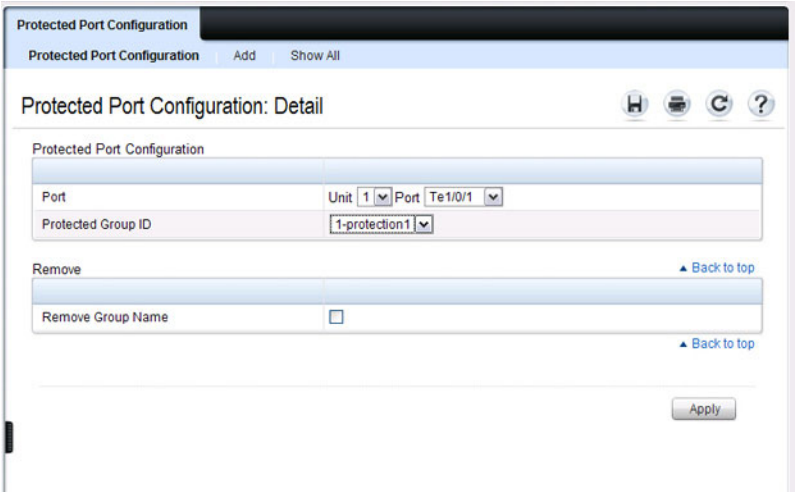
5 Click **Apply**.

6 Click **Protected Port Configuration** to return to the main page.

7 Select the port to add to the group.

8 Select the protected port group ID.

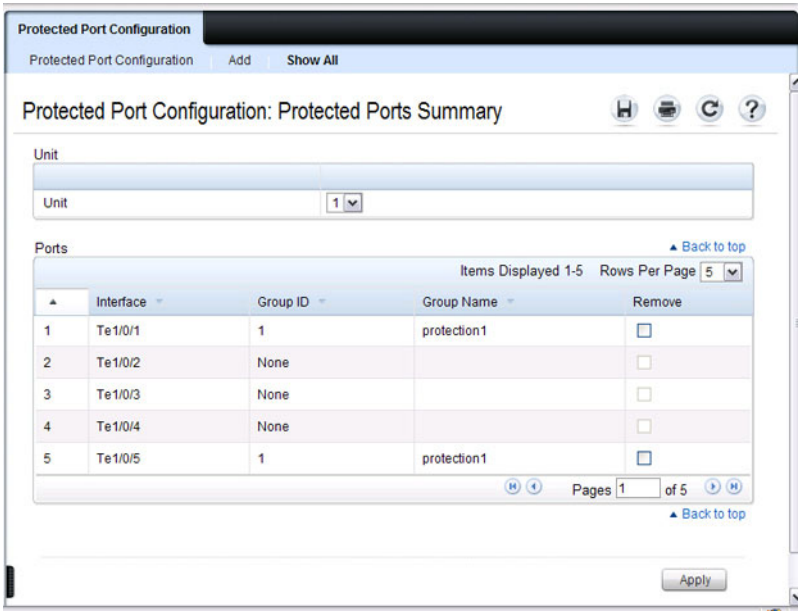
Figure 25-6. Add Protected Ports



9 Click **Apply**.

10 To view protected port group membership information, click **Show All**.

Figure 25-7. View Protected Port Information



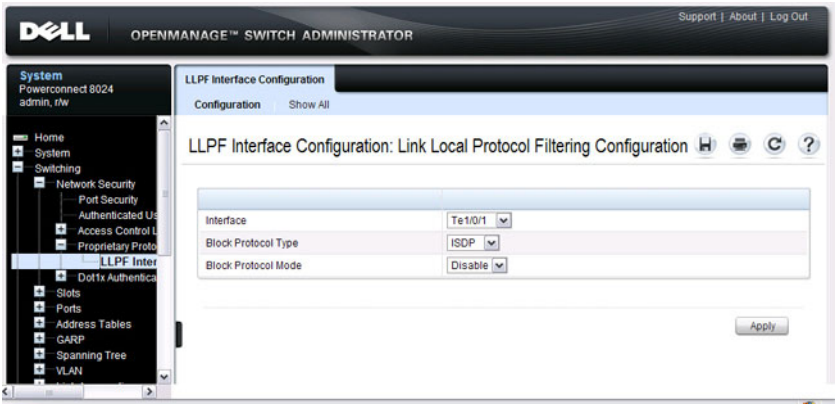
11 To remove a port from a protected port group, select the **Remove** check box associated with the port and click **Apply**.

LLPF Configuration

Use the **LLPF Interface Configuration** page to filter out various proprietary protocol data units (PDUs) and/or ISDP if problems occur with these protocols running on standards-based switches.

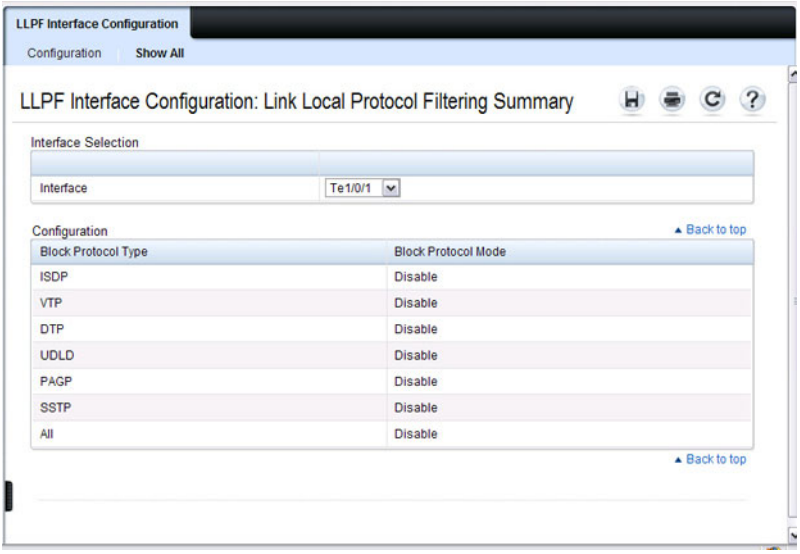
To display the **LLPF Interface Configuration** page, click **Switching → Network Security → Proprietary Protocol Filtering → LLPF Interface Configuration** the navigation menu.

Figure 25-8. LLPF Interface Configuration



To view the protocol types that have been blocked for an interface, click **Show All**.

Figure 25-9. LLPF Filtering Summary



Configuring Port-Based Traffic Control (CLI)

This section provides information about the commands you use to configure port-based traffic control settings. For more information about the commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Configuring Flow Control and Storm Control

Beginning in Privileged EXEC mode, use the following commands to configure the flow control and storm control features.

Command	Purpose
configure	Enter global configuration mode.
flowcontrol	Globally enable flow control.
interface *interface*	Enter interface configuration mode for the specified interface. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**.
	You can also specify a range of interfaces with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12.
storm-control broadcast [**level** *rate*]	Enable broadcast storm recovery mode on the interface and (optionally) set the threshold.
	rate — threshold as percentage of port speed. The percentage is converted to a PacketsPerSecond value based on a 512 byte average packet size.
storm-control multicast [**level** *rate*]	Enable multicast storm recovery mode on the interface and (optionally) set the threshold.
	rate — threshold as percentage of port speed. The percentage is converted to a PacketsPerSecond value based on a 512 byte average packet size.
storm-control unicast [**level** *rate*]	Enable unknown unicast storm recovery mode on the interface and (optionally) set the threshold.
	rate — threshold as percentage of port speed. The percentage is converted to a PacketsPerSecond value based on a 512 byte average packet size.

Command	Purpose
CTRL + Z	Exit to Privileged EXEC mode.
show interfaces detail *interface*	Display detailed information about the specified interface, including the flow control status.
show storm-control	View whether 802.3x flow control is enabled on the switch.
show storm-control [*interface* | **all**]	View storm control settings for all interfaces or the specified interface.

Configuring Protected Ports

Beginning in Privileged EXEC mode, use the following commands to add a name to a protected port group and add ports to the group.

Command	Purpose
configure	Enter global configuration mode.
switchport protected *groupid* **name** *name*	Specify a name for one of the three protected port groups. • *groupid* — Identifies which group the port is to be protected in. (Range: 0-2) • *name* — Name of the group. (Range: 0-32 characters)
interface *interface*	Enter interface configuration mode for the specified interface. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**.
switchport protected *groupid*	Add the interface to the specified protected port group.
CTRL + Z	Exit to Privileged EXEC mode.
show switchport protected	View protected group and port information.

Configuring LLPF

Beginning in Privileged EXEC mode, use the following commands to configure LLPF settings.

Command	Purpose
configure	Enter global configuration mode.
interface *interface*	Enter interface configuration mode for the specified interface. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**. You can also specify a range of interfaces with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12.
service-acl input {blockcdp | blockvtp | blockdtp | blockudld | blockpagp | blocksstp | blockall}	Use the appropriate keyword, or combination of keywords to block any (or all) of the following PDUs on the interface: • VTP • DTP • UDLD • PAgP • SSTP • All
CTRL + Z	Exit to Privileged EXEC mode.
show service-acl interface {*interface* | **all**}	View information about the blocked PDUs on the specified interface or all interfaces.

Port-Based Traffic Control Configuration Example

The commands in this example configure storm control, LLPF, and protected port settings for various interfaces on the switch.

The storm control configuration in this example sets thresholds on the switch so that if broadcast traffic occupies more than 10% on the bandwidth on any physical port, the interface blocks the broadcast traffic until the measured amount of this traffic drops below the threshold.

The LLPF configuration in this example disables all PAgP and VTP PDUs from being forwarded on any switch port or LAG.

The protected port configuration in this example prevents the clients connected to ports 3, 4, and 9 from being able to communicate with each other.

To configure the switch:

1. Configure storm control for broadcast traffic on all physical interfaces.

```
console(config)#interface range te1/0/1-24
console(config-if)#storm-control broadcast
level 10
```

2. Configure LLPF to block PAgP and VTP PDUs on all physical interfaces.

```
console(config-if)#service-acl blockpagp blockvtp
console(config-if)#exit
```

3. Specify a name for protected port group 0.

```
console(config)#protected 0 name clients
```

4. Add the ports to the protected port group.

```
console(config)#interface te1/0/3
console(config-if-Te1/0/3)#switchport protected 0
console(config-if-Te1/0/3)#exit
console(config)#interface te1/0/4
console(config-if-Te1/0/4)#switchport protected 0
console(config-if-Te1/0/4)#exit
console(config)#interface te1/0/9
console(config-if-Te1/0/9)#switchport protected 0
console(config-if-Te1/0/9)#exit
console(config)#exit
```

5 Verify the configuration.

```
console#show storm-control te1/0/1

        Bcast   Bcast   Mcast   Mcast   Ucast   Ucast
 Intf   Mode    Level   Mode    Level   Mode    Level
------ ------- ------- ------- ------- ------- -------
Te1/0/1 Enable    10    Enable    5     Disable   5

console#show service-acl interface te1/0/1

Protocol          Mode
----------------  ----------
CDP               Disabled
VTP               Enabled
DTP               Disabled
UDLD              Disabled
PAGP              Enabled
SSTP              Disabled
ALL               Disabled

console#show switchport protected 0
Name..................................... "clients"


Member Ports: Te1/0/1, Te1/0/2, Te1/0/3, Te1/0/4, Te1/0/9
```

26

Configuring L2 Multicast Features

This chapter describes the layer 2 multicast features on the PowerConnect 8000-series and 8100-series switches. The features this chapter describes include bridge multicast filtering, Internet Group Management Protocol (IGMP) snooping, Multicast Listener Discovery (MLD) snooping, and Multicast VLAN Registration (MVR).

The topics covered in this chapter include:

- L2 Multicast Overview
- Snooping Switch Restrictions
- Current MLD Snooping Functional Limitations Summary
- Default L2 Multicast Values
- Configuring L2 Multicast Features (Web)
- Configuring L2 Multicast Features (CLI)
- Case Study on a Real-World Network Topology

L2 Multicast Overview

Multicast traffic is traffic from one source that has multiple destinations. The L2 multicast features on the switch help control network flooding of Ethernet multicast and IP multicast traffic by keeping track of multicast group membership.

What Are the Multicast Bridging Features?

The PowerConnect 8000-series and 8100-series switches support bridge multicast filtering and bridge multicast forwarding. For Ethernet multicast traffic, the switch uses a database called the Layer 2 Multicast Forwarding Database (MFDB) to make forwarding decisions for packets that arrive with a multicast destination MAC address. By limiting multicasts to only certain ports in the switch, traffic is prevented from going to parts of the network where that traffic is unnecessary.

When a packet enters the switch, the destination MAC address is combined with the VLAN ID, and a search is performed in the Layer 2 MFDB. If no match is found, then the packet is either flooded to all ports in the VLAN or discarded, depending on the switch configuration. If a match is found, then the packet is forwarded only to the ports that are members of that multicast group.

You can create multicast bridging groups and specify the ports and LAGs that are members of each group. This allows L2 multicast traffic to be confined to the specified group.

L2 Multicast Forwarding Modes

You can configure the forwarding mode for each VLAN as one of the following:

- **Forward Unregistered**—Permits the forwarding of multicast frames for multicast groups that have no registrants (no hosts have joined the multicast group). Unregistered multicast streams are flooded on the VLAN. Registered multicast streams are forwarded to ports in the VLAN where hosts are registered.
- **Filter Unregistered**—Prohibits the forwarding of unregistered multicast frames, while continuing to forward registered multicast frames. Unregistered frames are dropped.
- **Forward All**—Floods all registered and unregistered multicast on the VLAN.

What Is IP Multicast Traffic?

IP multicast traffic is traffic that is destined to a host group. Host groups are identified by class D IP addresses, which range from 224.0.0.0 to 239.255.255.255.

When a packet with a broadcast or multicast destination IP address is received, the switch will forward a copy into each of the remaining network segments in accordance with the IEEE MAC Bridge standard. Eventually, the packet is made accessible to all nodes connected to the network.

This approach works well for broadcast packets that are intended to be seen or processed by all connected nodes. In the case of multicast packets, however, this approach could lead to less efficient use of network bandwidth,

particularly when the packet is intended for only a small number of nodes. Packets will be flooded into network segments where no node has any interest in receiving the packet.

What Is IGMP Snooping?

IGMP Snooping is a layer 2 feature that allows the switch to dynamically add or remove ports from IP multicast groups by listening to IGMP join and leave requests. By "snooping" the IGMP packets transmitted between hosts and routers, the IGMP Snooping feature enables the switch to forward IP multicast traffic more intelligently and help conserve bandwidth.

Based on the IGMP query and report messages, the switch forwards traffic only to the ports that request the multicast traffic. This prevents the switch from broadcasting the traffic to all ports and possibly affecting network performance. The switch uses the information in the IGMP packets as they are being forwarded throughout the network to determine which segments should receive packets directed to the group address.

IGMP Snooping Querier

When PIM and IGMP are enabled in a network with IP multicast routing, the IP multicast router acts as the IGMP querier. However, if the IP-multicast traffic in a VLAN needs to be Layer 2 switched only, an IP-multicast router is not required. The IGMP Snooping Querier can perform the IGMP snooping functions on the VLAN.

NOTE: Without an IP-multicast router on a VLAN, you must configure another switch as the IGMP querier so that it can send queries.

When the IGMP snooping querier is enabled, the IGMP snooping querier sends out periodic IGMP queries that trigger IGMP report messages from the switch that wants to receive IP multicast traffic. The IGMP snooping feature listens to these IGMP reports to establish appropriate forwarding.

What Is MLD Snooping?

In IPv4, Layer 2 switches can use IGMP snooping to limit the flooding of multicast traffic by dynamically configuring L2 interfaces so that multicast traffic is forwarded to only those interfaces associated with an IP multicast address. In IPv6, MLD snooping performs a similar function. With MLD

snooping, IPv6 multicast data is selectively forwarded to a list of ports that want to receive the data instead of being flooded to all ports in a VLAN. This list is constructed by snooping IPv6 multicast control packets.

MLD is a protocol used by IPv6 multicast routers to discover the presence of multicast listeners (nodes wishing to receive IPv6 multicast packets) on its directly-attached links and to discover which multicast packets are of interest to neighboring nodes. MLD is derived from IGMP; MLD version 1 (MLDv1) is equivalent to IGMPv2, and MLD version 2 (MLDv2) is equivalent to IGMPv3. MLD is a subprotocol of Internet Control Message Protocol version 6 (ICMPv6), and MLD messages are a subset of ICMPv6 messages.

The switch can snoop on both MLDv1 and MLDv2 protocol packets and bridge IPv6 multicast data based on destination IPv6 multicast MAC addresses. The switch can be configured to perform MLD snooping and IGMP snooping simultaneously.

What Is Multicast VLAN Registration?

IGMP snooping helps limit multicast traffic when member ports are in the same VLAN; however, when ports belong to different VLANs, a copy of the multicast stream is sent to each VLAN that has member ports in the multicast group. MVR eliminates the need to duplicate the multicast traffic when multicast group member ports belong to different VLANs.

MVR uses a dedicated multicast VLAN to forward multicast traffic over the L2 network. Only one MVLAN can be configured per switch, and it is used only for certain multicast traffic, such as traffic from an IPTV application, to avoid duplication of multicast streams for clients in different VLANs. Clients can dynamically join or leave the mutlicast VLAN without interfering with their membership in other VLANs.

MVR, like IGMP Snooping, allows a layer 2 switch to listen to IGMP messages to learn about multicast group membership.

There are two types of MVR ports: source and receiver.

- Source port is the port where multicast traffic is flowing to. It has to be the member of so called multicast VLAN.
- Receiver port is the port where listening host is connected to the switch. It can be the member of any VLAN, except multicast VLAN.

There are two configured learning modes of the MVR operation: dynamic and compatible.

- In the dynamic mode MVR learns existent multicast groups by parsing the IGMP queries from router on source ports and forwarding the IGMP joins from the hosts to the router.
- In the compatible mode MVR does not learn multicast groups, but they have to be configured by administrator and protocol does not forward joins from the hosts to the router. To work in this mode the IGMP router has to be configured to transmit required multicast streams to the network with the MVR switch.

Enabling MVR and IGMP on the Same Interface

MVR and IGMP snooping operate independently and could both be enabled on an interface. When both MVR and IGMP snooping are enabled, MVR listens to the IGMP join and report messages for static multicast group information, and IGMP snooping manages dynamic multicast groups.

When Are L3 Multicast Features Required?

In addition to L2 multicast features, the switch suports IP and IPv6 multicast features. You configure the IP/IPv6 multicast features if the switch functions as a multicast router that can route multicast traffic between VLAN routing interfaces. In this case, you must enable a multicast routing protocol on the switch, such as PIM-SM. For information about L3 multicast features, see "Managing IPv4 and IPv6 Multicast" on page 1157.

If you enable IGMP Snooping on the switch to listen to IGMP traffic, you do not need to enable IGMP, a layer 3 multicast protocol. If the switch functions as a multicast router, it is possible to enable both IGMP and IGMP Snooping so that the switch routes IGMP traffic between VLANs and examines the IGMP packets for join and leave information.

NOTE: If MVR is enabled, IP Multicast should be disabled. Multicast routing and MVR cannot coexist on a switch.

For information about configuring a PowerConnect 8000-series and 8100-series switches as a mutlicast router that also performs IGMP snooping, see "Configuring Multicast VLAN Routing With IGMP and PIM-SM" on page 1231.

What Are GARP and GMRP?

Generic Attribute Registration Protocol (GARP) is a general-purpose protocol that registers any network connectivity or membership-style information. GARP defines a set of switches interested in a given network attribute, such as VLAN ID or multicast address.

PowerConnect 8000-series and 8100-series switches can use GARP functionality for two applications:

- GARP VLAN Registration Protocol (GVRP) to help dynamically manage VLAN memberships on trunk ports.
- GARP Multicast Registration Protocol (GMRP) to help control the flooding of multicast traffic by keeping track of group membership information.

GVRP and GMRP use the same set of GARP Timers to specify the amount of time to wait before transmitting various GARP messages.

GMRP is similar to IGMP snooping in its purpose, but IGMP snooping is more widely used. GMRP must be running on both the host and the switch to function properly.

Snooping Switch Restrictions

Partial IGMPv3 and MLDv2 Support

The IGMPv3 and MLDv2 protocols allow multicast listeners to specify the list of hosts from which they want to receive the traffic. However the PowerConnect snooping switch does not track this information. IGMPv3/MLDv2 Report messages that have the group record type CHANGE_TO_INCLUDE_MODE with a null source list are treated as Leave messages. All other report messages are treated as IGMPv2/MLDv1 Report messages.

MAC Address-Based Multicast Group

The L2 multicast forwarding table consists of the IP Multicast group MAC address. For IPv4 multicast groups, 16 IP multicast group addresses map to the same multicast MAC address. For example, 224.1.1.1 and 225.1.1.1 map to the MAC address 01:00:5E:01:01:01, and IP addresses in the range [224-239].3.3.3 map to 01:00:5E:03:03:03. As a result, if a host requests 225.1.1.1, then it might receive multicast traffic of group 226.1.1.1 as well.

IGMP Snooping in a Multicast Router

IGMP snooping is a Layer 2 feature and is achieved by using the L2 multicast forwarding table. However when multicast routing is enabled on a PowerConnect switch, L2 multicast forwarding entries do not affect multicast data forwarding. Instead, corresponding IP multicast table entries need to be created to achieve similar behavior.

On a multicast router, for IGMP snooping to be functional, any multicast routing protocol needs to be operationally enabled on the routing interface. IGMP snooping also needs to be enabled on the VLAN corresponding to the routing interface. Note that IGMP snooping behavior will not be functional on VLANs that are not enabled for VLAN routing.

Topologies Where the Multicast Source Is Not Directly Connected to the Querier

If the multicast source is not directly connected to a multicast querier, it is likely that the multicast stream is dropped at intermediate snooping switches. This causes hosts in certain network segments not to receive the multicast

stream even when a Report message is successfully forwarded up to the querier. This problem is seen when multicast data forwarding behavior on snooping switches is configured to forward registered multicast groups and drop unregistered multicast groups.

To overcome this problem, the administrator should configure static multicast forwarding entries on intermediate snooping switches to forward the multicast stream to the multicast querier.

Alternatively, a better option is to reconfigure the topology to make sure multicast streams always reach the multicast querier without getting dropped by intermediate network elements. The Snooping on a Multicast Router section describes this problem.

Current MLD Snooping Functional Limitations Summary

This section outlines known functional limitations regarding MLD snooping in its ability to support multicast in a L2/L3 mixed topology. This section is intended to be read after the reader has acquired a thorough understanding of the connection setup/teardown steps that snooping undergoes, as outlined in the examples presented in the "Multicast Snooping Case Study" on page 736.

The following are known possible issues, depending upon the total number of unique multicast groups, their sources, and their data rates:

1. As described in "Topologies Where the Multicast Source Is Not Directly Connected to the Querier" on page 687, when operating in flood mode (see "L2 Multicast Forwarding Modes" on page 682), multicast streams towards a multicast (that is, PIM) router will be blocked by the multicast server's local layer 2 switch, inhibiting the router from discovering the stream.

Default L2 Multicast Values

All L2 multicast features are disabled by default. Details about the L2 multicast are in Table 26-1.

Table 26-1. L2 Multicast Defaults

Parameter	Default Value
Bridge Multicast Filtering	Disabled
IGMP Snooping mode	Disabled
MLD Snooping mode	Disabled
Bridge multicast group	None configured
Bridge multicast forwarding mode	Forward unregistered
IGMP/MLD snooping	Disabled on all interfaces
IGMP/MLD snooping auto-learn	Disabled
IGMP/MLD snooping host timeout	260 seconds
IGMP/MLD snooping multicast router timeout	300 seconds
IGMP/MLD snooping leave timeout	10 seconds
IGMP snooping querier	Disabled
IGMP version	v2
MLD version	v1
IGMP/MLD snooping querier query interval	60 seconds
IGMP/MLD snooping querier expiry interval	60 seconds
IGMP/MLD snooping VLAN querier	Disabled
VLAN querier election participate mode	Disabled
Snooping Querier VLAN Address	0.0.0.0
MVR running	Disabled
MVR multicast VLAN	1
MVR max multicast groups	256
MVR Global query response time	5 tenths of a second
MVR Mode	Compatible

Table 26-1. L2 Multicast Defaults (Continued)

Parameter	Default Value
GARP Leave Timer	60 centiseconds
GARP Leave All Timer	1000 centiseconds
GARP Join Timer	20 centiseconds
GMRP	Disabled globally and per-interface

Configuring L2 Multicast Features (Web)

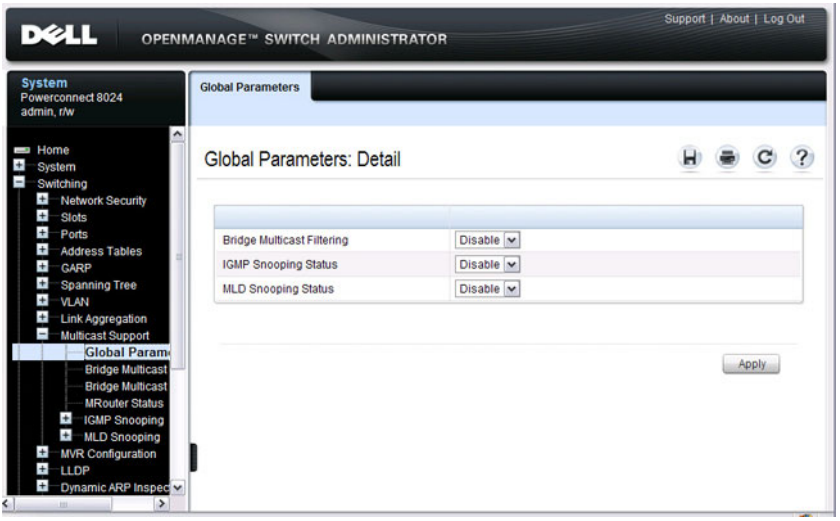
This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring L2 multicast features on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

Multicast Global Parameters

Use the **Multicast Global Parameters** page to enable or disable bridge multicast filtering, IGMP Snooping, or MLD Snooping on the switch.

To display the **Multicast Global Parameters** page, click **Switching** → **Multicast Support** → **Global Parameters** in the navigation menu.

Figure 26-1. Multicast Global Parameters

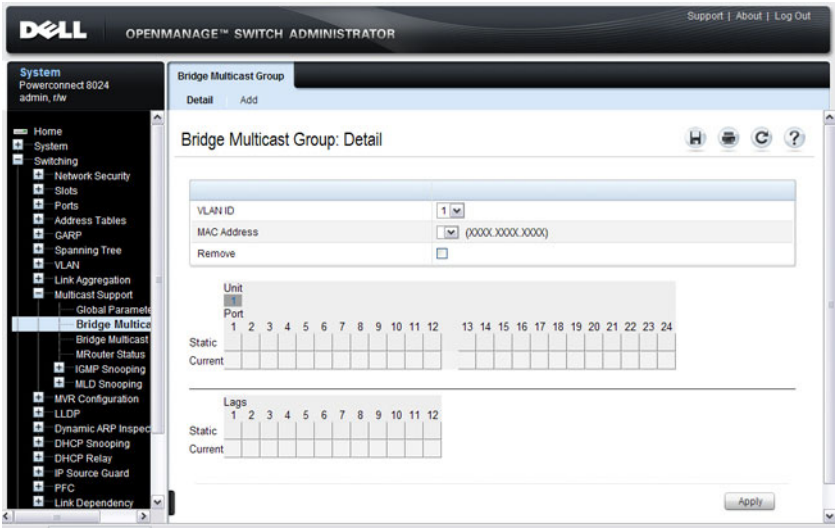


Bridge Multicast Group

Use the **Bridge Multicast Group** page to create new multicast service groups or to modify ports and LAGs assigned to existing multicast service groups. Attached interfaces display in the Port and LAG tables and reflect the manner in which each is joined to the Multicast group.

To display the **Bridge Multicast Group** page, click **Switching** → **Multicast Support** → **Bridge Multicast Group** in the navigation menu.

Figure 26-2. Bridge Multicast Group



Understanding the Port and LAG Member Tables

The **Bridge Multicast Group** tables display which Ports and LAGs are members of the multicast group, and whether they're static (S), dynamic (D), or forbidden (F). The tables have two rows: **Static** and **Current**. Only the **Static** row is accessible from this page. The **Current** row is updated when the **Static** row is changed and **Apply** is clicked.

The **Bridge Multicast Group** page contains two editable tables:

- **Unit and Ports** — Displays and assigns multicast group membership to ports. To assign membership, click in **Static** for a specific port. Each click toggles between S, F, and blank. See Table 26-2 for definitions.

- **LAGs** — Displays and assigns multicast group membership to LAGs. To assign membership, click in **Static** for a specific LAG. Each click toggles between S, F, and blank. See Table 26-2 for definitions.

Table 26-2 contains definitions for port/LAG IGMP management settings.

Table 26-2. Port/LAG IGMP Management Settings

Port Control	Definition
D	Dynamic: Indicates that the port/LAG was dynamically joined to the Multicast group (displays in the *Current* row).
S	Static: Attaches the port to the Multicast group as a static member in the *Static* row. Displays in the *Current* row once **Apply** is clicked.
F	Forbidden: Indicates that the port/LAG is forbidden entry into the Multicast group in the *Static* row. Displays in the *Current* row once **Apply** is clicked.
Blank	Blank: Indicates that the port is not attached to a Multicast group.

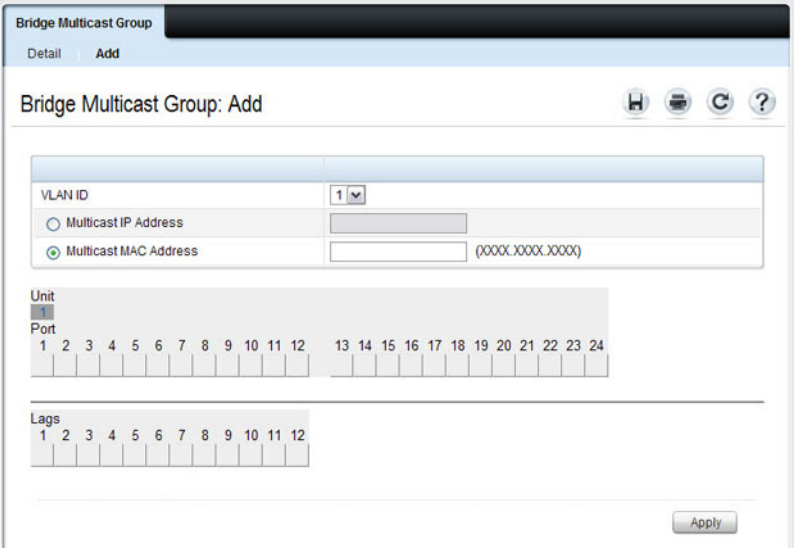
Adding and Configuring Bridge Multicast Address Groups

To configure a bridge multicast group:

1. From the **Bridge Multicast Group** page, click **Add**.

   The **Add Bridge Multicast Group** page displays.

Figure 26-3. Add Bridge Multicast Group



2 Select the ID of the VLAN to add to the multicast group or to modify membership for an existing group.

3 For a new group, specify the multicast group IP or MAC address associated with the selected VLAN.

4 In the **Bridge Multicast Group** tables, assign a setting by clicking in the **Static** row for a specific port/LAG. Each click toggles between S, F, and blank. (not a member).

5 Click **Apply**.

The bridge multicast address is assigned to the multicast group, ports/LAGs are assigned to the group (with the **Current** rows being updated with the **Static** settings), and the switch is updated.

Removing a Bridge Multicast Group

To delete a bridge multicast group:

1. Open the **Bridge Multicast Group** page.
2. Select the **VLAN ID** associated with the bridge multicast group to be removed from the drop-down menu.

   The **Bridge Multicast Address** and the assigned ports/LAGs display.
3. Check the **Remove** check box.
4. Click **Apply**.

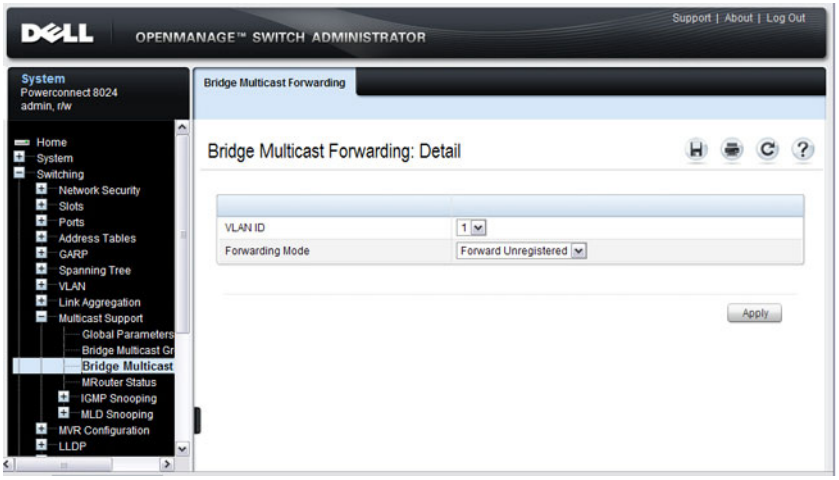
   The selected bridge multicast group is removed, and the device is updated.

Bridge Multicast Forwarding

Use the **Bridge Multicast Forwarding** page to enable attaching ports or LAGs to a switch that is attached to a neighboring Multicast switch. Once IGMP Snooping is enabled, multicast packets are forwarded to the appropriate port or VLAN.

To display the **Bridge Multicast Forwarding** page, click **Switching** → **Multicast Support** → **Bridge Multicast Forwarding** in the navigation menu.

Figure 26-4. Bridge Multicast Forwarding

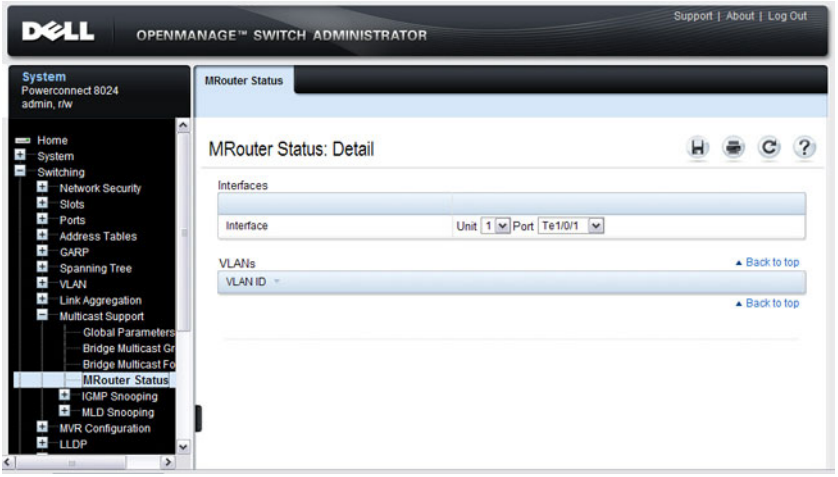


MRouter Status

Use the **MRouter Status** page to display the status of dynamically learned multicast router interfaces.

To access this page, click **Switching** → **Multicast Support** → **MRouter Status** in the navigation panel.

Figure 26-5. MRouter Status

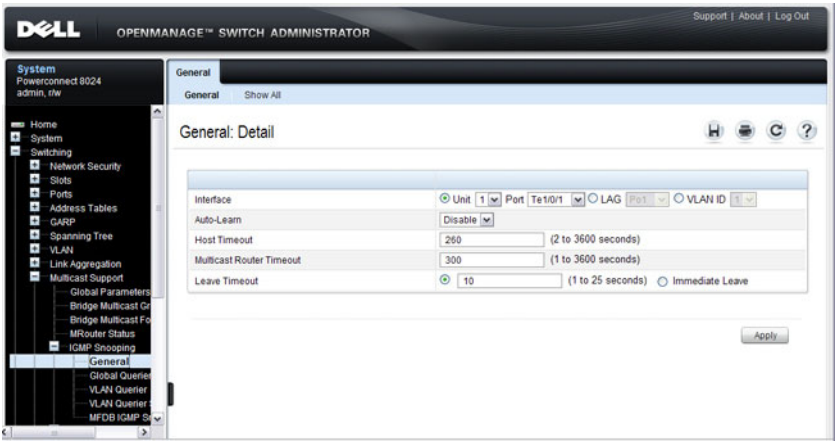


General IGMP Snooping

Use the **General** IGMP snooping page to configure IGMP snooping settings on specific ports and LAGs.

To display the **General** IGMP snooping page, click **Switching** → **Multicast Support** → **IGMP Snooping** → **General** in the navigation menu.

Figure 26-6. General IGMP Snooping



Modifying IGMP Snooping Settings for Multiple Ports, LAGs, or VLANs

To modify the IGMP snooping settings:

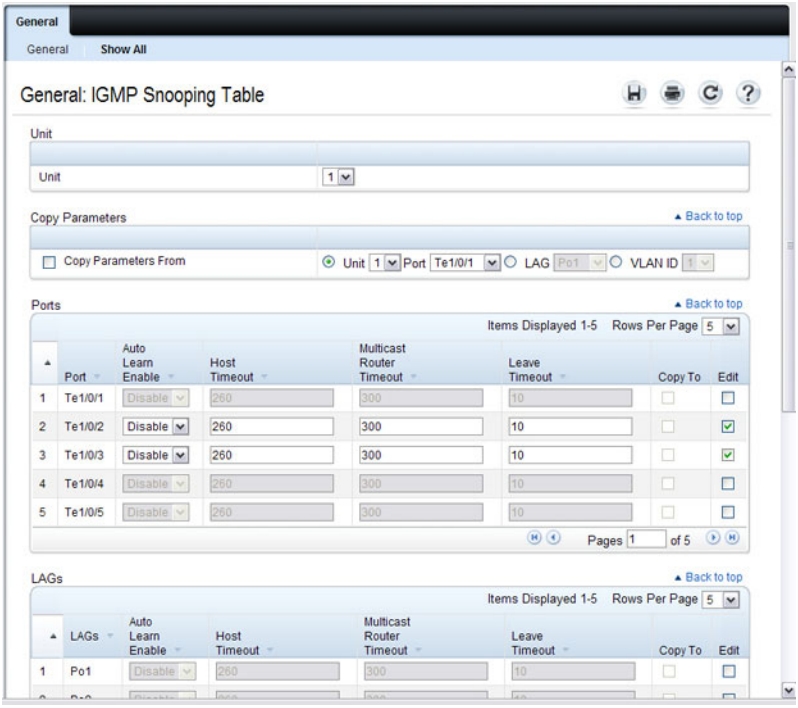
1. From the **General** IGMP snooping page, click **Show All**.

   The **IGMP Snooping Table** displays.

2. Select the **Edit** checkbox for each Port, LAG, or VLAN to modify.

   In Figure 26-7, ports 2 and 3 are to be modified.

Figure 26-7. Edit IGMP Snooping Settings



3 Edit the IGMP Snooping fields as needed.

4 Click **Apply**.

The IGMP Snooping settings are modified, and the device is updated.

Copying IGMP Snooping Settings to Multiple Ports, LAGs, or VLANs

To copy IGMP snooping settings:

1 From the **General** IGMP snooping page, click **Show All**.

The **IGMP Snooping Table** displays.

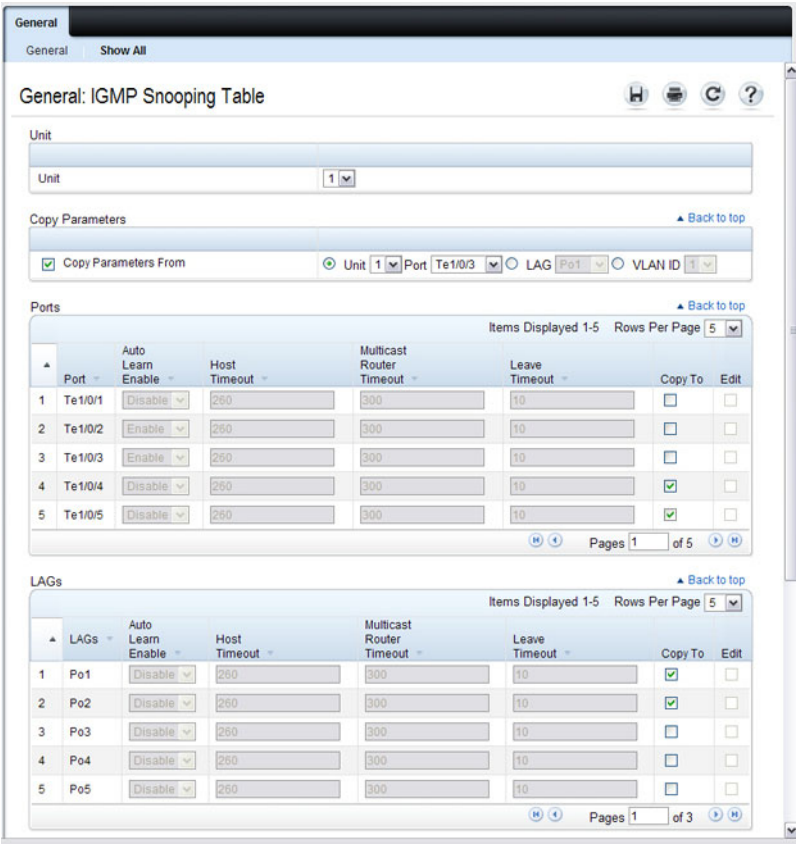
2 Select the **Copy Parameters From** checkbox.

3 Select a Unit/Port, LAG, or VLAN to use as the source of the desired parameters.

4 Select the **Copy To** checkbox for the Unit/Ports, LAGs, or VLANs that these parameters will be copied to.

In Figure 26-8, the settings for port 3 will be copied to ports 4 and 5 and LAGs 1 and 2.

Figure 26-8. Copy IGMP Snooping Settings



5 Click **Apply**.

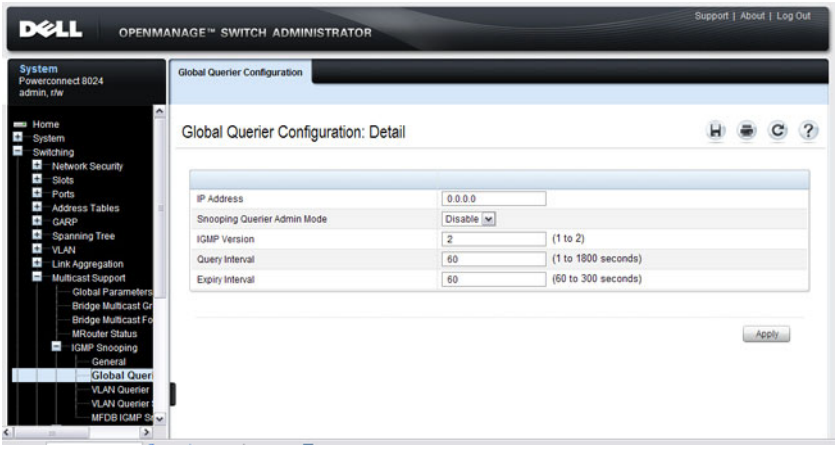
The IGMP Snooping settings are modified, and the device is updated.

Global Querier Configuration

Use the **Global Querier Configuration** page to configure IGMP snooping querier settings, such as the IP address to use as the source in periodic IGMP queries when no source address has been configured on the VLAN.

To display the **Global Querier Configuration** page, click **Switching** → **Multicast Support** → **IGMP Snooping** → **Global Querier Configuration** in the navigation menu.

Figure 26-9. Global Querier Configuration

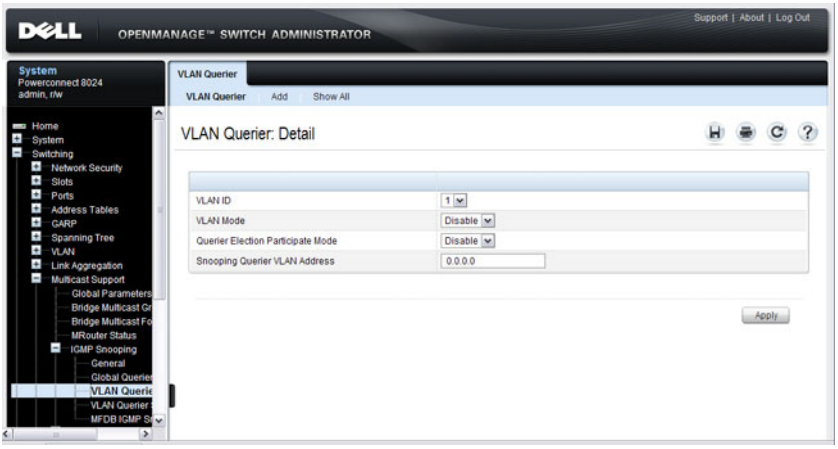


VLAN Querier

Use the **VLAN Querier** page to specify the IGMP Snooping Querier settings for individual VLANs.

To display the **VLAN Querier** page, click **Switching** → **Multicast Support** → **IGMP Snooping** → **VLAN Querier** in the navigation menu.

Figure 26-10. VLAN Querier



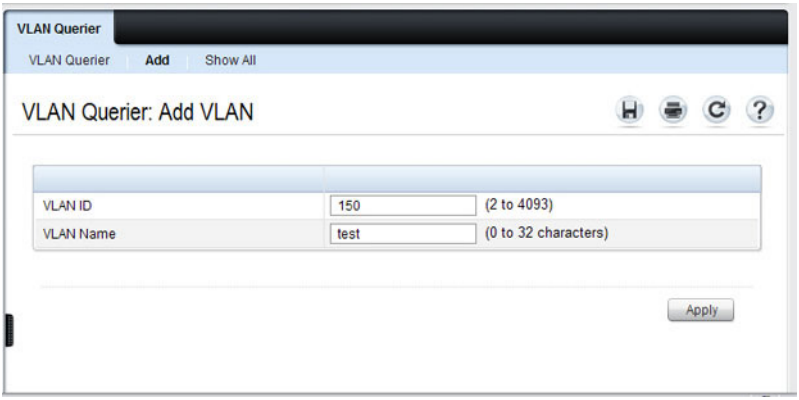
Adding a New VLAN and Configuring its VLAN Querier Settings

To configure a VLAN querier:

1. From the **VLAN Querier** page, click **Add**.

   The page refreshes, and the **Add VLAN** page displays.

Figure 26-11. Add VLAN Querier



2 Enter the VLAN ID and, if desired, an optional VLAN name.

3 Return to the **VLAN Querier** page and select the new VLAN from the **VLAN ID** menu.

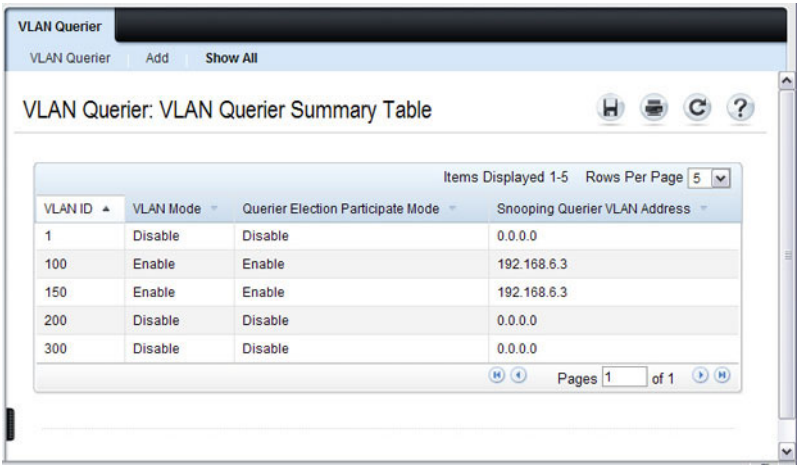
4 Specify the VLAN querier settings.

5 Click **Apply**.

The VLAN Querier settings are modified, and the device is updated.

To view a summary of the IGMP snooping VLAN querier settings for all VLANs on the switch, click **Show All**.

Figure 26-12. Add VLAN Querier

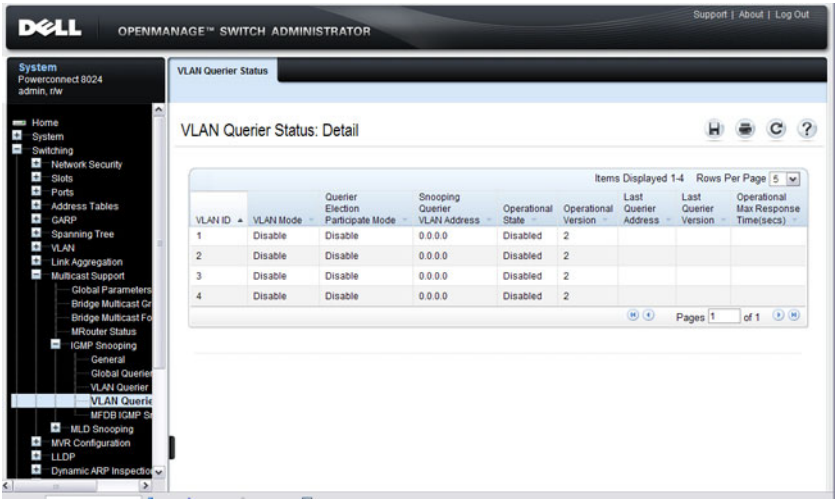


VLAN Querier Status

Use the **VLAN Querier Status** page to view the IGMP Snooping Querier settings for individual VLANs.

To display the **VLAN Querier Status** page, click **Switching → Multicast Support → IGMP Snooping → VLAN Querier Status** in the navigation menu.

Figure 26-13. IGMP Snooping VLAN Querier Status

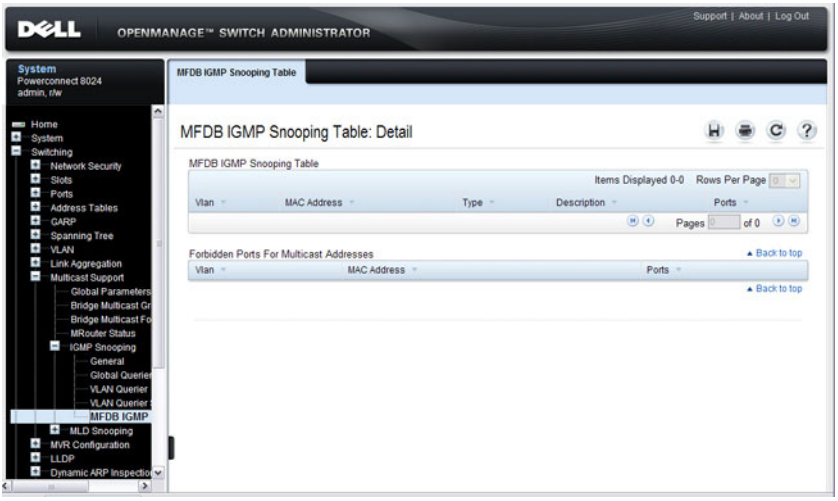


MFDB IGMP Snooping Table

Use the **MFDB IGMP Snooping Table** page to view the multicast forwarding database (MFDB) IGMP Snooping Table and Forbidden Ports settings for individual VLANs.

To display the **MFDB IGMP Snooping Table** page, click **Switching** → **Multicast Support** → **IGMP Snooping** → **MFDB IGMP Snooping Table** in the navigation menu.

Figure 26-14. MFDB IGMP Snooping Table

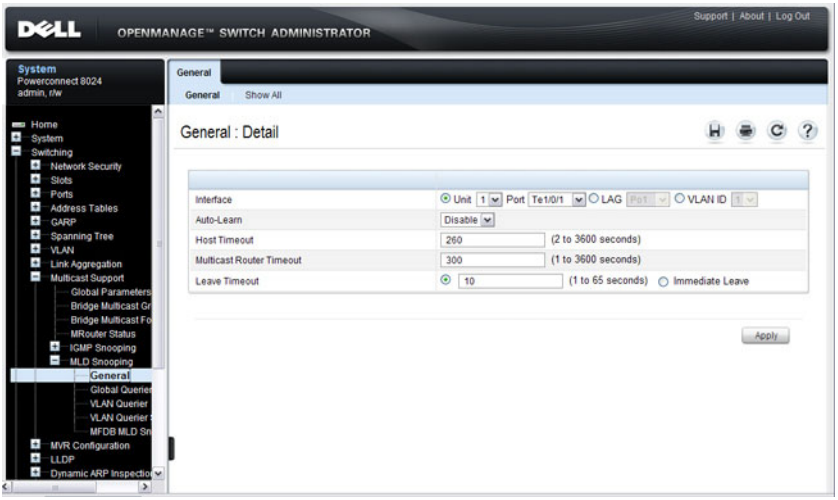


MLD Snooping General

Use the MLD Snooping **General** page to add MLD members.

To access this page, click **Switching** → **Multicast Support** → **MLD Snooping** → **General** in the navigation panel.

Figure 26-15. MLD Snooping General



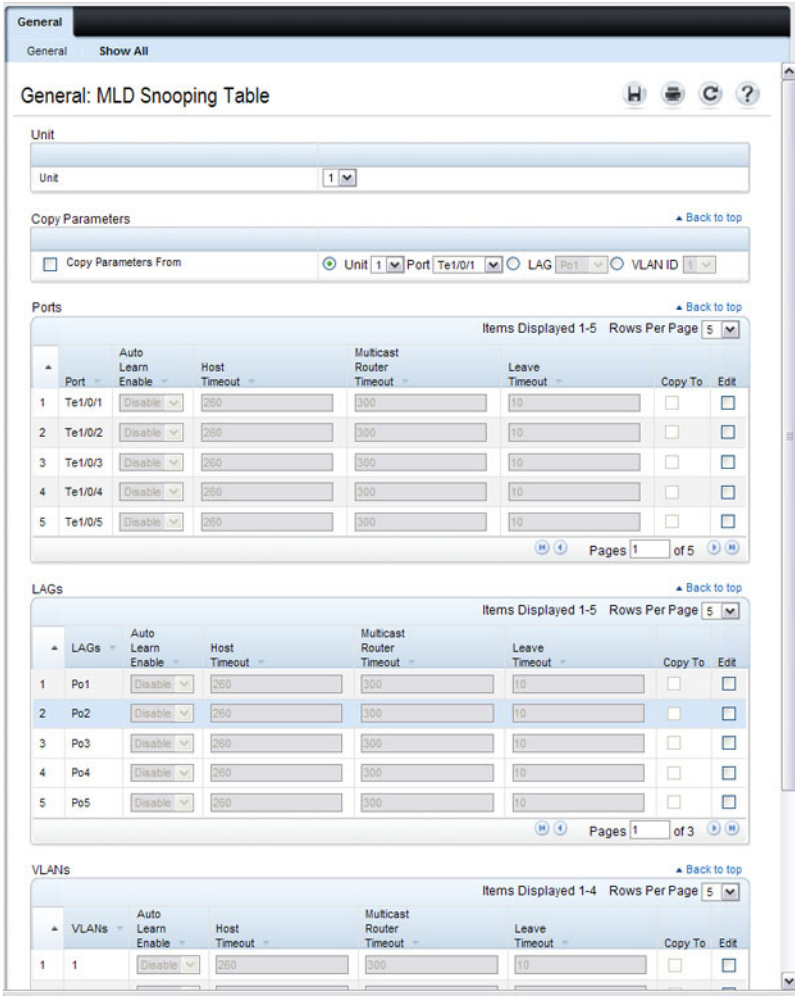
Modifying MLD Snooping Settings for Multiple Ports, LAGs, or VLANs

To configure MLD snooping:

1. From the **General** MLD snooping page, click **Show All**.

   The **MLD Snooping Table** displays.

Figure 26-16. MLD Snooping Table



2 Select the **Edit** checkbox for each Port, LAG, or VLAN to modify.

3 Edit the MLD Snooping fields as needed.

4 Click **Apply**.

The MLD Snooping settings are modified, and the device is updated.

Copying MLD Snooping Settings to Multiple Ports, LAGs, or VLANs

To copy MLD snooping settings:

1. From the **General** MLD snooping page, click **Show All**.

   The **MLD Snooping Table** displays.
2. Select the **Copy Parameters From** checkbox.
3. Select a Unit/Port, LAG, or VLAN to use as the source of the desired parameters.
4. Select the **Copy To** checkbox for the Unit/Ports, LAGs, or VLANs that these parameters will be copied to.
5. Click **Apply**.

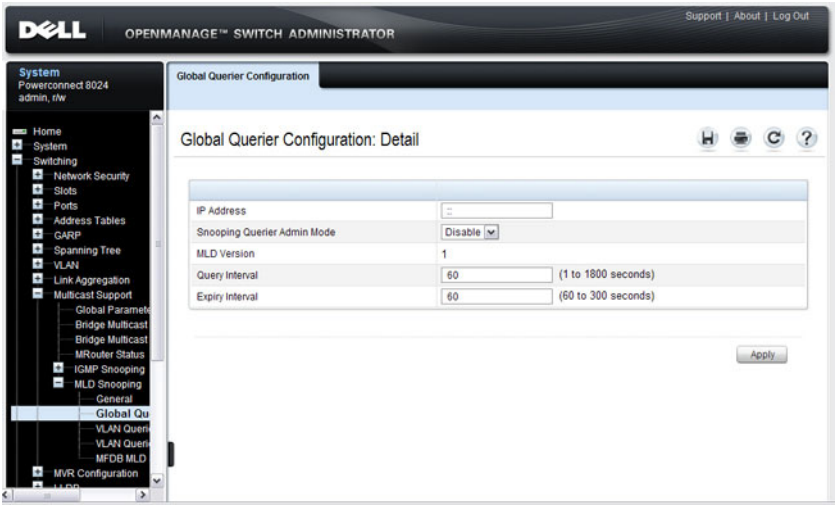
   The MLD Snooping settings are modified, and the device is updated.

MLD Snooping Global Querier Configuration

Use the MLD Snooping **Global Querier Configuration** page to configure the parameters for the MLD Snooping Querier.

To display the **Global Querier Configuration** page, click **Switching** → **Multicast Support** → **MLD Snooping** → **Global Querier Configuration** in the navigation menu.

Figure 26-17. MLD Snooping Global Querier Configuration

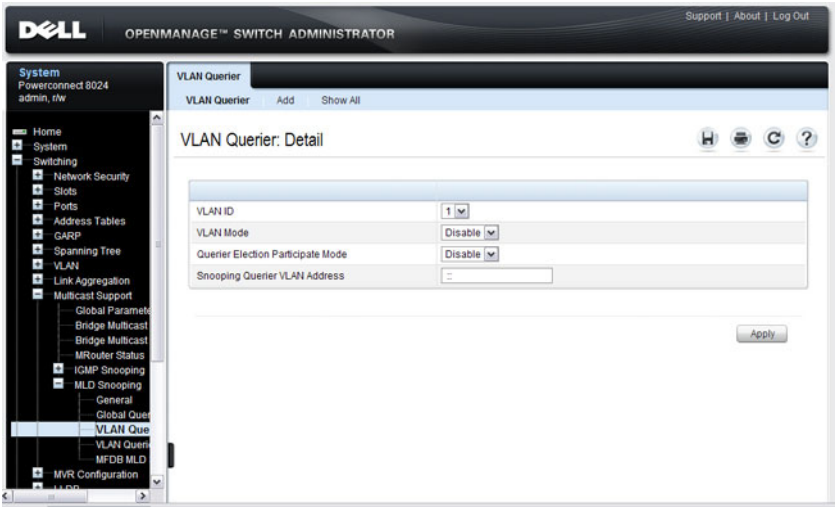


MLD Snooping VLAN Querier

Use the MLD Snooping **VLAN Querier** page to specify the MLD Snooping Querier settings for individual VLANs.

To display the MLD Snooping **VLAN Querier** page, click **Switching** → **Multicast Support** → **MLD Snooping** → **VLAN Querier** in the navigation menu.

Figure 26-18. MLD Snooping VLAN Querier



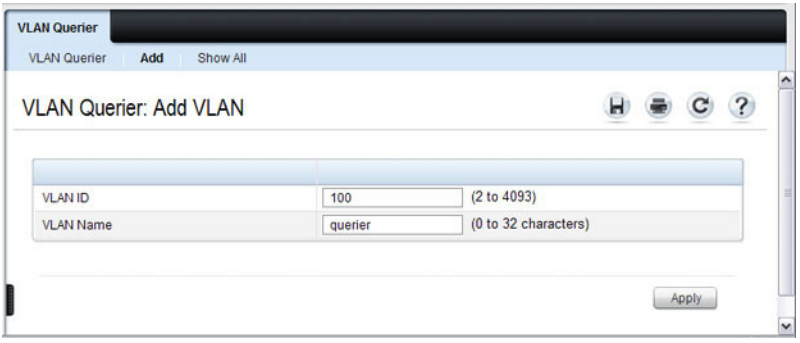
Adding a New VLAN and Configuring its MLD Snooping VLAN Querier Settings

To configure an MLD snooping VLAN querier:

1. From the **VLAN Querier** page, click **Add**.

   The page refreshes, and the **Add VLAN** page displays.

Figure 26-19. Add MLD Snooping VLAN Querier



2. Enter the VLAN ID and, if desired, an optional VLAN name.

3 Return to the **VLAN Querier** page and select the new VLAN from the **VLAN ID** menu.

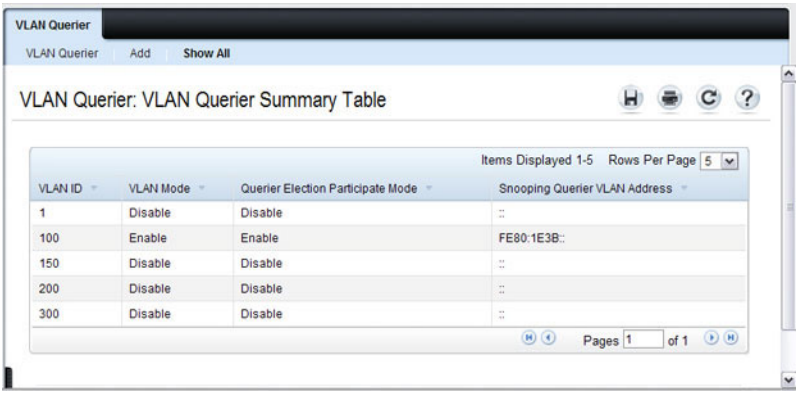
4 Specify the VLAN querier settings.

5 Click **Apply**.

The VLAN Querier settings are modified, and the device is updated.

To view a summary of the IGMP snooping VLAN querier settings for all VLANs on the switch, click **Show All**.

Figure 26-20. Add VLAN Querier

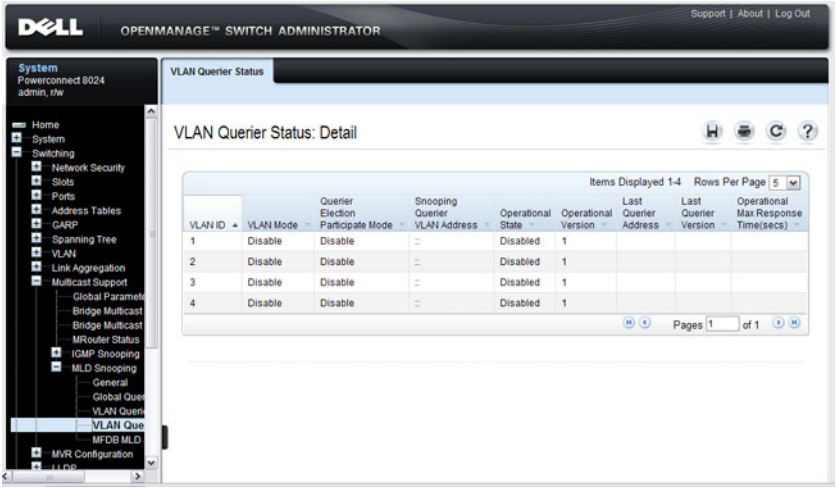


MLD Snooping VLAN Querier Status

Use the **VLAN Querier Status** page to view the MLD Snooping Querier settings for individual VLANs.

To display the **VLAN Querier Status** page, click **Switching → Multicast Support → MLD Snooping → VLAN Querier Status** in the navigation menu.

Figure 26-21. MLD Snooping VLAN Querier Status

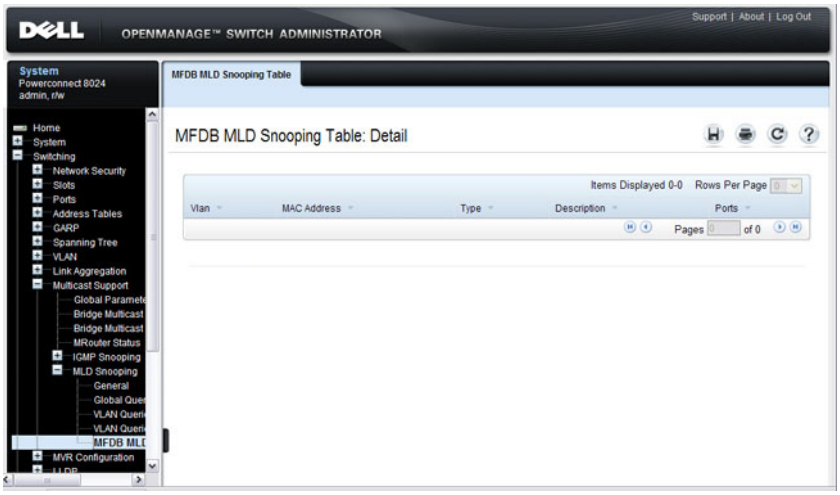


MFDB MLD Snooping Table

Use the **MFDB MLD Snooping Table** page to view the MFDB MLD Snooping Table settings for individual VLANs.

To display the **MFDB MLD Snooping Table** page, click **Switching** → **Multicast Support** → **MLD Snooping** → **MFDB MLD Snooping Table** in the navigation menu.

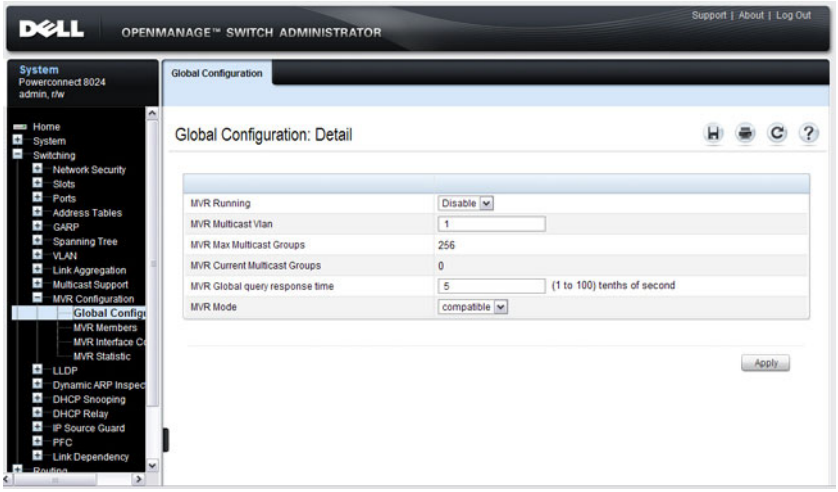
Figure 26-22. MFDB MLD Snooping Table



MVR Global Configuration

Use the **MVR Global Configuration** page to enable the MVR feature and configure global parameters. To display the **MVR Global Configuration** page, click **Switching → MVR Configuration → Global Configuration** in the navigation panel.

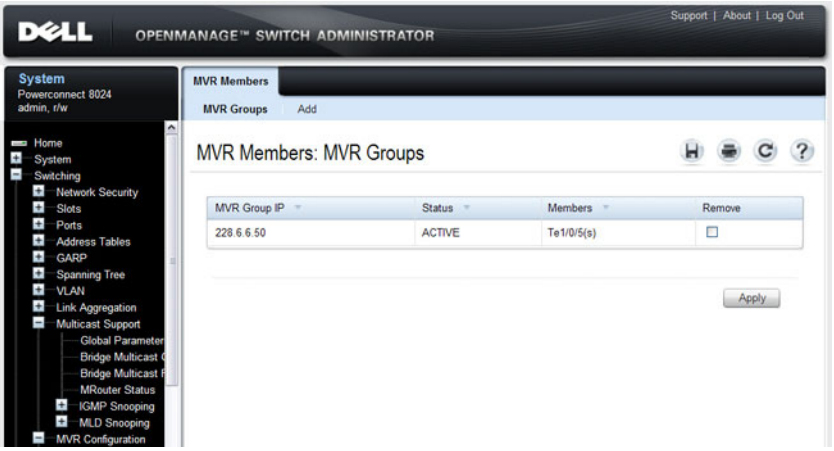
Figure 26-23. MVR Global Configuration



MVR Members

Use the MVR Members page to view and configure MVR group members. To display the MVR Members page, click **Switching** → **MVR Configuration** → **MVR Members** in the navigation panel.

Figure 26-24. MVR Members



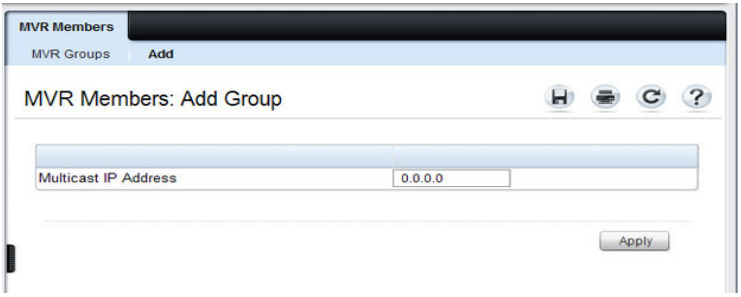
Adding an MVR Membership Group

To add an MVR membership group:

1. From the **MVR Membership** page, click **Add**.

   The MVR **Add Group** page displays.

Figure 26-25. MVR Member Group



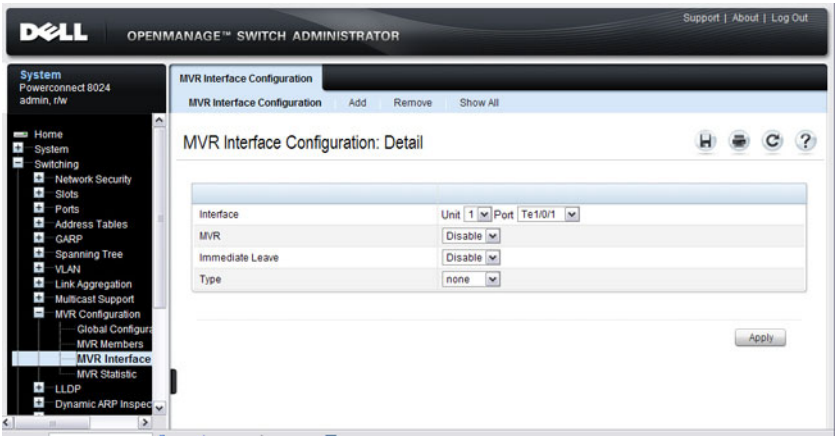
2 Specify the MVR group IP multicast address.

3 Click **Apply**.

MVR Interface Configuration

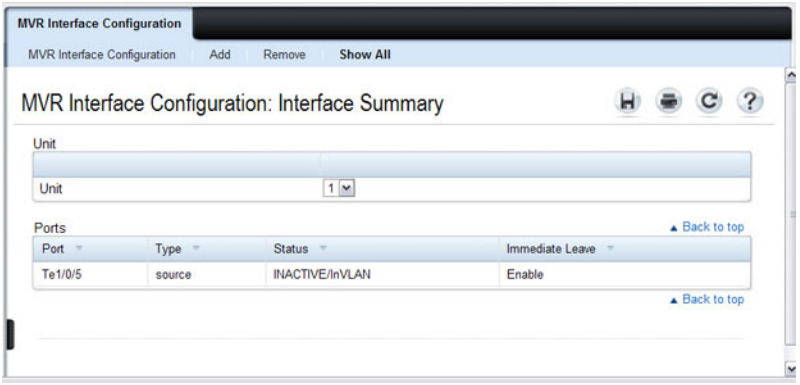
Use the MVR Interface Configuration page to enable MVR on a port, configure its MVR settings, and add the port to an MVR group. To display the MVR Interface Configuration page, click **Switching** → **MVR Configuration** → **MVR Interface Configuration** in the navigation panel.

Figure 26-26. MVR Interface Configuration



To view a summary of the MVR interface configuration, click **Show All**.

Figure 26-27. MVR Interface Summary

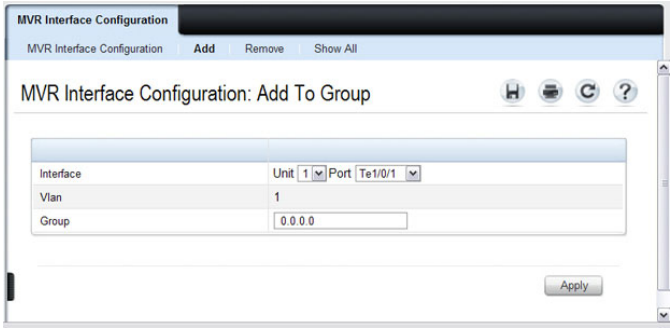


Adding an Interface to an MVR Group

To add an interface to an MVR group:

1. From the **MVR Interface** page, click **Add**.

Figure 26-28. MVR - Add to Group



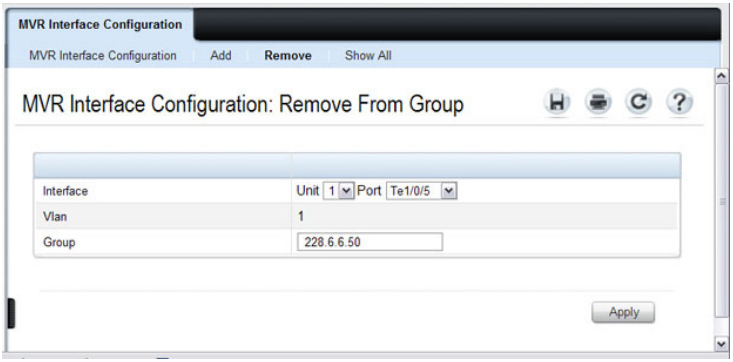
2. Select the interface to add to the MVR group.
3. Specify the MVR group IP multicast address.
4. Click **Apply**.

Removing an Interface from an MVR Group

To remove an interface from an MVR group:

1 From the **MVR Interface** page, click **Remove**.

Figure 26-29. MVR - Remove from Group



2 Select the interface to remove from an MVR group.

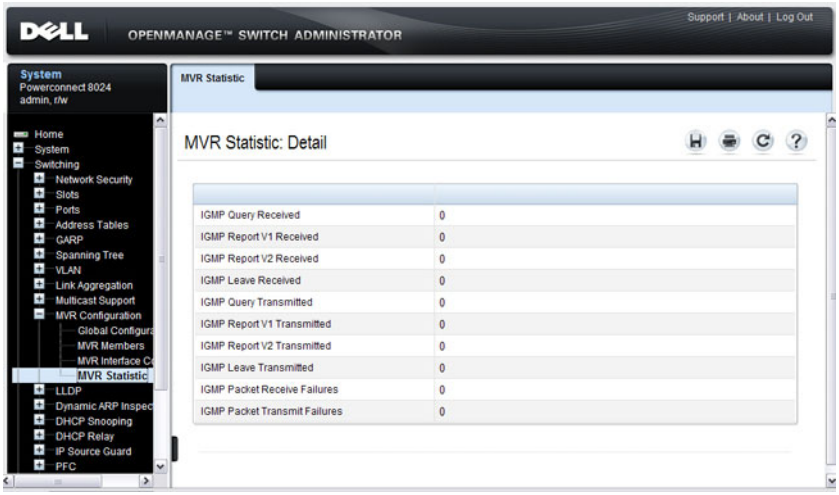
3 Specify the IP multicast address of the MVR group.

4 Click **Apply**.

MVR Statistics

Use the **MVR Statistics** page to view MVR statistics on the switch. To display the **MVR Statistics** page, click **Switching → MVR Configuration → MVR Statistics** in the navigation panel.

Figure 26-30. MVR Statistics

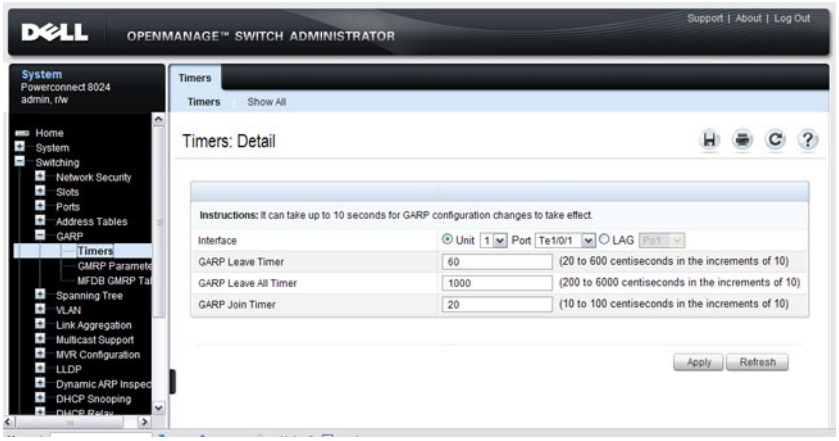


GARP Timers

The **Timers** page contains fields for setting the GARP timers used by GVRP and GMRP on the switch.

To display the **Timers** page, click **Switching** → **GARP** → **Timers** in the navigation panel.

Figure 26-31. GARP Timers

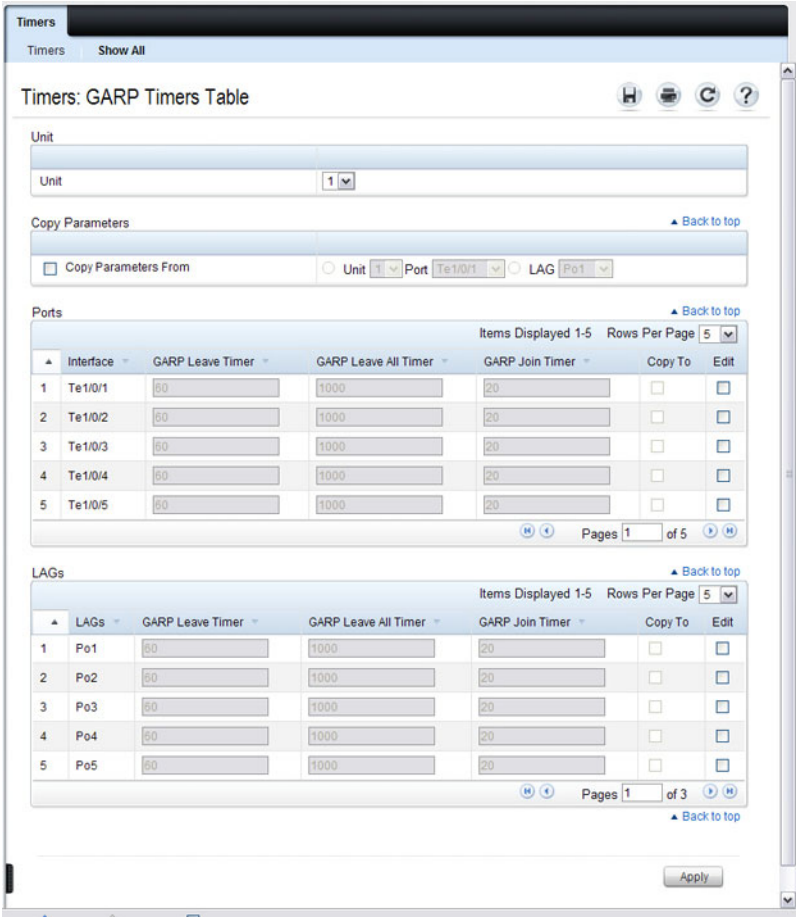


Configuring GARP Timer Settings for Multiple Ports

To configure GARP timers on multiple ports:

1. Open the **Timers** page.
2. Click **Show All** to display the **GARP Timers Table**.

Figure 26-32. Configure STP Port Settings



3 For each port or LAG to configure, select the check box in the **Edit** column in the row associated with the port.

4 Specify the desired timer values.

5 Click **Apply**.

Copying GARP Timer Settings From One Port to Others

To copy GARP timer settings:

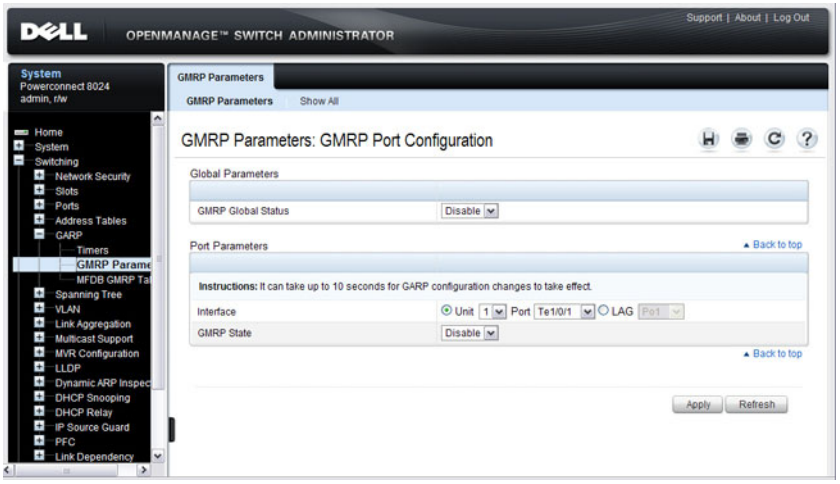
1. Select the **Copy Parameters From** check box, and select the port or LAG with the settings to apply to other ports or LAGs.
2. In the Ports or LAGs list, select the check box(es) in the **Copy To** column that will have the same settings as the port selected in the **Copy Parameters From** field.
3. Click **Apply** to copy the settings.

GMRP Parameters

Use the **GMRP Parameters** page to configure the administrative mode of GMRP on the switch and on each port or LAG.

To display the **GMRP Parameters** page, click **Switching → GARP → GMRP Parameters** in the navigation panel.

Figure 26-33. GMRP Parameters



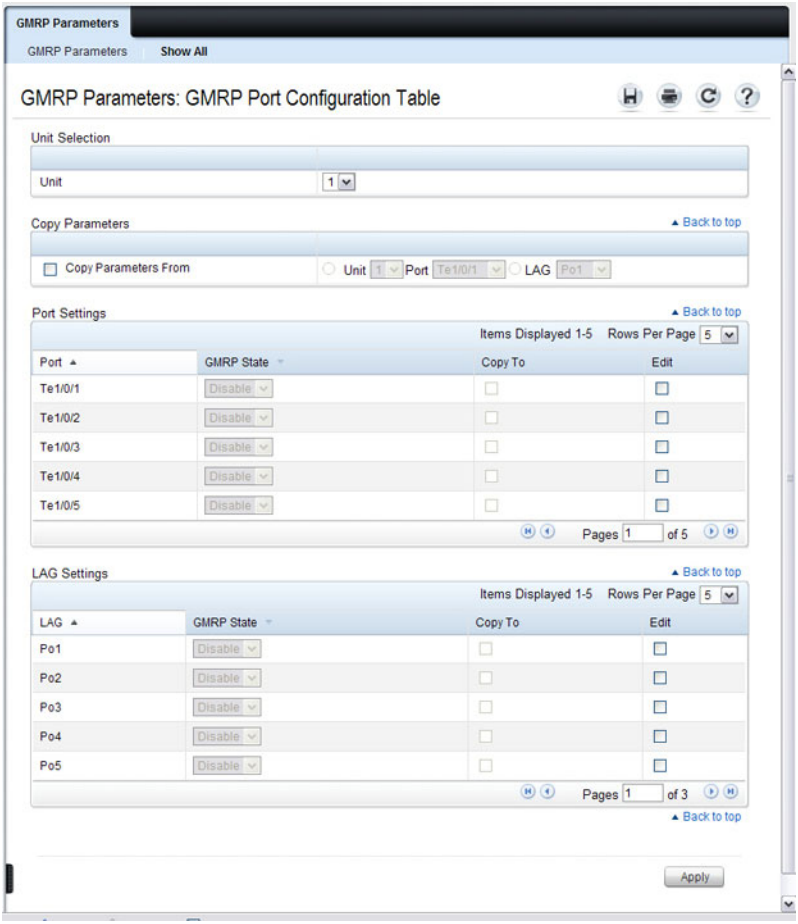
Configuring GMRP Parameters on Multiple Ports

To configure GMRP settings:

1. Open the **GMRP Parameters** page.

2 Click **Show All** to display the **GMRP Port Configuration Table**.

Figure 26-34. GMRP Port Configuration Table



3 For each port or LAG to configure, select the check box in the **Edit** column in the row associated with the port.

4 Specify the desired timer values.

5 Click **Apply**.

Copying Settings From One Port or LAG to Others

To copy GMRP settings:

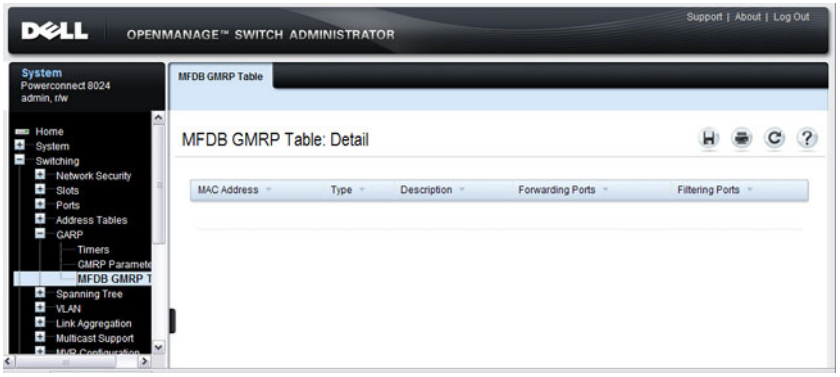
1. Select the **Copy Parameters From** check box, and select the port or LAG with the settings to apply to other ports or LAGs.
2. In the Ports or LAGs list, select the check box(es) in the **Copy To** column that will have the same settings as the port selected in the **Copy Parameters From** field.
3. Click **Apply** to copy the settings.

MFDB GMRP Table

Use the **MFDB GMRP Table** page to view all of the entries in the Multicast Forwarding Database that were created for the GMRP

To display the **MFDB GMRP Table** page, click **Switching** → **GARP** → **MFDB GMRP Table** in the navigation panel.

Figure 26-35. MFDB GMRP Table



Configuring L2 Multicast Features (CLI)

This section provides information about the commands you use to configure L2 multicast settings on the switch. For more information about the commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Configuring Bridge Multicasting

Beginning in Privileged EXEC mode, use the following commands to configure MAC address table features.

Command	Purpose
configure	Enter global configuration mode.
mac address-table multicast filtering	Enable multicast filtering on the switch.
mac address-table multicast static vlan *vlan-id* {*mac-multicast-address* | i*p-multicast-address*}	Register a MAC-layer Multicast address in the bridge table. • *mac-multicast-address* — MAC multicast address in the format xxxx.xxxx.xxxx. • *ip- multicast-address* — IP multicast address.
mac address-table multicast static vlan *vlan-id* {*mac-multicast-address* | i*p-multicast-address*} [**add**|**remove**] **interface** *interface-list*	Add ports and LAGs to the multicast group as static members. • **add** — Adds ports to the group. If no option is specified, this is the default option. • **remove** — Removes ports from the group. • *interface-list* — Specifies the interface type (**port-channel**, **tengigabitethernet**) and number. Separate nonconsecutive interfaces with a comma and no spaces; use a hyphen to designate a range of ports.

Command	Purpose
mac address-table multicast forbidden address vlan *vlan-id* {*mac-multicast-address* | *ip-multicast-address*} {**add** | **remove**} **interface** *interface-list*	Forbid adding a specific Multicast address to specific ports. • *mac-multicast-address* — MAC multicast address in the format xxxx.xxxx.xxxx. • *ip- multicast-address* — IP multicast address. • **add** — Adds ports to the group. If no option is specified, this is the default option. • **remove** — Removes ports from the group. • *interface-list* — Specifies the interface type (**port-channel**, **tengigabitethernet**) and number. Separate nonconsecutive interfaces with a comma and no spaces; use a hyphen to designate a range of ports.
mac address-table multicast forbidden forward-unregistered vlan *vlan-id*	Prevent the switch from forwarding traffic with unregistered multicast addresses on the specified VLAN. This command sets the forwarding mode to Filter Unregistered.
mac address-table multicast forward-all vlan *vlan-id*	Allow the switch to forward all multicast packets on the specified VLAN.
mac address-table multicast forward-unregistered vlan *vlan-id*	Allow the switch to forward packets with unregistered multicast addresses on the specified VLAN. **NOTE:** Do not use the **mac address-table multicast forbidden forward-unregistered** command with the **mac address-table multicast forward-unregistered** command on the same interface.
exit	Exit to Privileged EXEC mode.
show mac address-table multicast [**vlan** *vlan-id*] [**address** *mac-multicast-address* | *ip-multicast-address*] [**format ip** | **mac**]]	View entries in the multicast MAC address table.
show mac address-table filtering *vlan-id*	View the multicast filtering configuration information for the specified VLAN.

Configuring IGMP Snooping

Beginning in Privileged EXEC mode, use the following commands to configure IGMP snooping settings on the switch, ports, and LAGs.

Command	Purpose
configure	Enter global configuration mode.
ip igmp snooping	Globally enable IGMP snooping on the switch. (IGMP snooping is disabled by default.)
interface *interface*	Enter interface configuration mode for the specified port or LAG. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**. For a LAG, the interface type is **port-channel**. You can also specify a range of ports with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12.
ip igmp snooping	Enable IGMP snooping (auto-learn) on the interface.
ip igmp snooping host-time-out *time-out*	Specify the host time-out value for the interface. If an IGMP report for a multicast group is not received in the number of seconds specified by the *time-out* value, this interface is deleted from the member list of that multicast group.
ip igmp snooping leavetime-out {***time-out*** **| immediate-leave**}	Specify the leave time-out value for an interface. If an IGMP report for a multicast group is not received within the number of seconds specified by the leave-time-out period after an IGMP leave was received from a specific interface, the current interface is deleted from the member list of that multicast group. • time-out — Specifies the leave-time-out in seconds. (Range: 1 - 25) • immediate-leave — Specifies that the interface should be removed immediately from the members list after receiving IGMP Leave.

Command	Purpose
ip igmp snooping mrouter-time-out *time-out*	Specify the multicast router time-out value for an interface. This command sets the number of seconds to wait to age out an automatically-learned multicast router port.
CTRL + Z	Exit to Privileged EXEC mode.
show ip igmp snooping	View IGMP snooping settings configured on the switch.
show ip igmp snooping interface *interface*	View the IGMP snooping settings for a specific port or LAG.

Configuring IGMP Snooping on VLANs

Beginning in Privileged EXEC mode, use the following commands to configure IGMP snooping settings on VLANs.

Command	Purpose
configure	Enter global configuration mode.
vlan database	Enter VLAN Configuration mode.
ip igmp snooping *vlan-id*	Enable IGMP snooping on the specified VLAN.
ip igmp snooping groupmembership-interval *vlan-id seconds*	Specify the host time-out value for the specified VLAN. If an IGMP report for a multicast group is not received in the number of seconds specified by the *seconds* value, this VLAN is deleted from the member list of that multicast group.
ip igmp snooping maxresponse *vlan-id seconds*	Specify the leave time-out value for the VLAN. If an IGMP report for a multicast group is not received within the number of seconds configured with this command after an IGMP leave was received from a specific interface, the current VLAN is deleted from the member list of that multicast group.
ip igmp snooping fast-leave *vlan-id*	Enables IGMP snooping fast-leave mode on the specified VLAN. Enabling fast-leave allows the switch to immediately remove the layer 2 LAN interface from its forwarding table entry upon receiving an IGMP leave message for that multicast group without first sending out MAC-based general queries to the interface.

Command	Purpose
ip igmp snooping mcrtexpiretime *vlan-id seconds*	Specify the multicast router time-out value for to associate with a VLAN. This command sets the number of seconds to wait to age out an automatically-learned multicast router port.
CTRL + Z	Exit to Privileged EXEC mode.
show ip igmp snooping groups	Shows IGMP snooping configuration on all VLANs.
show ip igmp snooping vlan *vlan-id*	View the IGMP snooping settings on the VLAN.

Configuring IGMP Snooping Querier

Beginning in Privileged EXEC mode, use the following commands to configure IGMP snooping querier settings on the switch and on VLANs.

Command	Purpose
configure	Enter global configuration mode.
ip igmp snooping querier [**vlan** *vlan-id*] [**address** *ip-address*]	Enable the IGMP snooping querier on the switch or on the VLAN specified with the *vlan-id* parameter. Use the optional *ip-address* parameter to specify the IP address that the snooping querier switch should use as the source address when generating periodic queries.
ip igmp snooping querier query-interval *interval-count*	Set the IGMP snooping querier query interval time, which is the amount of time in seconds that the switch waits before sending another periodic query. The range is 1–1800 seconds.
ip igmp snooping querier timer expiry *seconds*	Set the IGMP snooping querier timer expiration period. This is the time period, in seconds, that the switch remains in non-querier mode after it has discovered that there is a multicast querier in the network.
ip igmp snooping querier version *version*	Set the IGMP version of the query that the switch sends periodically. The *version* range is 1–2.
vlan database	Enter VLAN Configuration mode.
ip igmp snooping querier *vlan-id*	Enable the IGMP snooping querier on the specified VLAN.

Command	Purpose
ip igmp snooping querier election participate *vlan-id*	Allow the IGMP snooping querier to participate in the querier election process when it discovers the presence of another querier in the VLAN. When this mode is enabled, if the snooping querier finds that the other querier source address is more than the snooping querier address, it stops sending periodic queries. If the snooping querier wins the election, then it continues sending periodic queries.
CTRL + Z	Exit to Privileged EXEC mode.
show ip igmp snooping querier [detail |vlan *vlan-id*]	View IGMP snooping querier settings configured on the switch, on all VLANs, or on the specified VLAN.

Configuring MLD Snooping

Beginning in Privileged EXEC mode, use the following commands to configure MLD snooping settings on the switch, ports, and LAGs.

Command	Purpose
configure	Enter global configuration mode.
ipv6 mld snooping	Enable MLD snooping on the switch.
interface *interface*	Enter interface configuration mode for the specified port or LAG. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**. For a LAG, the interface type is **port-channel**. You can also specify a range of ports with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12.
ipv6 mld snooping	Enable MLD snooping (auto-learn) on the interface.
ipv6 mld snooping groupmembership-interval *seconds*	Specify the host time-out value for the interface. If an MLD report for a multicast group is not received in the number of seconds specified by the *seconds* value, this interface is deleted from the member list of that multicast group.

Command	Purpose
ipv6 mld snooping maxresponse *seconds*	Specify the leave time-out value for an interface. If an MLD report for a multicast group is not received within the number of *seconds* specified by the leave-time-out period after an MLD leave was received from a specific interface, the current interface is deleted from the member list of that multicast group.
ipv6 mld snooping immediate-leave	Enables MLD snooping immediate-leave mode on the interface. Enabling immediate-leave allows the switch to immediately remove the layer 2 LAN interface from its forwarding table entry upon receiving an MLD leave message for that multicast group without first sending out MAC-based general queries to the interface.
ipv6 mld snooping mcrtrexpiretime *seconds*	Specify the multicast router time-out value for an interface. This command sets the number of seconds to wait to age out an automatically-learned multicast router port.
CTRL + Z	Exit to Privileged EXEC mode.
show ipv6 mld snooping	View MLD snooping settings configured on the switch.
show ipv6 mld snooping interface *interface*	View the MLD snooping settings for a specific port, LAG, or VLAN.

Configuring MLD Snooping on VLANs

Beginning in Privileged EXEC mode, use the following commands to configure MLD snooping settings on VLANs.

Command	Purpose
configure	Enter global configuration mode.
vlan database	Enter VLAN Configuration mode.
ipv6 mld snooping *vlan-id*	Enable MLD snooping on the specified VLAN.
ipv6 mld snooping groupmembership-interval *vlan-id seconds*	Specify the host time-out value for the specified VLAN. If an MLD report for a multicast group is not received in the number of seconds specified by the *seconds* value, this VLAN is deleted from the member list of that multicast group.

Command	Purpose
ipv6 mld snooping maxresponse *vlan-id seconds*	Specify the leave time-out value for the VLAN. If an MLD report for a multicast group is not received within the number of seconds configured with this command after an MLD leave was received from a specific interface, the current VLAN is deleted from the member list of that multicast group.
ipv6 mld snooping fast-leave *vlan-id*	Enables MLD snooping fast-leave mode on the specified VLAN. Enabling fast-leave allows the switch to immediately remove the layer 2 LAN interface from its forwarding table entry upon receiving an MLD leave message for that multicast group without first sending out MAC-based general queries to the interface.
ipv6 mld snooping mcrtexpiretime *vlan-id seconds*	Specify the multicast router time-out value for to associate with a VLAN. This command sets the number of seconds to wait to age out an automatically-learned multicast router port.
CTRL + Z	Exit to Privileged EXEC mode.
show ipv6 mld snooping vlan *vlan-id*	View the MLD snooping settings on the VLAN.

Configuring MLD Snooping Querier

Beginning in Privileged EXEC mode, use the following commands to configure MLD snooping querier settings on the switch and on VLANs.

Command	Purpose
configure	Enter global configuration mode.
ipv6 mld snooping querier	Enable the MLD snooping querier on the switch.
vlan database	Enter VLAN Configuration mode
ipv6 mld snooping querier *vlan-id* [**address** *ipv6-address*]	Enable the MLD snooping querier on VLAN specified with the *vlan-id* parameter. Use the optional *ip-address* parameter to specify the IP address that the snooping querier switch should use as the source address when generating periodic queries.

Command	Purpose
ipv6 mld snooping querier election participate *vlan-id*	Allow the MLD snooping querier to participate in the querier election process when it discovers the presence of another querier in the VLAN. When this mode is enabled, if the snooping querier finds that the other querier source address is more than the snooping querier address, it stops sending periodic queries. If the snooping querier wins the election, then it continues sending periodic queries.
exit	Exit to Global Configuration mode.
ipv6 mld snooping querier address *ipv6-address*	Specify the IP address that the snooping querier switch should use as the source address when generating periodic queries.
ipv6 mld snooping querier query-interval *interval-count*	Set the MLD snooping querier query interval time, which is the amount of time in seconds that the switch waits before sending another periodic query. The range is 1–1800 seconds.
ipv6 mld snooping querier timer expiry *seconds*	Set the MLD snooping querier timer expiration period. This is the time period, in seconds, that the switch remains in non-querier mode after it has discovered that there is a multicast querier in the network.
exit	Exit to Privileged EXEC mode.
show ipv6 mld snooping querier [detail |vlan *vlan-id*]	View MLD snooping querier settings configured on the switch, on all VLANs, or on the specified VLAN.

Configuring MVR

Beginning in Privileged EXEC mode, use the following commands to configure MVR features on the switch.

Command	Purpose
configure	Enter global configuration mode.
mvr	Enable MVR on the switch.
mvr vlan *vlan-id*	Set the VLAN to use as the multicast VLAN for MVR.
mvr querytime *time*	Set the MVR query response time. The value for *time* is in units of tenths of a second.

Command	Purpose
mvr mode {compatible | dynamic}	Specify the MVR mode of operation.
mvr group *mcast-address* [*groups*]	Add an MVR membership group. • *mcast-address*—The group IP multicast address • *group*—Specifies the number of contiguous groups
interface *interface*	Enter interface configuration mode for the specified port. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**. You can also specify a range of ports with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12.
mvr	Enable MVR on the port.
mvr immediate	Enable MVR immediate leave mode on the port.
mvr type {source | receiver}	Specify the MVR port type.
mvr vlan *vlan-id* **group** *mcast-address*	Allow the port to participate in the specified MVR group. The *vlan-id* parameter is the ID of the MVR multicast VLAN.
CTRL + Z	Exit to Privileged EXEC mode.
show ip dhcp snooping [interfaces]	View the DHCP snooping global and per port configuration.
show ip dhcp snooping binding [{**static | dynamic**}] [**interface** *port*] [**vlan** *vlan-id*]	View the entries in the DHCP snooping bindings database.
show mvr	View information about the administrative mode of MVR.
show mvr members	View information about MVR groups and their members.
show mvr interface *interface*	View information about the MVR configuration for a specific port.
show mvr traffic	View information about IGMP traffic in the MVR table.

Configuring GARP Timers and GMRP

Beginning in Privileged EXEC mode, use the following commands to configure the GARP timers and to control the administrative mode GMRP on the switch and per-interface.

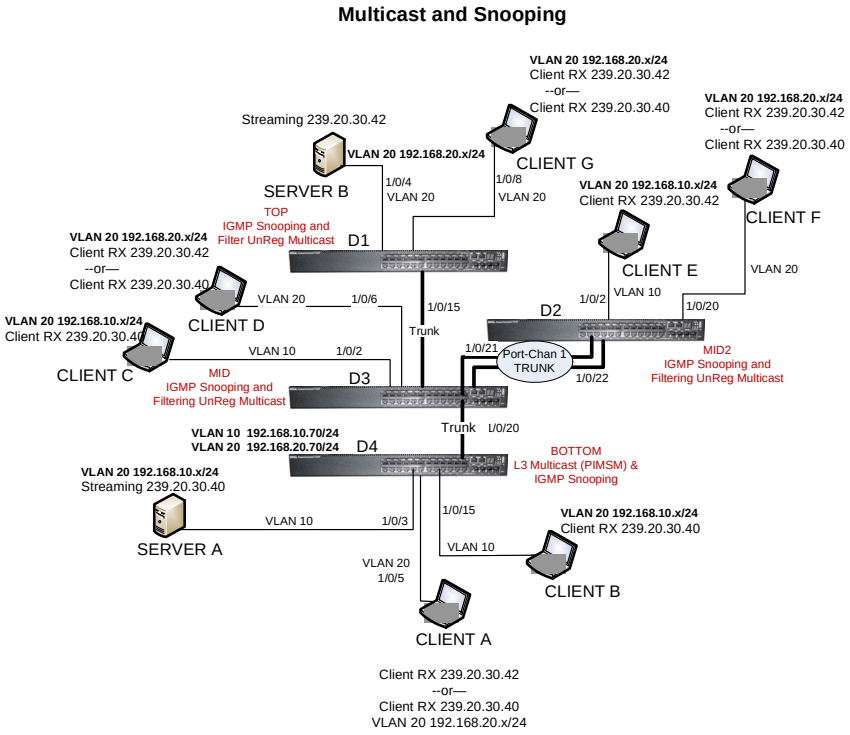
Command	Purpose
configure	Enter global configuration mode.
garp timer {**join** | **leave** | **leaveall**} *timer_value*	Adjust the GARP application join, leave, and leaveall GARP timer values The *timer_value* variable is in centiseconds. The range is 10-100 for **join**, 20-600 for **leave**, and 200-6000 for **leaveall**.
gmrp enable	Enable GMRP globally on the switch.
interface *interface*	Enter interface configuration mode for the specified port or LAG. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**. For a LAG, the interface type is **port-channel**. You can also specify a range of ports with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12.
gmrp enable	Enable GMRP on the interface or range of interfaces.
exit	Exit to Global Config mode.
exit	Exit to Privileged EXEC mode.
show gmrp configuration	View the administrative status of GMRP on the switch and all interfaces.

Case Study on a Real-World Network Topology

Multicast Snooping Case Study

Figure 26-36 shows the topology that the scenarios in this case study use.

Figure 26-36. Case Study Topology



The topology in Figure 26-36 includes the following elements:

- Snooping Switches: D1, D2, D3 with snooping on VLANs 10, 20
- Multicast Router: D4 with PIM-SM and snooping on VLANs 10, 20
- Multicast Listeners: Client A-G
- Multicast Sources: Server A – 239.20.30.40, Server B – 239.20.30.42
- Subnets: VLAN 10 – 192.168.10.70, VLAN 20 – 192.168.10.70

- Router-attached ports: D3 – 1/0/20, D2 – PortChannel1, D1 – 1/0/15

Snooping Within a Subnet

In the example network topology, the multicast source and listeners are in the same subnet VLAN 20 – 192.168.20.70/24. D1 and D3 are configured to drop unregistered multicast traffic. D4 sends periodic queries on VLAN 10, and these queries are forwarded to D1, D2, and D3 via trunk links. Snooping switches D1, D2, and D3 forward these queries to clients G, F, and D, respectively.

Multicast Source and Listener directly connected to a snooping switch: Server B → Client G

1. Client G sends a report for 239.20.30.42.
2. The report is forwarded to multicast router D4 via D1 – 1/0/15 and D3 – 1/0/20.
3. A forwarding entry is created by D1 for VLAN 20, 239.20.30.42 – 1/0/8, 1/0/15.
4. Client G receives the multicast stream from Server B.
5. D3 receives the multicast stream but is dropped because 239.20.30.42 is not registered on D3.
6. Client D sends a report for 239.20.30.42.
7. The report is forwarded to multicast router D4 via D3 – 1/0/20.
8. A forwarding entry is created by D3 for VLAN 20, 239.20.30.42 – 1/0/6, 1/0/20.
9. Client D receives the multicast stream from Server B.
10. Client F does not receive the multicast stream because it did not respond to queries from D4.

Multicast Source and Listener connected by intermediate snooping switches: Server B → Client D

1. Client D sends a report for 239.20.30.42.
2. The report is forwarded to multicast router D4 via D3 – 1/0/20.
3. A forwarding entry is created by D3 for VLAN20, 239.20.30.42 – 1/0/6, 1/0/20.

4 Client D will not receive the multicast stream from Server B because it is dropped at D1, where 239.20.30.42 is not a registered group.

5 The administrator creates a static multicast forwarding entry VLAN 20, 239.20.30.42 – 1/0/15 on D1.

6 Client G will not receive the Server B multicast stream because it did not request it, and the static multicast forwarding entry does not have D1 – 1/0/8 as part of the forwarding list.

7 Server B multicast stream reaches D3.

8 Client D receives the multicast stream from Server B.

9 Client F does not receive the multicast stream because it did not respond to queries from D4.

Snooping on a Multicast Router

In the example network topology, consider Client B and Server A. Both are in the same subnet VLAN10 – 192.168.10.70/24. Server A is a source for multicast stream 239.20.30.40. Snooping switches D1, D2, and D3 are configured to drop unregistered multicast traffic. D4 sends periodic queries on VLAN 10 and VLAN 20, and these queries reach D1, D2, and D3 via trunk links, which in turn forward them in VLAN 10 and VLAN 20 to reach their respective attached clients.

Multicast Source and Listener directly connected to Multicast Router on the same routing VLAN: Server A → Client B

1 Multicast stream 239.20.30.40 is flooded to Client B, and trunk link 1/0/20 on D4 is configured to flood unregistered multicast traffic.

2 Because multicast routing is enabled on D4 VLAN 10, an IP multicast table entry is created to include D4 – 1/0/15, D4 – 1/0/20 as part of the L2 forwarding list members.

3 Client B sends a report for 239.20.30.40.

4 The IP multicast table entry is modified to include only D4 – 1/0/15 as the layer 2 forwarding list member.

5 Client B receives multicast data.

6 The multicast stream is not forwarded to D3 on trunk link 1/0/20 because no other clients requested this data.

Multicast Source directly connected to Multicast Router, and Listener connected to a different routing VLAN via intermediate snooping switches: Server A → Client F

Clients A, D and F are in the same subnet VLAN20 - 192.168.20.70/24. Server A is in a different subnet VLAN10 – 192.168.10.70/24.

1. Client F sends a report for 239.20.30.40.
2. A multicast forwarding entry is created on D2 VLAN20, 239.20.30.40 – 1/0/20, PortChannel1.
3. The Client F report message is forwarded to D3-PortChannel1 (multicast router attached port).
4. A multicast forwarding entry is created on D3 VLAN 20, 239.20.30.40 – PortChannel1, 1/0/20.
5. The Client F report message is forwarded to D4 via D3 – 1/0/20 (multicast router attached port).
6. An IP multicast routing entry is created on D4 VLAN 10 – VLAN 20 with the L3 outgoing port list as VLAN 20 – 1/0/20.
7. The multicast stream is routed to D3.
8. The multicast stream is forwarded to listener Client F using forwarding entries created on D3 and D2.
9. Clients A and D do not receive the Server A multicast stream because they did not send a report.

Multicast Source connected to Multicast Router via intermediate snooping switches, and Listener directly connected to multicast router in a different routing interface: Server B → Client B

Server A and Clients B, C, and E are on the same subnet VLAN10 – 192.168.10.70/24. Server B is in a different subnet VLAN20 – 192.168.20.70/24.

1. Client B sends a report for 239.20.30.42.
2. Multicast Router D4 learns group 239.20.30.42 but cannot identify the source because Server B traffic is dropped by D1 and D3 as it is not a registered group for the snooping switches.
3. Client B cannot receive traffic from Server B.

4 The administrator creates a static multicast forwarding entry on D1 VLAN 20, 239.20.30.42 – 1/0/15 and on D3 VLAN 20, 239.20.30.42 – 1/0/20.

5 The multicast stream from Server B reaches D4 via trunk links because it is a statically registered group on D1 and D3.

6 An IP multicast routing entry is created on D4 VLAN 20 – VLAN 10 with the L3 outgoing port list as VLAN 10 – 1/0/15.

7 Client B receives multicast data from Server B.

8 Server A and Clients C and E do not receive Server B data because no report messages were sent requesting Server B traffic.

Multicast Source and Listener connected to Multicast Router via intermediate snooping switches and are part of different routing VLANs: Server B → Client E

Clients E, B, and C are on the same subnet VLAN10 – 192.168.10.70/24. Server B is in a different subnet VLAN20 – 192.168.20.70/24.

1 Client E sends a report for 239.20.30.42.

2 A multicast forwarding entry is created on D2 VLAN10, 239.20.30.42 – 1/0/2, PortChannel 1.

3 The report from Client E is forwarded to D3 via D2 – PortChannel 1.

4 A multicast forwarding entry is created on D3 VLAN10, 239.20.30.42 – PortChannel 1, 1/0/20.

5 The report from Client E is forwarded to D4 via D3 – 1/0/20.

6 Multicast Router D4 learns group 239.20.30.42 but cannot identify the source because Server B traffic is dropped by D1 and D3. Switches D1 and D3 drop the traffic because it is not a registered group for the snooping switches.

7 Client E cannot receive traffic from Server B.

8 The administrator creates a static multicast forwarding entry on D1 VLAN 20, 239.20.30.42 – 1/0/15 and on D3 VLAN 20, 239.20.30.42 – 1/0/20.

9 The multicast stream from Server B reaches D4 via trunk links because it is a statically registered group on D1 and D3.

10 An IP multicast routing entry is created on D4 VLAN 20 – VLAN 10 with the L3 outgoing port list as VLAN 10 – 1/0/20.

11 Client E receives multicast data from Server B.

12 Clients B and C do not receive Server B data because no report messages were sent requesting Server B traffic.

27

Snooping and Inspecting Traffic

This chapter describes Dynamic Host Configuration Protocol (DHCP) Snooping, IP Source Guard (IPSG), and Dynamic ARP Inspection (DAI), which are layer 2 security features that examine traffic to help prevent accidental and malicious attacks on the switch or network.

The topics covered in this chapter include:

- Traffic Snooping and Inspection Overview
- Default Traffic Snooping and Inspection Values
- Configuring Traffic Snooping and Inspection (Web)
- Configuring Traffic Snooping and Inspection (CLI)
- Traffic Snooping and Inspection Configuration Examples

Traffic Snooping and Inspection Overview

DHCP Snooping is a security feature that monitors DHCP messages between a DHCP client and DHCP server to filter harmful DHCP messages and to build a bindings database. The IPSG and DAI features use the DHCP Snooping bindings database to help enforce switch and network security.

IP Source Guard allows the switch to drop incoming packets that do not match a binding in the bindings database. Dynamic ARP Inspection allows the switch to drop ARP packets whose sender MAC address and sender IP address do not match an entry in the DHCP snooping bindings database.

What Is DHCP Snooping?

Dynamic Host Configuration Protocol (DHCP) Snooping is a security feature that monitors DHCP messages between a DHCP client and DHCP server to accomplish the following tasks:

- Filter harmful DHCP messages
- Build a bindings database with entries that consist of the following information:
  - MAC address
  - IP address
  - VLAN ID
  - Client port

Entries in the bindings database are considered to be authorized network clients.

DHCP snooping can be enabled on VLANs, and the trust status (trusted or untrusted) is specified on individual physical ports or LAGS that are members of a VLAN. When a port or LAG is configured as untrusted, it could potentially be used to launch a network attack. DHCP servers must be reached through trusted ports.

DHCP snooping enforces the following security rules:

- DHCP packets from a DHCP server (DHCPOFFER, DHCPACK, DHCPNAK, DHCPRELEASEQUERY) are dropped if they are received on an untrusted port.
- DHCPRELEASE and DHCPDECLINE messages are dropped if the MAC addresses in the snooping database, but the binding's interface is other than the interface where the message was received.
- On untrusted interfaces, the switch drops DHCP packets with a source MAC address that does not match the client hardware address. This is a configurable option.

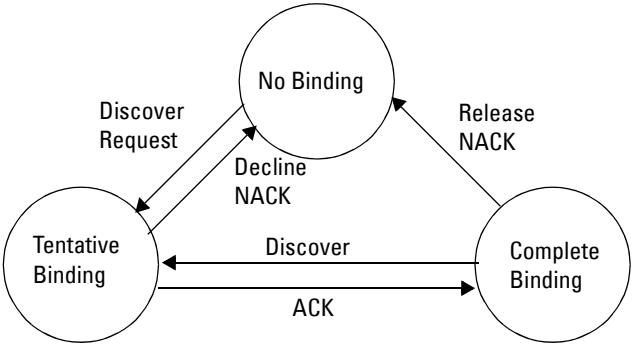
How Is the DHCP Snooping Bindings Database Populated?

The DHCP snooping application uses DHCP messages to build and maintain the binding's database. DHCP snooping creates a tentative binding from DHCP DISCOVER and REQUEST messages. Tentative bindings tie a client to a port (the port where the DHCP client message was received). Tentative bindings are completed when DHCP snooping learns the client's IP address from a DHCP ACK message on a trusted port. DHCP snooping removes bindings in response to DECLINE, RELEASE, and NACK messages. The DHCP snooping application ignores the ACK messages as a reply to the DHCP Inform messages received on trusted ports. You can also enter static bindings into the binding database.

When a switch learns of new bindings or loses bindings, the switch immediately updates the entries in the database. The switch also updates the entries in the binding file. The frequency at which the file is updated is based on a configurable delay, and the updates are batched.

If the absolute lease time of the snooping database entry expires, that entry is removed. Make sure the system time is consistent across the reboots. Otherwise, the snooping entries will not expire properly. If a host sends a DHCP release while the switch is rebooting, when the switch receives the DHCP discovery or request, the client's binding goes to the tentative binding as shown in Figure 27-1.

Figure 27-1. DHCP Binding



The binding database includes data for clients only on untrusted ports.

DHCP Snooping and VLANs

DHCP snooping forwards valid DHCP client messages received on non-routing VLANs. The message is forwarded on all trusted interfaces in the VLAN.

DHCP snooping can be configured on switching VLANs and routing VLANs. When a DHCP packet is received on a routing VLAN, the DHCP snooping application applies its filtering rules and updates the bindings database. If a client message passes filtering rules, the message is placed into the software forwarding path where it may be processed by the DHCP relay agent, the local DHCP server, or forwarded as an IP packet.

DHCP Snooping Logging and Rate Limits

The DHCP snooping application processes incoming DHCP messages. For DHCPRELEASE and DHCPDECLINE messages, the application compares the receive interface and VLAN with the client interface and VLAN in the bindings database. If the interfaces do not match, the application logs the event and drops the message. For valid client messages, DHCP snooping compares the source MAC address to the DHCP client hardware address. When there is a mismatch, DHCP snooping drops the packet and generates a log message if logging of invalid packets is enabled.

If DHCP relay co-exists with DHCP snooping, DHCP client messages are sent to DHCP relay for further processing.

To prevent DHCP packets from being used as a DoS attack when DHCP snooping is enabled, the snooping application enforces a rate limit for DHCP packets received on interfaces. DHCP snooping monitors the receive rate on each interface separately. If the receive rate exceeds a configurable limit, DHCP snooping brings down the interface. Administrative intervention is necessary to enable the port, either by using the **no shutdown** command in Interface Config mode or on the **Switching → Ports → Port Configuration** page.

What Is IP Source Guard?

IPSG is a security feature that filters IP packets based on source ID. This feature helps protect the network from attacks that use IP address spoofing to compromise or overwhelm the network.

The source ID may be either the source IP address or a {source IP address, source MAC address} pair. You can configure:

- Whether enforcement includes the source MAC address
- Static authorized source IDs

The DHCP snooping bindings database and static IPSG entries identify authorized source IDs. IPSG can be enabled on physical and LAG ports.

If you enable IPSG on a port where DHCP snooping is disabled or where DHCP snooping is enabled but the port is trusted, all IP traffic received on that port is dropped depending on the admin-configured IPSG entries.

IPSG and Port Security

IPSG interacts with port security, also known as port MAC locking, (see "Port Security (Port-MAC Locking)" on page 495) to enforce the source MAC address. Port security controls source MAC address learning in the layer 2 forwarding database (MAC address table). When a frame is received with a previously unlearned source MAC address, port security queries the IPSG feature to determine whether the MAC address belongs to a valid binding.

If IPSG is disabled on the ingress port, IPSG replies that the MAC is valid. If IPSG is enabled on the ingress port, IPSG checks the bindings database. If the MAC address is in the bindings database and the binding matches the VLAN the frame was received on, IPSG replies that the MAC is valid. If the MAC is not in the bindings database, IPSG informs port security that the frame is a security violation.

In the case of an IPSG violation, port security takes whatever action it normally takes upon receipt of an unauthorized frame. Port security limits the number of MAC addresses to a configured maximum. If the limit n is less than the number of stations m in the bindings database, port security allows only n stations to use the port. If $n > m$, port security allows only the stations in the bindings database. For information about configuring the Port Security feature, see "Configuring Port and System Security" on page 457.

What is Dynamic ARP Inspection?

Dynamic ARP Inspection (DAI) is a security feature that rejects invalid and malicious ARP packets. DAI prevents a class of man-in-the-middle attacks where an unfriendly station intercepts traffic for other stations by poisoning the ARP caches of its unsuspecting neighbors. The malicious attacker sends ARP requests or responses mapping another station's IP address to its own MAC address.

When DAI is enabled, the switch drops ARP packets whose sender MAC address and sender IP address do not match an entry in the DHCP snooping bindings database. You can optionally configure additional ARP packet validation.

When DAI is enabled on a VLAN, DAI is enabled on the interfaces (physical ports or LAGs) that are members of that VLAN. Individual interfaces are configured as trusted or untrusted. The trust configuration for DAI is independent of the trust configuration for DHCP snooping.

Optional DAI Features

If the network administrator has configured the option, DAI verifies that the sender MAC address equals the source MAC address in the Ethernet header. There is a configurable option to verify that the target MAC address equals the destination MAC address in the Ethernet header. This check applies only to ARP responses, since the target MAC address is unspecified in ARP requests. You can also enable IP address checking. When this option is enabled, DAI drops ARP packets with an invalid IP address. The following IP addresses are considered invalid:

- 0.0.0.0
- 255.255.255.255
- all IP multicast addresses
- all class E addresses (240.0.0.0/4)
- loopback addresses (in the range 127.0.0.0/8)

Why Is Traffic Snooping and Inspection Necessary?

DHCP Snooping, IPSG, and DAI are security features that can help protect the switch and the network against various types of accidental or malicious attacks. It might be a good idea to enable these features on ports that provide network access to hosts that are in physically unsecured locations or if network users connect nonstandard hosts to the network.

For example, if an employee unknowingly connects a workstation to the network that has a DHCP server, and the DHCP server is enabled, hosts that attempt to acquire network information from the legitimate network DHCP server might obtain incorrect information from the rogue DHCP server. However, if the workstation with the rogue DHCP server is connected to a port that is configured as untrusted and is a member of a DHCP Snooping-enabled VLAN, the port discards the DHCP server messages.

Default Traffic Snooping and Inspection Values

DHCP snooping is disabled globally and on all VLANs by default. Ports are untrusted by default.

Table 27-1. Traffic Snooping Defaults

Parameter	Default Value
DHCP snooping mode	Disabled
DHCP snooping VLAN mode	Disabled on all VLANs
Interface trust state	Disabled (untrusted)
DHCP logging invalid packets	Disabled
DHCP snooping rate limit	No limit
DHCP snooping burst interval	1 second
DHCP snooping binding database storage	Local
DHCP snooping binding database write delay	300 seconds
Static DHCP bindings	None configured
IPSG mode	Disabled on all interfaces
IPSG port security	Disabled on all interfaces

Table 27-1. Traffic Snooping Defaults (Continued)

Parameter	Default Value
Static IPSG bindings	None configured
DAI validate source MAC	Disabled
DAI validate destination MAC	Disabled
DAI validate IP	Disabled
DAI trust state	Disabled (untrusted)
DAI Rate limit	15 packets per second
DAI Burst interval	1 second
DAI mode	Disabled on all VLANs
DAI logging invalid packets	Disabled
DAI ARP ACL	None configured
DAI Static flag	Disabled (validation by ARP ACL and DHCP snooping binding database)

Configuring Traffic Snooping and Inspection (Web)

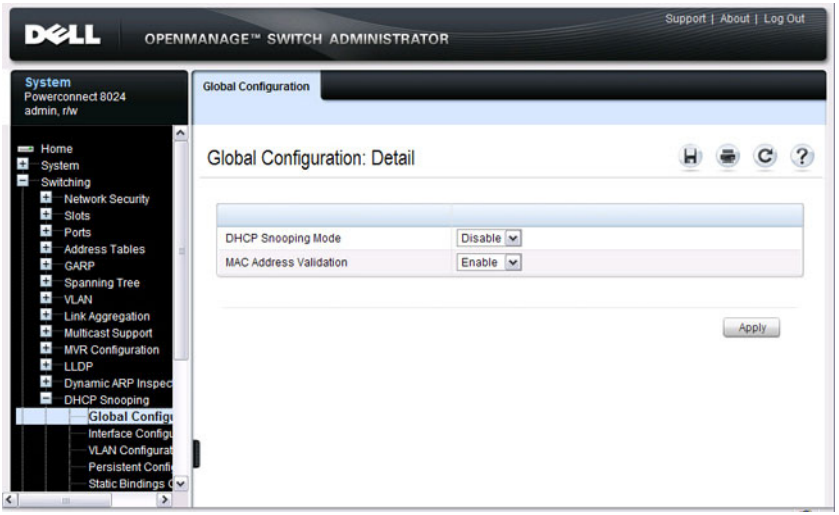
This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring DHCP snooping, IPSG, and DAI features on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

DHCP Snooping Configuration

Use the **DHCP Snooping Configuration** page to control the DHCP Snooping mode on the switch and to specify whether the sender MAC Address for DHCP Snooping must be validated.

To access the **DHCP Snooping Configuration** page, click **Switching** → **DHCP Snooping** → **Global Configuration** in the navigation panel.

Figure 27-2. DHCP Snooping Configuration

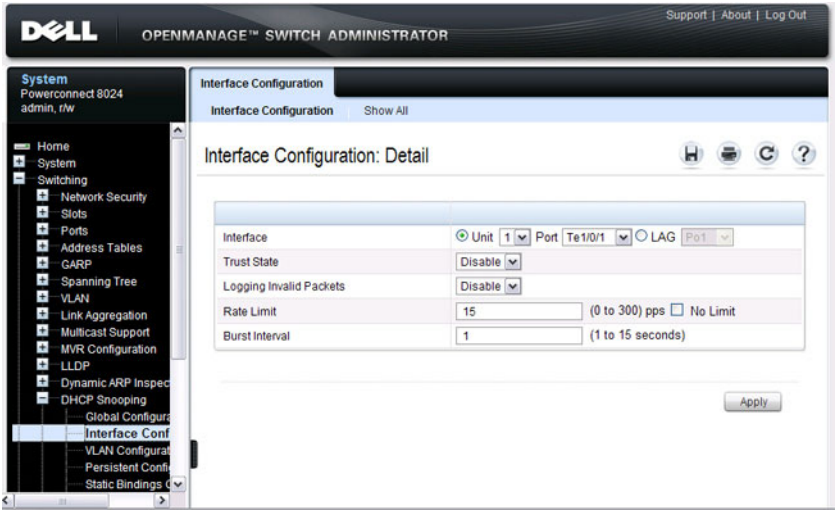


DHCP Snooping Interface Configuration

Use the **DHCP Snooping Interface Configuration** page to configure the DHCP Snooping settings on individual ports and LAGs.

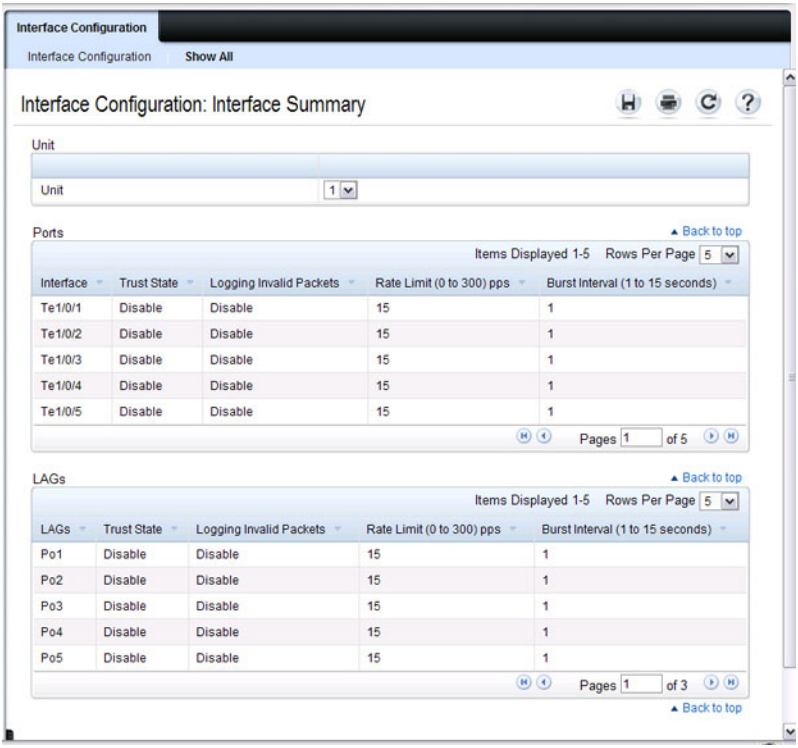
To access the **DHCP Snooping Interface Configuration** page, click **Switching → DHCP Snooping → Interface Configuration** in the navigation panel.

Figure 27-3. DHCP Snooping Interface Configuration



To view a summary of the DHCP snooping configuration for all interfaces, click **Show All**.

Figure 27-4. DHCP Snooping Interface Configuration Summary

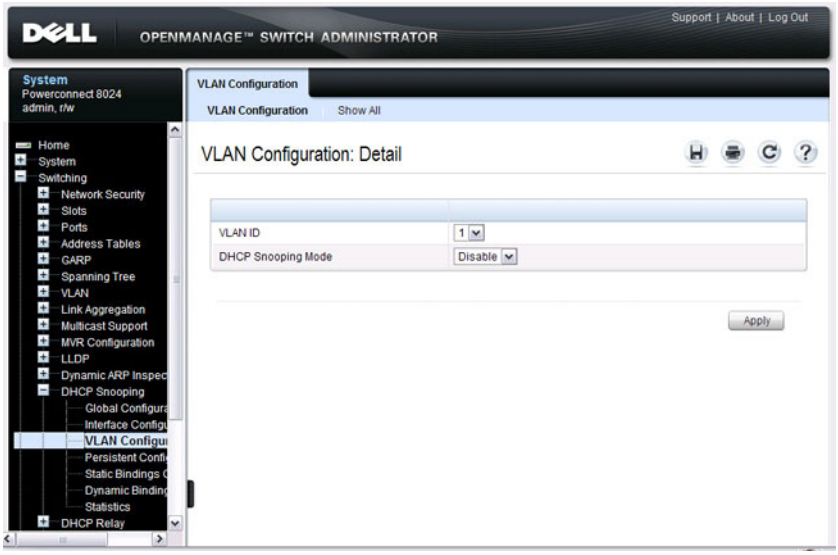


DHCP Snooping VLAN Configuration

Use the **DHCP Snooping VLAN Configuration** page to control the DHCP snooping mode on each VLAN.

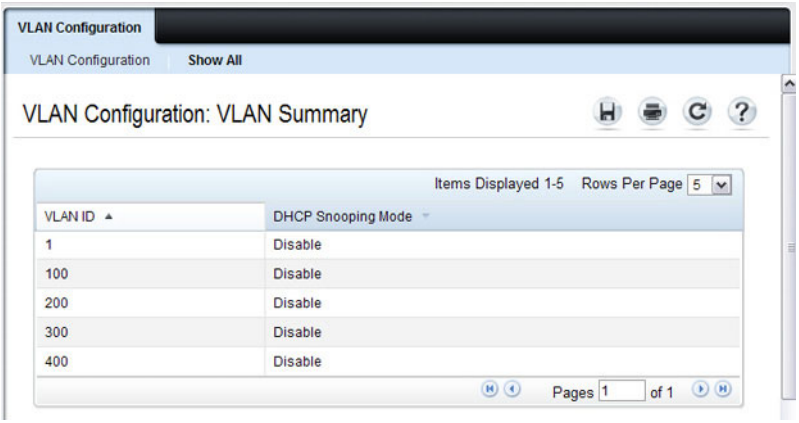
To access the **DHCP Snooping VLAN Configuration** page, click **Switching → DHCP Snooping → VLAN Configuration** in the navigation panel.

Figure 27-5. DHCP Snooping VLAN Configuration



To view a summary of the DHCP snooping status for all VLANs, click **Show All**.

Figure 27-6. DHCP Snooping VLAN Configuration Summary

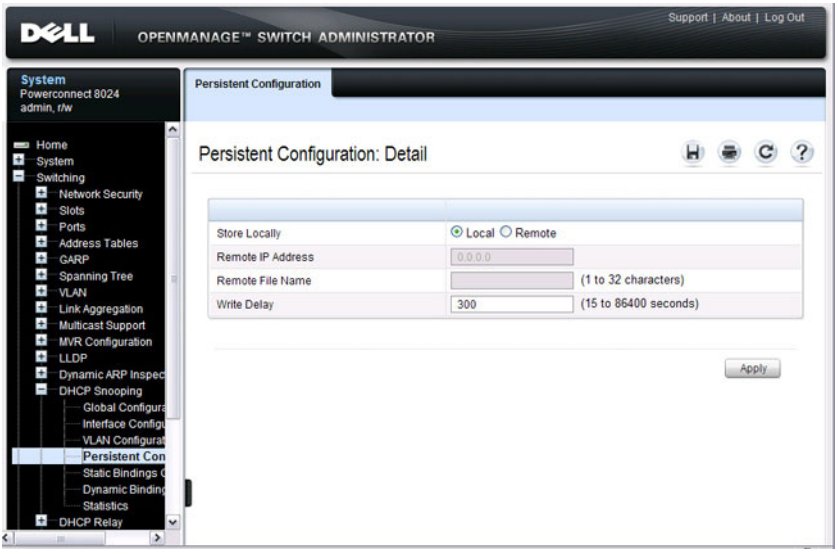


DHCP Snooping Persistent Configuration

Use the **DHCP Snooping Persistent Configuration** page to configure the persistent location of the DHCP snooping database. The bindings database can be stored locally on the switch or on a remote system somewhere else in the network. The switch must be able to reach the IP address of the remote system to send bindings to a remote database.

To access the **DHCP Snooping Persistent Configuration** page, click **Switching → DHCP Snooping → Persistent Configuration** in the navigation panel.

Figure 27-7. DHCP Snooping Persistent Configuration

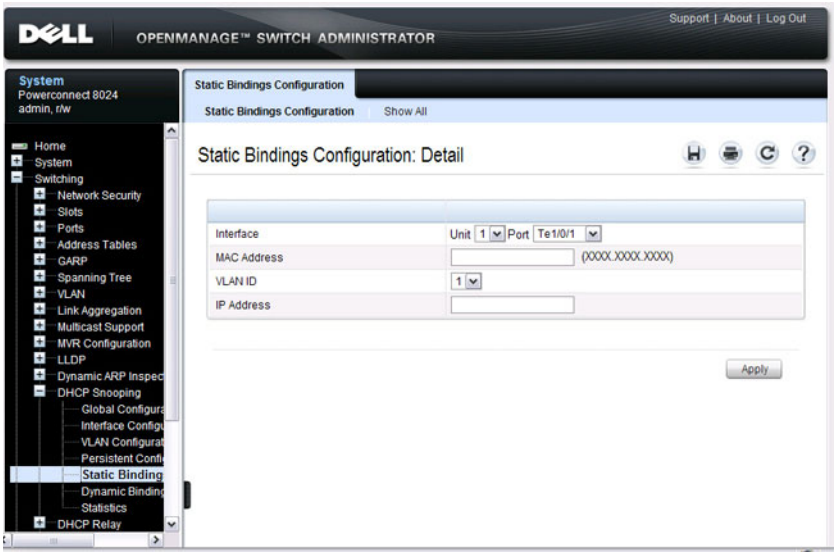


DHCP Snooping Static Bindings Configuration

Use the **DHCP Snooping Static Bindings Configuration** page to add static DHCP bindings to the binding database.

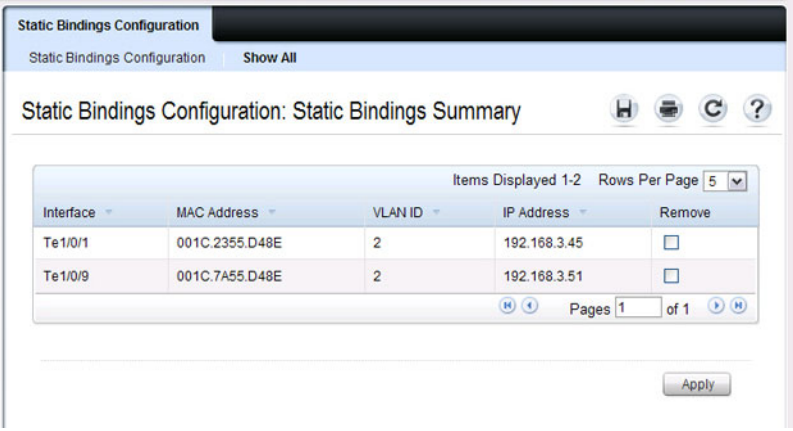
To access the **DHCP Snooping Static Bindings Configuration** page, click **Switching → DHCP Snooping → Static Bindings Configuration** in the navigation panel.

Figure 27-8. DHCP Snooping Static Bindings Configuration



To view a summary of the DHCP snooping status for all VLANs, click **Show All**.

Figure 27-9. DHCP Snooping Static Bindings Summary



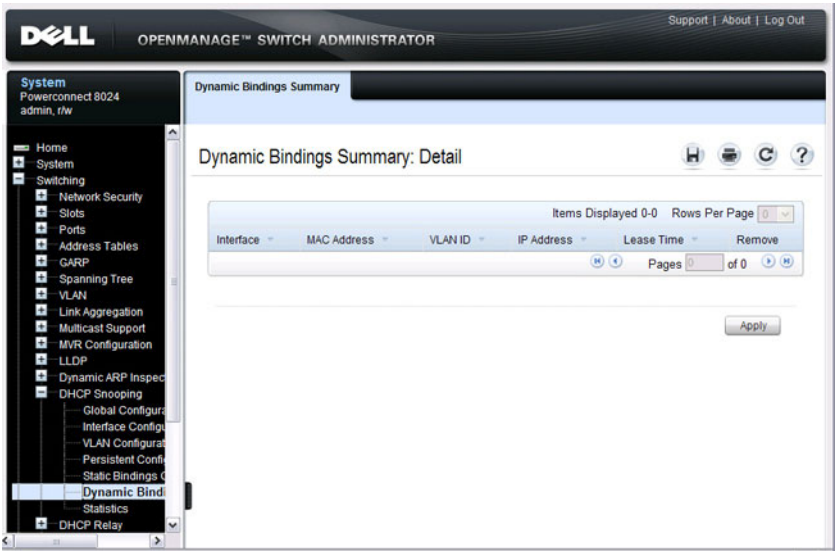
To remove a static binding, select the **Remove** checkbox associated with the binding and click **Apply**.

DHCP Snooping Dynamic Bindings Summary

The **DHCP Snooping Dynamic Bindings Summary** lists all the DHCP snooping dynamic binding entries learned on the switch ports.

To access the **DHCP Snooping Dynamic Bindings Summary** page, click **Switching** → **DHCP Snooping** → **Dynamic Bindings Summary** in the navigation panel.

Figure 27-10. DHCP Snooping Dynamic Bindings Summary

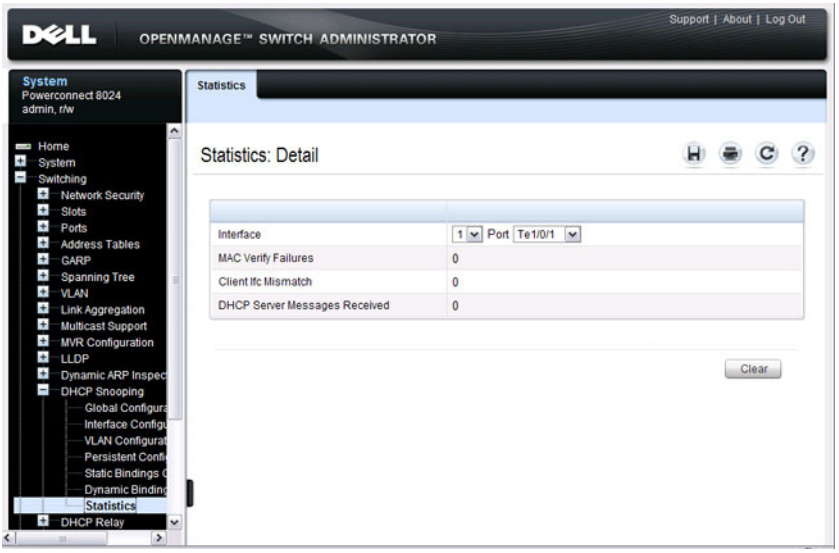


DHCP Snooping Statistics

The **DHCP Snooping Statistics** page displays DHCP snooping interface statistics.

To access the **DHCP Snooping Statistics** page, click **Switching** → **DHCP Snooping** → **Statistics** in the navigation panel.

Figure 27-11. DHCP Snooping Statistics

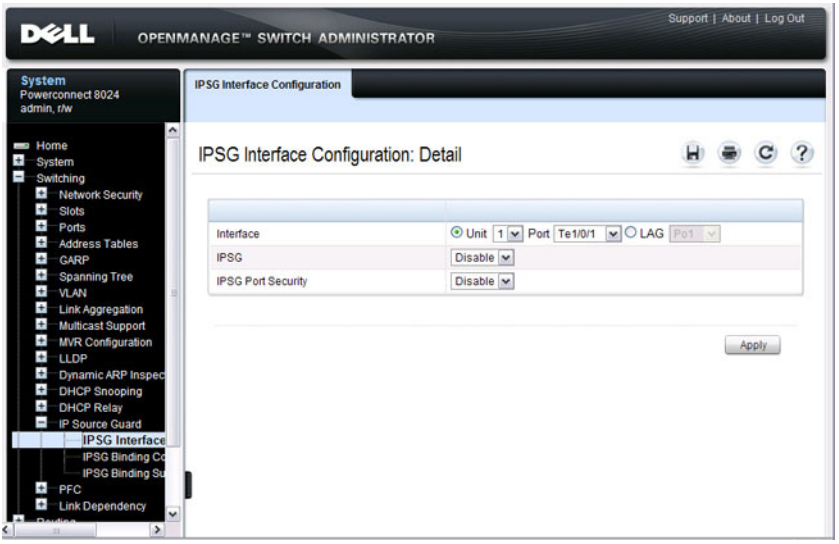


IPSG Interface Configuration

Use the **IPSG Interface Configuration** page to configure IPSG on an interface.

To access the **IPSG Interface Configuration** page, click **Switching → IP Source Guard → IPSG Interface Configuration** in the navigation panel.

Figure 27-12. IPSG Interface Configuration

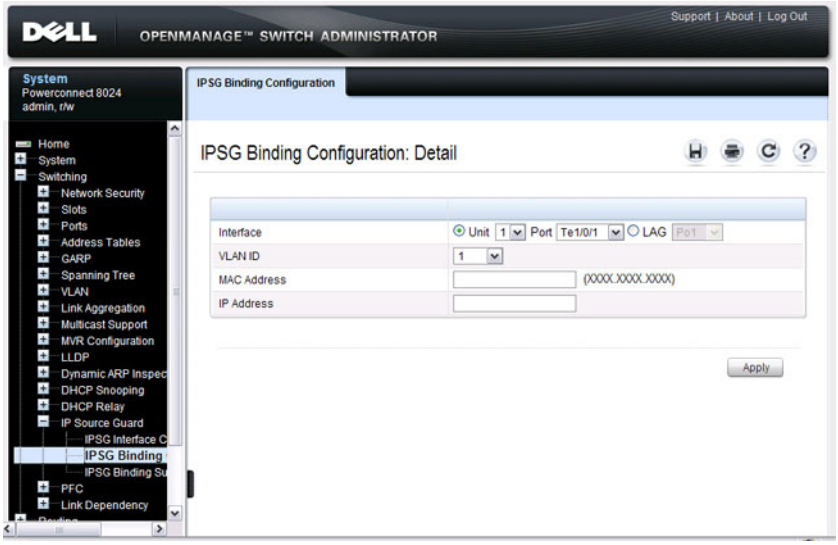


IPSG Binding Configuration

Use the **IPSG Binding Configuration** page displays DHCP snooping interface statistics.

To access the **IPSG Binding Configuration** page, click **Switching → IP Source Guard → IPSG Binding Configuration** in the navigation panel.

Figure 27-13. IPSG Binding Configuration

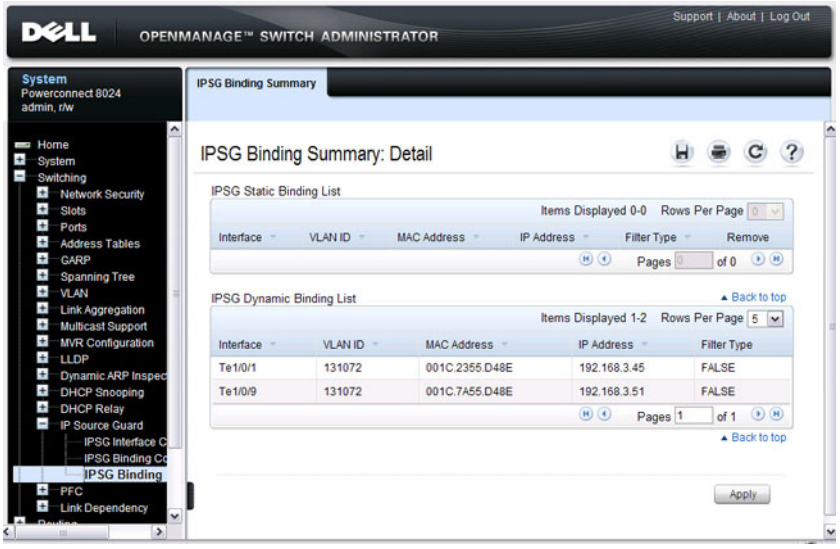


IPSG Binding Summary

The **IPSG Binding Summary** page displays the IPSG Static binding list and IPSG dynamic binding list (the static bindings configured in Binding configuration page).

To access the **IPSG Binding Summary** page, click **Switching → IP Source Guard → IPSG Binding Summary** in the navigation panel.

Figure 27-14. IPSG Binding Summary

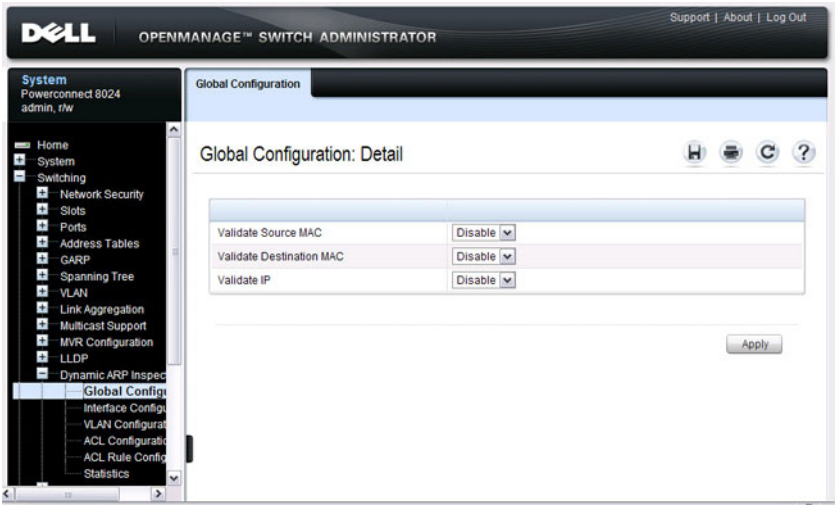


DAI Global Configuration

Use the **DAI Configuration** page to configure global DAI settings.

To display the **DAI Configuration** page, click **Switching** → **Dynamic ARP Inspection** → **Global Configuration** in the navigation panel.

Figure 27-15. Dynamic ARP Inspection Global Configuration

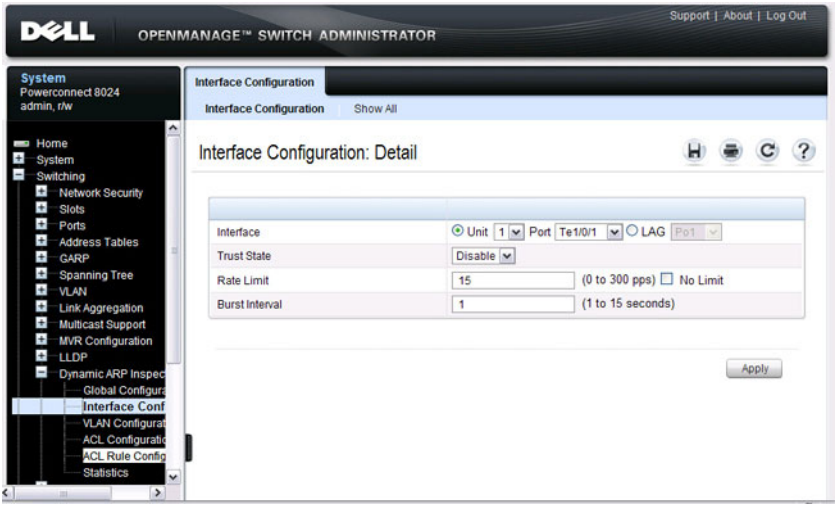


DAI Interface Configuration

Use the **DAI Interface Configuration** page to select the DAI Interface for which information is to be displayed or configured.

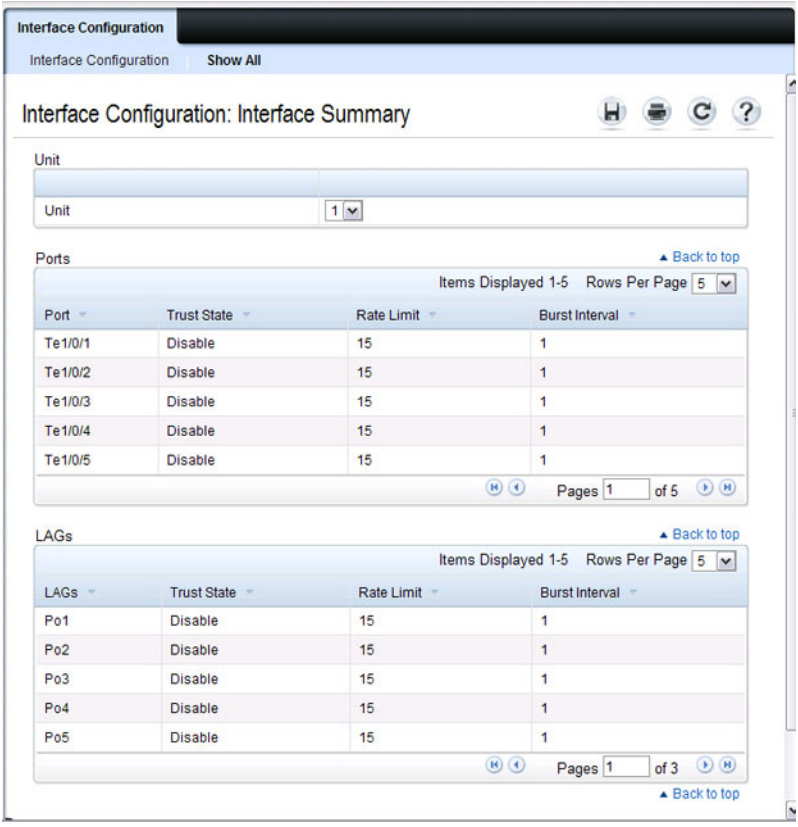
To display the **DAI Interface Configuration** page, click **Switching → Dynamic ARP Inspection → Interface Configuration** in the navigation panel.

Figure 27-16. Dynamic ARP Inspection Interface Configuration



To view a summary of the DAI status for all interfaces, click **Show All**.

Figure 27-17. DAI Interface Configuration Summary

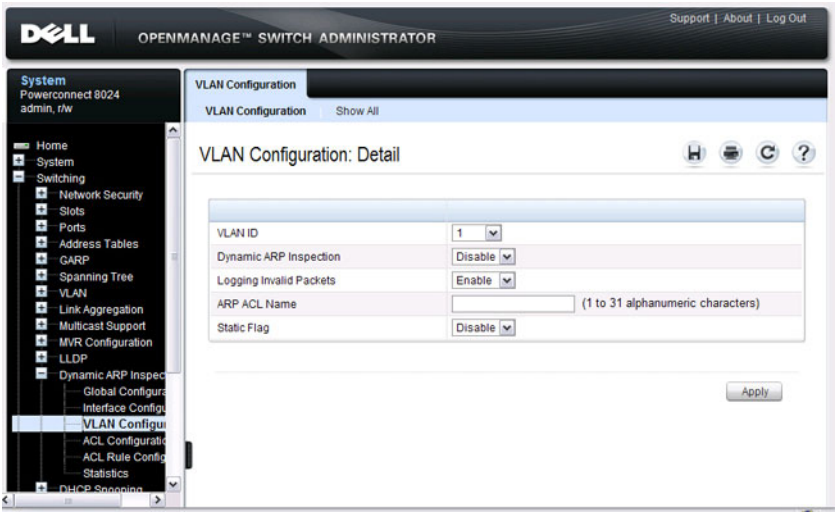


DAI VLAN Configuration

Use the DAI **VLAN Configuration** page to select the VLANs for which information is to be displayed or configured.

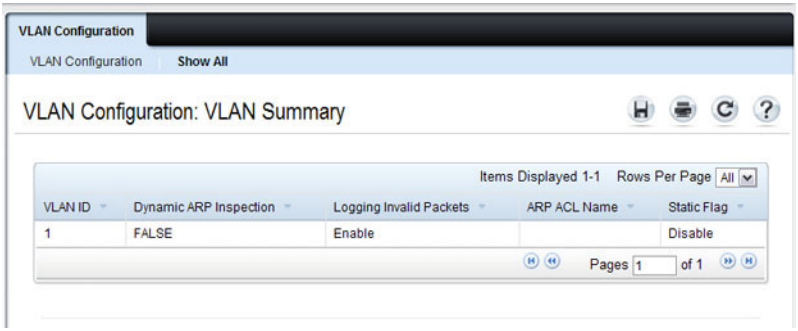
To display the DAI **VLAN Configuration** page, click **Switching** → **Dynamic ARP Inspection** → **VLAN Configuration** in the navigation panel.

Figure 27-18. Dynamic ARP Inspection VLAN Configuration



To view a summary of the DAI status for all VLANs, click **Show All**.

Figure 27-19. Dynamic ARP Inspection VLAN Configuration Summary

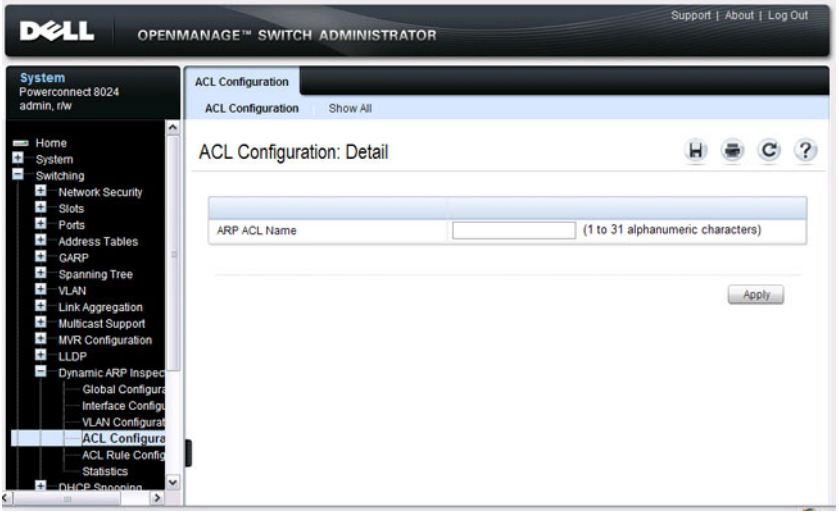


DAI ACL Configuration

Use the DAI **ACL Configuration** page to add or remove ARP ACLs.

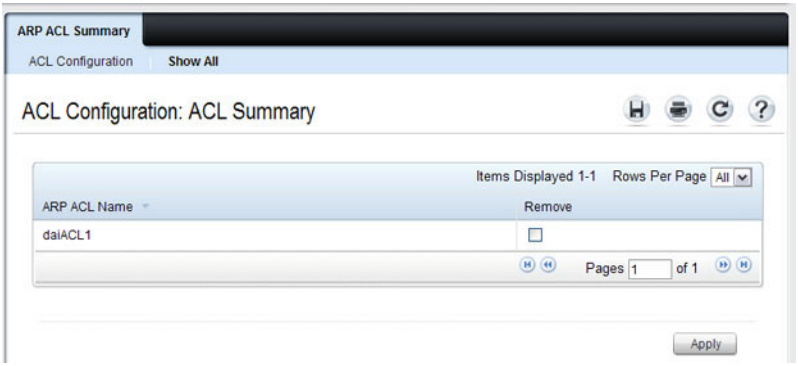
To display the DAI **ACL Configuration** page, click **Switching** → **Dynamic ARP Inspection** → **ACL Configuration** in the navigation panel.

Figure 27-20. Dynamic ARP Inspection ACL Configuration



To view a summary of the ARP ACLs that have been created, click **Show All**.

Figure 27-21. Dynamic ARP Inspection ACL Summary



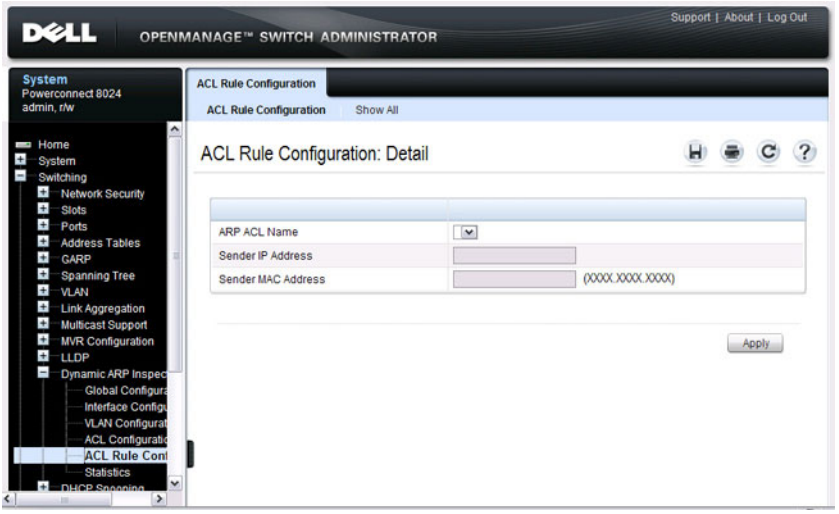
To remove an ARP ACL, select the **Remove** checkbox associated with the ACL and click **Apply**.

DAI ACL Rule Configuration

Use the DAI ARP **ACL Rule Configuration** page to add or remove DAI ARP ACL Rules.

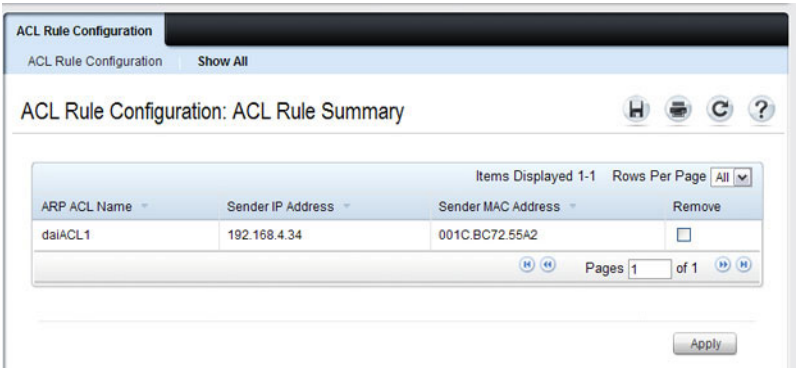
To display the DAI ARP **ACL Rule Configuration** page, click **Switching** → **Dynamic ARP Inspection** → **ACL Rule Configuration** in the navigation panel.

Figure 27-22. Dynamic ARP Inspection Rule Configuration



To view a summary of the ARP ACL rules that have been created, click **Show All**.

Figure 27-23. Dynamic ARP Inspection ACL Rule Summary



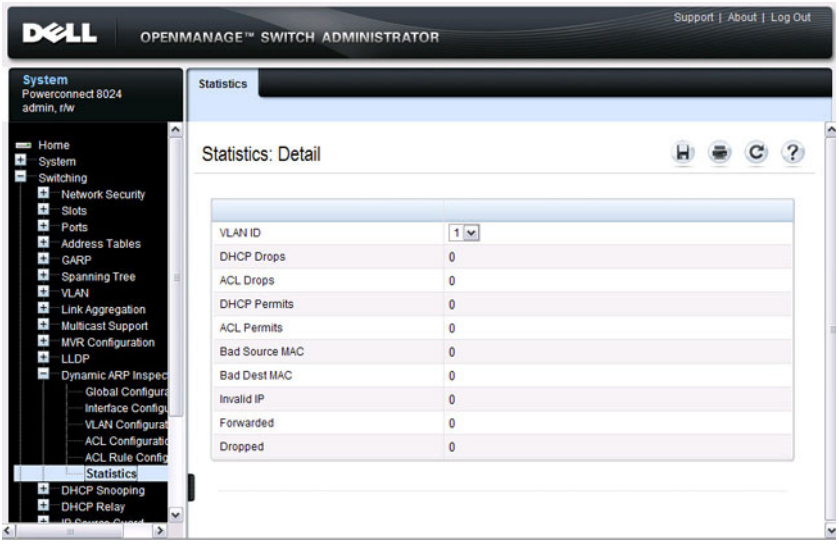
To remove an ARP ACL rule, select the **Remove** checkbox associated with the rule and click **Apply**.

DAI Statistics

Use the DAI **Statistics** page to display the statistics per VLAN.

To display the DAI **Statistics** page, click **Switching** → **Dynamic ARP Inspection** → **Statistics** in the navigation panel.

Figure 27-24. Dynamic ARP Inspection Statistics



Configuring Traffic Snooping and Inspection (CLI)

This section provides information about the commands you use to configure DHCP snooping, IPSG, and DAI settings on the switch. For more information about the commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Configuring DHCP Snooping

Beginning in Privileged EXEC mode, use the following commands to configure and view DHCP snooping settings.

Command	Purpose
configure	Enter global configuration mode.
ip dhcp snooping	Enable DHCP snooping on the switch.
ip dhcp snooping verify mac-address	Enable the verification of the source MAC address with the client MAC address in the received DHCP message.
ip dhcp snooping log-invalid	Enable the logging of DHCP messages filtered by the DHCP Snooping application.
ip dhcp snooping binding mac-address vlan vlan-id ip-address interface interface	Configure a static binding in the DHCP snooping static bindings database. • *mac-address*—The client's MAC address. • *vlan-id*—The number of the VLAN the client is authorized to use. • *ip-address*—The IP address of the client. • *interface*—The interface on which the client is authorized. The form is unit/port.
ip dhcp snooping database {**local** | **tftp:**//*hostIP/filename* }	Configure the persistent storage location of the DHCP snooping database. • *hostIP*—The IP address of the remote host. • *filename*—The name of the file for the database on the remote host.
ip dhcp snooping database write-delay *seconds*	Configure the interval, in seconds, at which the DHCP Snooping database will be stored in persistent storage. The number of seconds can range from 15–86400.

Command	Purpose
ip dhcp snooping limit {**none** | **rate** *rate* [**burst interval** *seconds*]}	Configure the maximum rate of DHCP messages allowed on the switch at any given time. • *rate*—The maximum number of packets per second allowed (Range: 0–300 pps). • *seconds*—The time allowed for a burst (Range: 1–15 seconds).
interface *interface*	Enter interface configuration mode for the specified port or LAG. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**. For a LAG, the interface type is **port-channel**. You can also specify a range of ports with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12.
ip dhcp snooping trust	Configure the interface (or range of interfaces) as a trusted port. DHCP server messages are not filtered on trusted ports.
exit	Exit to Global Configuration mode.
interface [range] vlan *vlan id*	Enter interface configuration mode for the specified VLAN or range of VLANs.
CTRL + Z	Exit to Privileged EXEC mode.
show ip dhcp snooping [interfaces]	View the DHCP snooping global and per port configuration.
show ip dhcp snooping binding [{**static** | **dynamic**}] [**interface** *port*] [**vlan** *vlan-id*]	View the entries in the DHCP snooping bindings database.
show ip dhcp snooping database	View information about the persistent database configuration.
show ip dhcp snooping statistics	View the DHCP snooping statistics.
clear ip dhcp snooping statistics	Reset the DHCP snooping statistics to zero.

Configuring IP Source Guard

Beginning in Privileged EXEC mode, use the following commands to configure IPSG settings on the switch.

Command	Purpose
configure	Enter global configuration mode.
interface *interface*	Enter interface configuration mode for the specified port or LAG. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**. For a LAG, the interface type is **port-channel**.
	You can also specify a range of ports with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12.
ip verify source [**port-security**]	Enable IPSG on the port or LAG to prevent packet forwarding if the source IP address in the packet is not in the DHCP snooping binding database. Use the option **port-security keyword** to also prevent packet forwarding if the sender MAC address is not in forwarding database table or the DHCP snooping binding database. \
	NOTE: To enforce filtering based on the source MAC address, port security must also be enabled on the interface by using the **port security** command in Interface Configuration mode.
exit	Exit to Global Config mode.
ip verify binding *mac_addr* **vlan** *vlan_id ipaddr* **interface** *interface*	Configure a static binding for IPSG.
exit	Exit to Privileged EXEC mode.
show ip verify interface *interface*	View IPSG parameters for a specific port or LAG. The *interface* parameter includes the interface type (**tengigabitethernet** or **port-channel**) and number.
show ip verify source [**interface** *interface*]	View IPSG bindings configured on the switch or on a specific port or LAG.
show ip source binding	View IPSG bindings.

Configuring Dynamic ARP Inspection

Beginning in Privileged EXEC mode, use the following commands to configure DAI settings on the switch.

Command	Purpose
configure	Enter global configuration mode.
ip arp inspection vlan vlan-range [**logging**]	Enable Dynamic ARP Inspection on a single VLAN or a range of VLANs. Use the **logging** keyword to enable logging of invalid packets.
ip arp inspection validate {[**src-mac**] [**dst-mac**] [**ip**]}	Enable additional validation checks like source MAC address validation, destination MAC address validation, or IP address validation on the received ARP packets. Each command overrides the configuration of the previous command. For example, if a command enables source MAC address and destination validations and a second command enables IP address validation only, the source MAC address and destination MAC address validations are disabled as a result of the second command. • **src-mac**—For validating the source MAC address of an ARP packet. • **dst-mac**—For validating the destination MAC address of an ARP packet. • **ip**—For validating the IP address of an ARP packet.
arp access-list *acl-name*	Create an ARP ACL with the specified name (1–31 characters) and enter ARP Access-list Configuration mode for the ACL.
permit ip host *sender-ip* **mac host** *sender-mac*	Configure a rule for a valid IP address and MAC address combination used in ARP packet validation. • *sender-ip* — Valid IP address used by a host. • *sender-mac* —Valid MAC address in combination with the above sender-ip used by a host.
exit	Exit to Global Config mode.

Command	Purpose
ip arp inspection filter *acl-name* **vlan** *vlan-range* [**static**]	Configure the ARP ACL to be used for a single VLAN or a range of VLANs to filter invalid ARP packets. Use the **static** keyword to indicate that packets that do not match a permit statement are dropped without consulting the DHCP snooping bindings.
interface *interface*	Enter interface configuration mode for the specified port or LAG. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**. For a LAG, the interface type is **port-channel**. You can also specify a range of ports with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12.
ip arp inspection limit {**none** | **rate** *pps* [**burst interval** *seconds*]}	Configure the rate limit and burst interval values for an interface.Use the keyword **none** to specify that the interface is not rate limited for Dynamic ARP Inspection. • **none** — To set no rate limit. • *pps* — Packets per second (Range: 0–300). • *seconds* — The number of seconds (Range: 1–15).
ip arp inspection trust	Specify that the interface as trusted for Dynamic ARP Inspection.
CTRL + Z	Exit to Privileged EXEC mode.
show ip arp inspection interfaces [*interface*]	View the Dynamic ARP Inspection configuration on all the DAI-enabled interfaces or for the specified interface.
show ip arp inspection vlan [*vlan-range*]	View the Dynamic ARP Inspection configuration on the specified VLAN(s). This command also displays the global configuration values for source MAC validation, destination MAC validation and invalid IP validation.
show ip arp inspection statistics [**vlan** *vlan-range*]	View the statistics of the ARP packets processed by Dynamic ARP Inspection for the switch or for the specified VLAN(s).
show arp access-list [*acl-name*]	View all configured ARP ACL and their rules, or use the ACL name to view information about that ARP ACL only.

Traffic Snooping and Inspection Configuration Examples

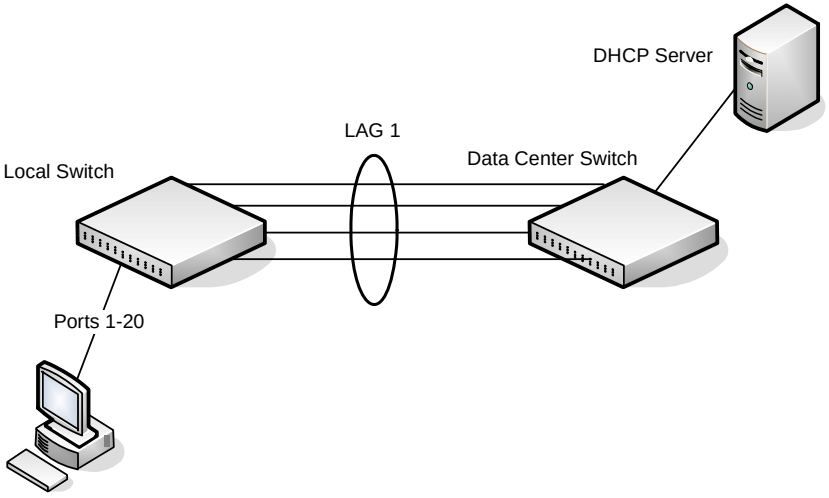
This section contains the following examples:

- Configuring DHCP Snooping
- Configuring IPSG

Configuring DHCP Snooping

In this example, DHCP snooping is enabled on VLAN 100. Ports 1-20 connect end users to the network and are members of VLAN 100. These ports are configured to limit the maximum number of DHCP packets with a rate limit of 100 packets per second. LAG 1, which is also a member of VLAN 100 and contains ports 21-24, is the trunk port that connects the switch to the data center, so it is configured as a trusted port.

Figure 27-25. DHCP Snooping Configuration Topology



The commands in this example also enforce rate limiting and remote storage of the bindings database. The switch has a limited amount of storage space in NVRAM and flash memory, so the administrator specifies that the DHCP snooping bindings database is stored on an external TFTP server.

To configure the switch:

1. Enable DHCP snooping on VLAN 100.

```
console#config
console(config)#ip dhcp snooping vlan 100
```

2. Configure LAG 1, which includes ports 21-24, as a trusted port. All other interfaces are untrusted by default.

```
console(config)#interface port-channel 1
console(config-if-Po1)#ip dhcp snooping trust
console(config-if-Po1)#exit
```

3. Enter interface configuration mode for all untrusted interfaces (ports 1-20) and limit the number of DHCP packets that an interface can receive to 100 packets per second. LAG 1 is a trusted port and keeps the default value for rate limiting (unlimited).

```
console(config)#interface range te1/0/1-20
console(config-if)#ip dhcp snooping limit rate 100
console(config-if)#exit
```

4. Specify that the DHCP snooping database is to be stored remotely in a file called dsDb.txt on a TFTP server with and IP address of 10.131.11.1.

```
console(config)#ip dhcp snooping database
tftp://10.131.11.1/dsDb.txt
console(config)#exit
```

5. Enable DHCP snooping for the switch

```
console(config)#ip dhcp snooping
```

6. View DHCP snooping information.

```
console#show ip dhcp snooping

DHCP snooping is Enabled
DHCP snooping source MAC verification is enabled
DHCP snooping is enabled on the following VLANs:
100

 Interface    Trusted     Log Invalid Pkts
-----------  ----------   ----------------
```

Configuring IPSG

This example builds on the previous example and uses the same topology shown in Figure 27-25. In this configuration example, IP source guard is enabled on ports 1-20. DHCP snooping must also be enabled on these ports. Additionally, because the ports use IP source guard with source IP and MAC address filtering, port security must be enabled on the ports as well.

To configure the switch:

1. Enter interface configuration mode for the host ports and enable IPSG.

```
console(config)#interface range te1/0/1-20
console(config-if)#ip verify source port-security
```

2. Enable port security on the ports.

```
console(config-if)#port security
```

3. View IPSG information.

```
console#show ip verify source

Interface   Filter    IP Address     MAC Address         Vlan
---------  -------    ----------
Te1/0/1    ip-mac     192.168.3.45   00:1C:23:55:D4:8E   100
Te1/0/2    ip-mac     192.168.3.40   00:1C:23:12:44:B6   100
Te1/0/3    ip-mac     192.168.3.33   00:1C:23:AA:B8:01   100
Te1/0/4    ip-mac     192.168.3.18   00:1C:23:67:D3:CC   100
Te1/0/5    ip-mac     192.168.3.49   00:1C:23:55:1B:6E   100
--More-- or (q)uit
```

28

Configuring Link Aggregation

This chapter describes how to create and configure link aggregation groups (LAGs), which are also known as port channels.

The topics covered in this chapter include:

- Link Aggregation Overview
- Default Link Aggregation Values
- Configuring Link Aggregation (Web)
- Configuring Link Aggregation (CLI)
- Link Aggregation Configuration Examples

Link Aggregation Overview

Link Aggregation allows one or more full-duplex (FDX) Ethernet links of the same speed to be aggregated together to form a LAG. This allows the switch to treat the LAG as if it is a single link.

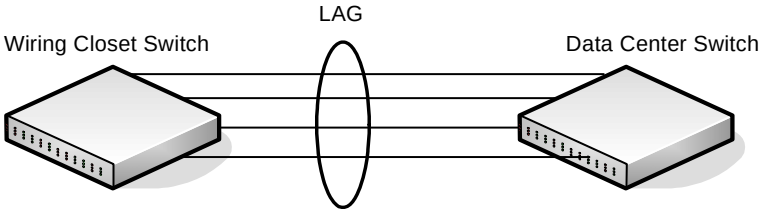
The PowerConnect 8000-series and 8100-series switches support industry-standard LAGs that adhere to the IEEE 802.3ad specification. The maximum number of LAGs that may be configured is limited to the maximum number of ports possible in the switch stack or stand-alone switch divided by two. This allows for a flexible configuration of LAGs where LAGs may have up to eight ports or as few as two ports. You can configure LAGs until all ports in the system are assigned to a LAG.

Assignment of interfaces to dynamic LAGs is based on a maximum of 144 interfaces assigned to dynamic LAGs, a maximum of 72 dynamic LAGs and a maximum of 8 interfaces per dynamic LAG. For example, 72 LAGs may be assigned 2 interfaces each, or 18 LAGs may be assigned 8 interfaces each.

Each PowerConnect 8000-series and 8100-series switches supports a maximum of 7 static or dynamic LAGs. Each LAG can consist of up to eight 10 Gbps ports. When eight 10 Gbps ports are configured as a LAG, the maximum bandwidth for the single, logical interface is 80 Gbps.

Figure 28-1 shows an example of a switch in the wiring closet connected to a switch in the data center by a LAG that consists of four physical 10 Gbps links. The LAG provides full-duplex bandwidth of 40 Gbps between the two switches.

Figure 28-1. LAG Configuration



Why Are Link Aggregation Groups Necessary?

The primary purpose of LAGs is to increase the overall bandwidth between two switches. This is accomplished by effectively aggregating multiple ports together that act as a single, logical connection between the two switches.

LAGs also provide redundancy. If a link fails, traffic is automatically redistributed across the remaining links.

What Is the Difference Between Static and Dynamic Link Aggregation?

Link aggregation can be configured as either dynamic or static. Dynamic configuration is supported using the IEEE 802.3ad standard, which is known as Link Aggregation Control Protocol (LACP). Static configuration is used when connecting a PowerConnect 8000-series and 8100-series switches to an external Gigabit Ethernet switch that does not support LACP.

One advantage of LACP is that the protocol enables the switch to confirm that the external switch is also configured for link aggregation. When using static configuration, a cabling or configuration mistake involving the 8000-series and 8100-series switch or the external switch could go undetected and thus cause undesirable network behavior. Both static and dynamic LAGs (via LACP) can detect physical link failures within the LAG and continue forwarding traffic through the other connected links within that same LAG. LACP can also detect switch or port failures that do not result in loss of link.

This provides a more resilient LAG. Best practices suggest using dynamic link aggregation instead of static link aggregation.When a port is added to a LAG as a static member, it neither transmits nor receives LACP PDUs.

What is LAG Hashing?

PowerConnect 8000-series and 8100-series switches support configuration of hashing algorithms for each LAG interface. The hashing algorithm is used to distribute traffic load among the physical ports of the LAG while preserving the per-flow packet order.

The hashing algorithm uses various packet attributes to determine the outgoing physical port.

The switch supports the following set of packet attributes to be used for hash computation:

- Source MAC, VLAN, EtherType, and incoming port.
- Destination MAC, VLAN, EtherType, and incoming port.
- Source IP and Source TCP/UDP port numbers.
- Destination IP and Destination TCP/UDP port numbers.
- Source/Destination MAC, VLAN, EtherType, and incoming port.
- Source/Destination IP and Source/Destination TCP/UDP port numbers.
- Enhanced hashing mode

Enhanced hashing mode has following advantages:

- MODULO-N operation based on the number of ports in the LAG.
- Packet attributes selection based on the packet type. For L2 packets, Source and Destination MAC address are used for hash computation. For IP packets, Source IP, Destination IP address, TCP/UDP ports are used.
- Non-Unicast traffic and Unicast traffic is hashed using a common hash algorithm.
- Excellent load balancing performance.

How Do LAGs Interact with Other Features?

From a system perspective, a LAG is treated just as a physical port, with the same configuration parameters for administrative enable/disable, spanning tree port priority, path cost as may be for any other physical port.

VLAN

When members are added to a LAG, they are removed from all existing VLAN membership. When members are removed from a LAG they are added back to the VLANs that they were previously members of as per the configuration file. Note that a port's VLAN membership can still be configured when it's a member of a LAG. However this configuration is only actually applied when the port leaves the LAG.

The LAG interface can be a member of a VLAN complying with IEEE 802.1Q.

STP

Spanning tree does not maintain state for members of a LAG, but the Spanning Tree does maintain state for the LAG interface. As far as STP is concerned, members of a LAG do not exist. (Internally, the STP state of the LAG interface is replicated for the member links.)

When members are deleted from a LAG they become normal links, and spanning tree maintains their state information.

Statistics

Statistics are maintained for all LAG interfaces as they are done for the physical ports, besides statistics maintained for individual members as per the 802.3ad MIB statistics.

LAG Configuration Guidelines

Ports to be aggregated must be configured so that they are compatible with the link aggregation feature and with the partner switch to which they connect.

Ports to be added to a LAG must meet the following requirements:

- Interface must be a physical Ethernet link.
- Each member of the LAG must be running at the same speed and must be in full duplex mode.
- The port cannot be a mirrored port

The following are the interface restrictions

- The configured speed of a LAG member cannot be changed.
- An interface can be a member of only one LAG.

Default Link Aggregation Values

The LAGs on the switch are created by default, but no ports are members. Table 28-1 summarizes the default values for the MAC address table.

Table 28-1. MAC Address Table Defaults

Parameter	Default Value
LACP system priority	1
LACP port priority	1
LACP timeout	Long
LAG hash algorithm type	Source IP and source TCP/UDP port

Configuring Link Aggregation (Web)

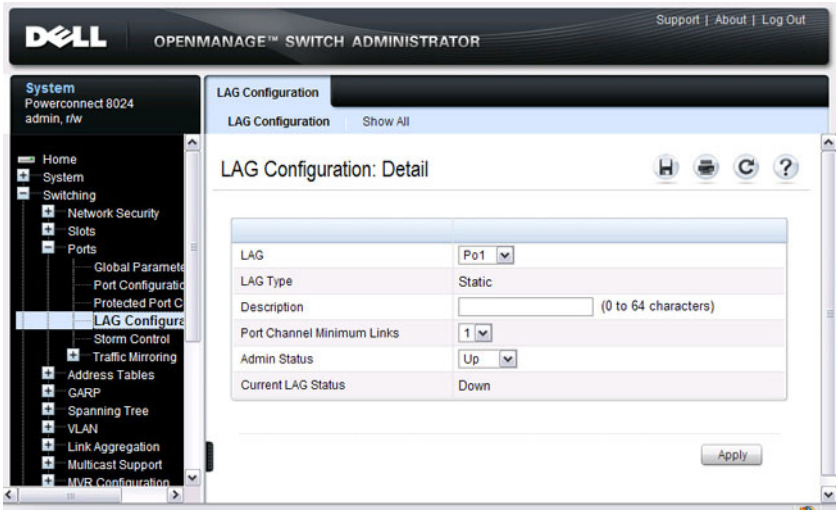
This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring LAGs on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

LAG Configuration

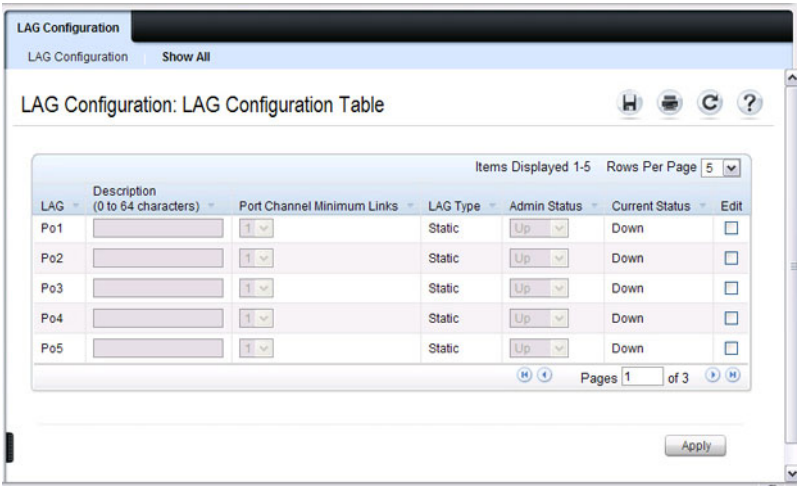
Use the **LAG Configuration** page to set the name and administrative status (up/down) of a LAG.

To display the **LAG Configuration** page, click **Switching** → **Ports** → **LAG Configuration** in the navigation panel.

Figure 28-2. LAG Configuration



To view or edit settings for multiple LAGs, click **Show All**.

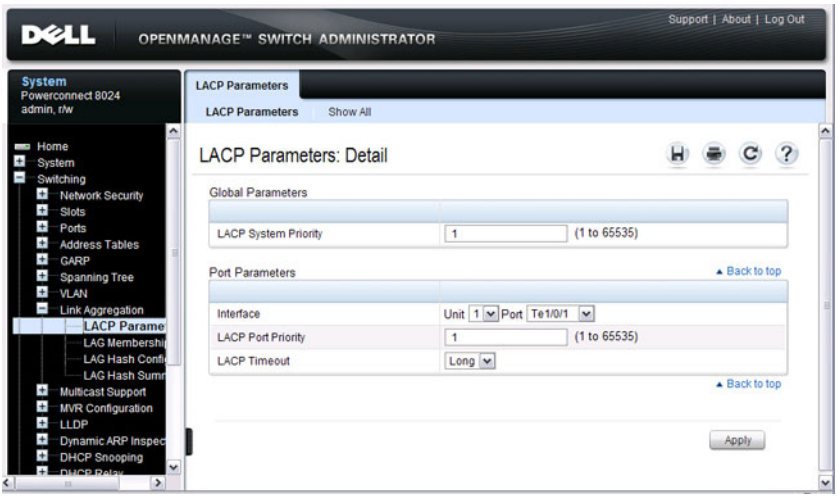


LACP Parameters

Dynamic link aggregation is initiated and maintained by the periodic exchanges of LACP PDUs. Use the **LACP Parameters** page to configure LACP LAGs.

To display the **LACP Parameters** page, click **Switching** → **Link Aggregation** → **LACP Parameters** in the navigation panel.

Figure 28-3. LACP Parameters



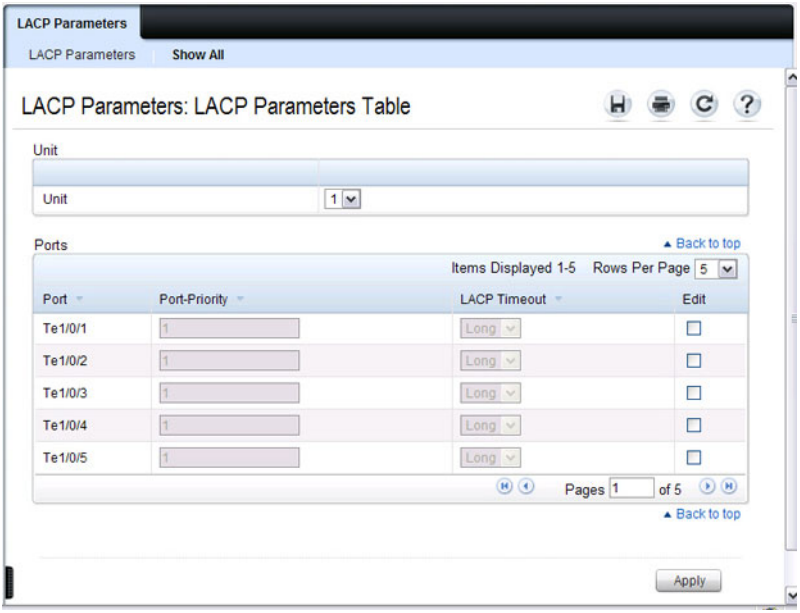
Configuring LACP Parameters for Multiple Ports

To configure LACP settings:

1. Open the **LACP Parameters** page.
2. Click **Show All**.

   The **LACP Parameters Table** page displays.

Figure 28-4. LACP Parameters Table



3 Select the **Edit** check box associated with each port to configure.

4 Specify the LACP port priority and LACP timeout for each port.

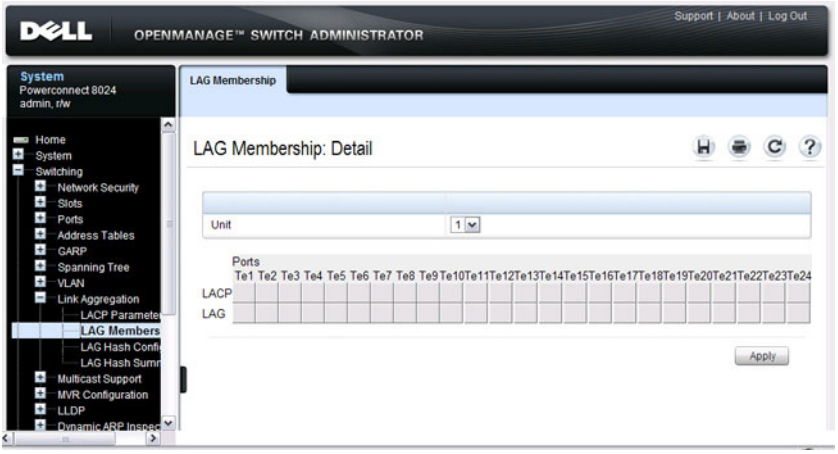
5 Click **Apply**.

LAG Membership

Your switch supports 48 LAGs per system, and eight ports per LAG. Use the **LAG Membership** page to assign ports to static and dynamic LAGs.

To display the **LAG Membership** page, click **Switching** → **Link Aggregation** → **LAG Membership** in the navigation panel.

Figure 28-5. LAG Membership



Adding a Port to a Static LAG

To add a static LAG member:

1. Open the **LAG Membership** page.
2. Click in the **LAG** row to toggle the port to the desired LAG.

   The LAG number displays for that port. The LAG number increases each time you click until the number reaches the maximum LAG number and then returns to blank (no LAG assigned).
3. Click **Apply**.

   The port is assigned to the selected LAG, and the device is updated.

Adding a LAG Port to a Dynamic LAG by Using LACP

To add a dynamic LAG member:

1. Open the **LAG Membership** page.
2. Click in the **LACP** row to toggle the desired LAG port to **L**.

   NOTE: The port must be assigned to a LAG before it can be aggregated to an LACP.

3. Click **Apply**.

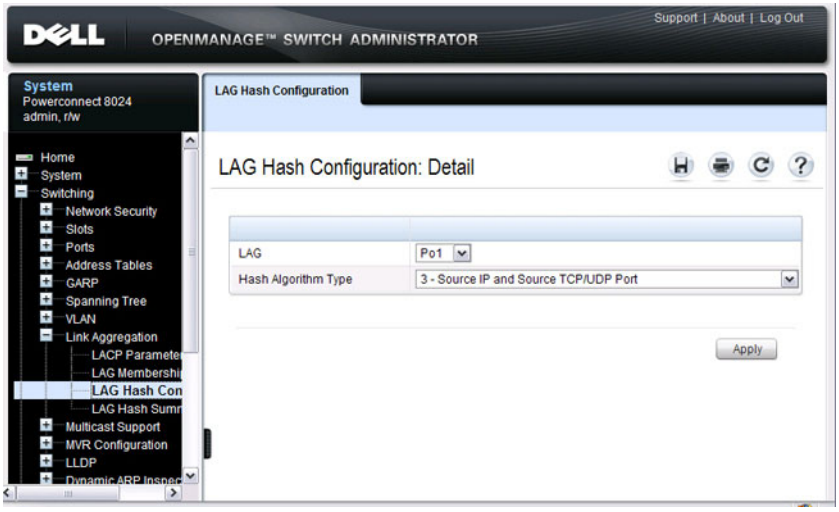
   The LAG port is added as a dynamic LAG member to the selected LAG.

LAG Hash Configuration

Use the LAG hash algorithm to set the traffic distribution mode on the LAG. You can set the hash type for each LAG.

To display the **LAG Hash Configuration** page, click **Switching** → **Link Aggregation** → **LAG Hash Configuration** in the navigation panel.

Figure 28-6. LAG Hash Configuration

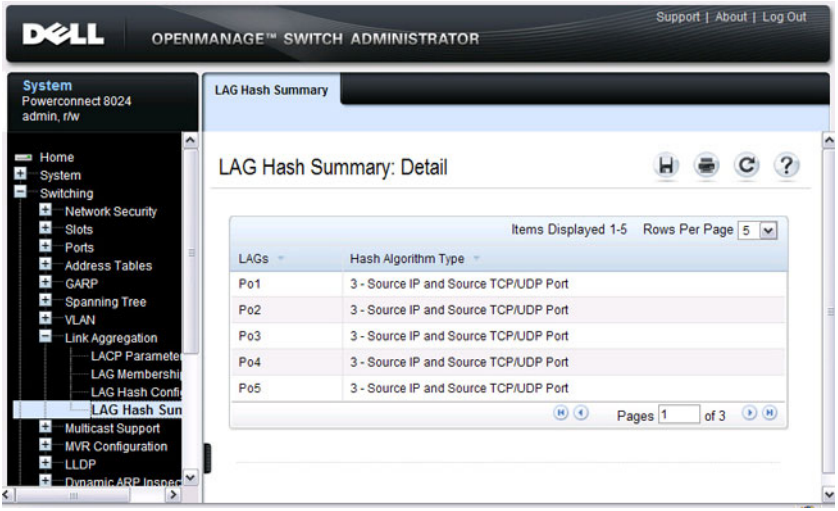


LAG Hash Summary

The **LAG Hash Summary** page lists the channels on the system and their assigned hash algorithm type.

To display the **LAG Hash Summary** page, click **Switching** → **Link Aggregation** → **LAG Hash Summary** in the navigation panel.

Figure 28-7. LAG Hash Summary



Configuring Link Aggregation (CLI)

This section provides information about the commands you use to configure link aggregation settings on the switch. For more information about the commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Configuring LAG Characteristics

Beginning in Privileged EXEC mode, use the following commands to configure a few of the available LAG characteristics. Many of the commands described in "Configuring Port Characteristics (CLI)" on page 451 are also applicable to LAGs.

Command	Purpose
configure	Enter global configuration mode.
interface *interface*	Enter interface configuration mode for the specified LAG. The *interface* variable includes the interface type, which is **port-channel**, and the LAG number, for example **port-channel 3**.
	You can also specify a range of LAGs with the **interface range port-channel** command, for example, **interface range port-channel 3-6** configures LAGs 3, 4, 5, and 6.
description *description*	Configure a description for the LAG or range of LAGs
port-channel min-links *minimum*	Set the minimum number of links that must be up in order for the port channel interface to be declared up.
exit	Exit to Global Config mode.
CTRL + Z	Exit to Privileged EXEC mode.
show interfaces description port-channel *port-channel number*	View the configured description for the specified LAG.
show interfaces port-channel [*port-channel number*]	View LAG information for the specified LAG or for all LAGs.

Configuring Link Aggregation Groups

Beginning in Privileged EXEC mode, use the following commands to add ports as LAG members and to configure the LAG hashing mode.

Command	Purpose
configure	Enter global configuration mode.
interface *interface*	Enter interface configuration mode for the specified port. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**. You can also specify a range of ports with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12.
channel-group *port-channel-number* **mode** {**on** | **auto**}	Add the port(s) to the LAG specified with the *port-channel-number* value. Use the auto keyword to add the port(s) as dynamic members, or use **on** to specify that the LAG membership is static. • *port-channel-number* — Number of a valid port-channel for the current port to join. • **on** — Forces the port to join a channel without LACP (static LAG). • **active** — Forces the port to join a channel with LACP (dynamic LAG).
exit	Exit to Global Config mode.
interface port-channel *number*	Enter interface configuration mode for the specified LAG. You can also specify a range of LAGs to configure with the **interface range port-channel** command, for example, **interface range port-channel 1-3,10** configures LAGs 1, 2, 3, and 10.

Command	Purpose
hashing-mode *mode*	Set the hashing algorithm on the LAG. The *mode* value is a number from 1 to 7. The numbers correspond to the following algorithms: • 1 — Source MAC, VLAN, EtherType, source module, and port ID • 2 — Destination MAC, VLAN, EtherType, source module, and port ID • 3 — Source IP and source TCP/UDP port • 4 — Destination IP and destination TCP/UDP port • 5 — Source/destination MAC, VLAN, EtherType, and source MODID/port • 6 — Source/destination IP and source/destination TCP/UDP port • 7 — Enhanced hashing mode
CTRL + Z	Exit to Privileged EXEC mode.
show interfaces port-channel [*port-channel number*]	View LAG information for the specified LAG or for all LAGs.
show statistics port-channel *port-channel-number*	View interface statistics for the specified LAG.

Configuring LACP Parameters

Beginning in Privileged EXEC mode, use the following commands to configure system and per-port LACP parameters.

Command	Purpose
configure	Enter global configuration mode.
lacp system-priority *value*	Set the Link Aggregation Control Protocol priority for the switch. the priority value range is 1–65535.
interface port-channel *number*	Enter interface configuration mode for the specified LAG. You can also specify a range of LAGs to configure with the **interface range port-channel** command, for example, **interface range port-channel 1-3,10** configures LAGs 1, 2, 3, and 10.
lacp port-priority *value*	Set the Link Aggregation Control Protocol priority for the port or range of ports. The priority value range is 1–65535.
lacp timeout {long | short}	Specify whether to wait a long or short time between LACP PDU transmissions.
exit	Exit to Privileged EXEC mode.
show lacp *interface*	View LACP parameters for an Ethernet interface or a LAG. The *interface* parameter includes the interface type (**tengigabitethernet** or **port-channel**) and number.

Link Aggregation Configuration Examples

This section contains the following examples:

- Configuring Dynamic LAGs
- Configuring Static LAGs

NOTE: The examples in this section show the configuration of only one switch. Because LAGs involve physical links between two switches, the LAG settings and member ports must be configured on both switches.

Configuring Dynamic LAGs

The commands in this example show how to configure a static LAG on a switch. The LAG number is 1, and the member ports are 1, 2, 3, 6, and 7.

To configure the switch:

1. Enter interface configuration mode for the ports that are to be configured as LAG members.

   ```
   console(config)#interface range te1/0/1-3,te1/0/6-7
   ```

2. Add the ports to LAG 2 with LACP.

   ```
   console(config-if)#channel-group 1 mode active
   ```

3. View information about LAG 1.

   ```
   console#show interfaces port-channel 1

   Channel  Ports              Hash Algorithm  Ch-Type  min-links
   -------  -----------------  --------------  -------  ---------
   Po1      Inactive: Te1/0/1, 3               Dynamic  1
            Te1/0/2, Te1/0/3,
            Te1/0/6, Te1/0/7
   ```

Configuring Static LAGs

The commands in this example show how to configure a static LAG on a switch. The LAG number is 2, and the member ports are 10, 11, 14, and 17.

To configure the switch:

1. Enter interface configuration mode for the ports that are to be configured as LAG members.

   ```
   console(config)#interface range te1/0/10-12,
   te1/0/14,te1/0/17
   ```

2. Add the ports to LAG 2 without LACP.

   ```
   console(config-if)#channel-group 2 mode on
   ```

3. View information about LAG 2.

   ```
   console#show interfaces port-channel 2

   Channel  Ports               Hash Algorithm  Ch-Type  min-links
   -------  -----------------   --------------  -------  ---------
   Po2      Inactive: Te1/0/10, 3               Static   1
            Te1/0/11, Te1/0/12,
            Te1/0/14, Te1/0/17
   ```

29

Configuring Data Center Bridging Features

This chapter describes how to manage the features developed for use in data center environments but often used in a variety of 10G applications.

The topics covered in this chapter include:

- Data Center Bridging Technology Overview
- Priority Flow Control (PFC)
- DCB Capability Exchange (DCBX)
- FIP Snooping
- Enhanced Transmission Selection (ETS)

Data Center Bridging Technology Overview

The PowerConnect 8024/8024F switches support Data Center Bridging (DCB) features to increase the reliability of Ethernet-based networks in the data center.

The PC81xx switches support PFC, ETS, and DCBX capability exchange, with the ability to autoconfigure from a peer switch. The PC80xx switches support FIP Snooping, DCBX capability exchanges and ETS proxy, and the ability to autoconfigure from a peer switch. The PC80xx also can be manually configured to support bandwidth sharing among traffic classes.

The Ethernet enhancements that DCB provides are well suited for Fibre Channel over Ethernet (FCoE) environments and iSCSI applications.

Table 29-1 provides a summary of the features this chapter describes.

Table 29-1. Data Center Features

Feature	Description
PFC	Provides a way to distinguish which traffic on a physical link is paused when congestion occurs based on the priority of the traffic.

Table 29-1. Data Center Features (Continued)

Feature	Description
DCBx	Allows DCB devices to exchange configuration information, using type-length-value (TLV) information elements over LLDP, with directly connected peers.
FIP Snooping	Inspects and monitors FIP frames and applies policies based upon the L2 header information in those frames
ETS	Supports the ETS configuration and Application Priority TLVs, which are accepted from auto-upstream devices and propagated to auto-downstream devices. The 8024 switch only transports the ETS TLVs and does not configure itself with received ETS information. ETS configuration parameters may be entered manually on the PC8024. The PC8132/PC8164 switches support the automatic configuration of ETS parameters.

Default DCB Values

Table 29-2 lists the default values for the DCB features that this chapter describes.

Table 29-2. Default Port-Based Traffic Control Values

Feature	Default
PFC	Disabled, no priority classifications are configured.
DCBx version	Auto detect
FIP snooping	Disabled globally and on all VLANs
FC map value	0x0efc00
FIP snooping port mode	Host facing (not FCF facing)

Priority Flow Control

Ordinarily, when flow control is enabled on a physical link, it applies to all traffic on the link. When congestion occurs, the hardware sends pause frames that temporarily suspend traffic flow to help prevent buffer overflow and dropped frames.

PFC provides a means of pausing individual priorities within a single physical link. By pausing the congested priority or priorities independently, protocols that are highly loss-sensitive can share the same link with traffic that has different loss tolerances.

This feature is used in networks where the traffic has differing loss tolerances. For example, Fibre Channel traffic is highly sensitive to traffic loss. If a link contains both loss-sensitive data and other less loss-sensitive data, the loss-sensitive data should use a no-drop priority that is enabled for flow control.

Priorities are differentiated by the priority field of the IEEE 802.1Q VLAN header, which identifies an IEEE 802.1p priority value. These priority values must be mapped to internal class-of-service (CoS) values.

The PFC feature allows you to specify the CoS values that should be paused (due to greater loss sensitivity) instead of dropped when congestion occurs on a link. Unless configured as no-drop, all CoS priorities are considered non-pausable ("drop") when priority-based flow control is enabled until no-drop is specifically turned on.

PFC Operation and Behavior

PFC uses a control packet newly defined in IEEE 802.1Qbb and, therefore, is not compatible with IEEE 802.3 Annex 31B flow control. An interface that is configured for PFC is automatically disabled for flow control. When PFC is disabled on an interface, the flow control configuration for the interface becomes active. Any flow control frames received on a PFC configured interface are ignored.

Each priority is configured as either *drop* or *no-drop*. If a priority that is designated as no-drop is congested, the priority is paused. Drop priorities do not participate in pause. You must configure the same no-drop priorities across the network in order to ensure end-to-end lossless behavior.

Operator configuration of PFC is used only when the port is configured in a manual role. When interoperating with other equipment in a manual role, the peer equipment must be configured with identical PFC priorities and VLAN assignments. Interfaces not enabled for PFC ignore received PFC frames. Ports configured in auto-upstream or auto-downstream roles receive their PFC configuration from the configuration source and ignore any manually configured information.

NOTE: This feature is configurable on physical full duplex interfaces only. To enable PFC on a LAG interface, the member interfaces must have the same configuration.

When PFC is disabled, the interface defaults to the IEEE 802.3 flow control setting for the interface. PFC is disabled by default.

If you enable priority-based flow control for a particular priority value on an interface, ensure that Voice VLAN tagging is enabled on the interface so that the 802.1p priority values are carried through the network (see "Voice VLAN" on page 545). Additionally, make sure that 802.1p priority values are mapped to CoS values (see "Configuring Class-of-Service" on page 1133).

PFC can be configured using the web interface and the command line interface.

Configuring PFC Using the Web Interface

This section provides information about the OpenManage Switch Administrator pages to use to view and configure PFC on PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

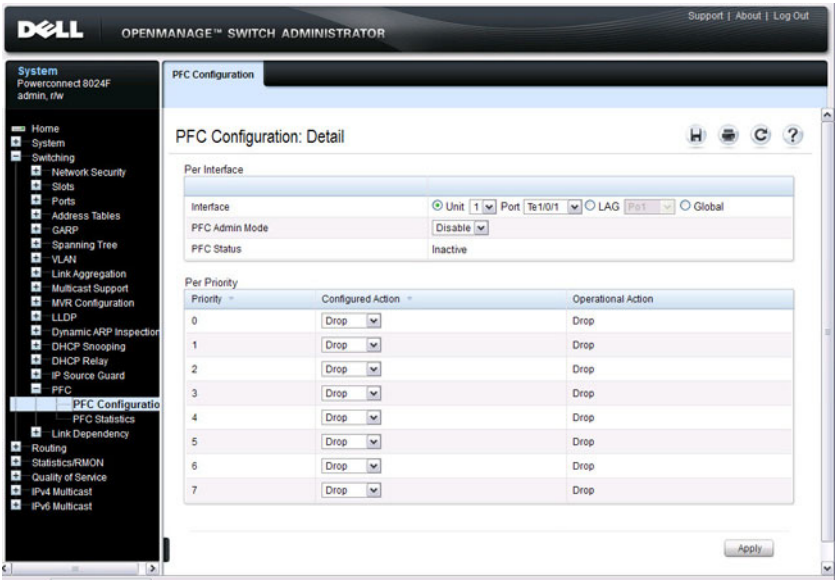
NOTE: FIP snooping can be enabled, configured, and monitored only by using the CLI.

PFC Configuration Page

Use the PFC Configuration page to enable priority flow control on one or more interfaces and to configure which priorities are subject to being paused to prevent data loss.

To display the **PFC Configuration** page, click **Switching → PFC → PFC Configuration** in the navigation menu.

Figure 29-1. PFC Configuration

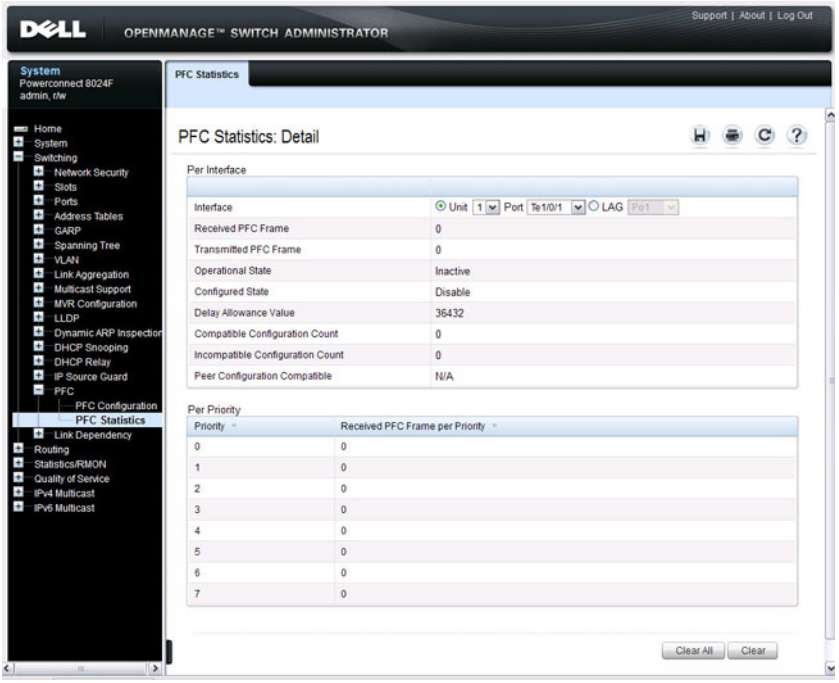


PFC Statistics Page

Use the **PFC Statistics** page to view the PFC statistics for interfaces on the switch.

To display the **PFC Statistics** page, click **Switching** → **PFC** → **PFC Statistics** in the navigation menu.

Figure 29-2. PFC Statistics



Configuring PFC Using the CLI

Beginning in Privileged EXEC mode, use the following commands to configure PFC.

NOTE: If DCBx is enabled and the switch is set to autoconfigure from a DCBX peer, configuring PFC is not necessary because the DCBx protocol automatically configures the PFC parameters.

Command	Purpose
configure	Enter global configuration mode.

Command	Purpose
interface *interface*	Enter interface configuration mode for the specified interface. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**. You can also specify a range of interfaces with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12.
datacenter-bridging	Enter the Data Center Bridging mode. PFC commands are issued from within this mode.
priority-flow-control mode on	Enable PFC on the interface(s) **NOTE:** It is unnecessary to set the priority flow control to enable if the **lldp dcbx port-role auto-down** or **lldp dcbx port-role auto-up** command has already been applied.
priority-flow-control priority *priority-id* {**drop** | **no-drop**}	Use the **no-drop** option to enable the priority group for lossless behavior. To enable lossy behavior, use the **drop** form of the command. *priority-id* — Specify the IEEE 802.1p priority value (range: 0–7) **NOTE:** Only two queues can be set to no-drop at one time.
CTRL + Z	Exit to Privileged EXEC mode.
show interfaces datacenter-bridging [*interface* | **port-channel** *port-channel-id*]	Display the datacenter-bridging configuration, status and counters for a given interface.
clear priority-flow-control statistics [*interface* | **port-channel** *port-channel-id*]	Clear all PFC statistics or the PFC statistics for the specified interface.

PFC Configuration Example

The network in this example handles both data and voice traffic. Because the voice traffic is time sensitive, it requires a higher priority than standard data traffic. The voice traffic uses VLAN 100 and has an 802.1p priority of 5, which is mapped to hardware queue 4. IP phones are connected to ports 3, 5, and 10, so PFC is enabled on these ports with 802.1p priority 5 traffic as no-drop. The configuration also enables VLAN tagging so that the 802.1p priority is identified. This example assumes the voice VLAN (VLAN 100) has already been configured.

△ **CAUTION:** All ports may be briefly shutdown when modifying either flow control (FC) or PFC settings. PFC uses a control packet defined in 802.1Qbb and is not compatible with 802.3x FC.

1. Map 802.1p priority 5 to traffic class 4. The following command changes the priority to traffic class mapping to be one-to-one, based upon the default switch settings. For lossless service, a priority must be mapped one-to-one to a traffic class. For more information about traffic classes, see "Configuring Class-of-Service" on page 1133.

```
console#configure
console(config)#classofservice dot1p-mapping 5 4
```

2. Enter Interface Configuration mode for ports 3, 5, and 10, and then enter Data Center Bridging mode for these ports.

```
console(config)#interface range te1/0/3, te1/0/5,
   te1/0/10
console(config-if)#datacenter-bridging
```

3. Enable PFC and configure traffic marked with 802.1p priority 5 to be paused rather than dropped when congestion occurs.

```
console(config-dcb)#priority-flow-control mode on
console(config-dcb)#priority-flow-control
   priority 5 no-drop
console(config-dcb)#exit
```

4. Enable VLAN tagging on the ports so the 802.1p priority is identified.

```
console(config-if)#switchport mode general
console(config-if)#switchport general allowed vlan
   add 100 tagged
console(config-if)#exit
```

DCB Capability Exchange

The Data Center Bridging Exchange Protocol (DCBx) is used by DCB devices to exchange configuration information with directly connected peers. DCBx uses type-length-value (TLV) information elements over LLDP to exchange information, so LLDP must be enabled on the port to enable the information exchange. By default, LLDP is enabled on all ports. For more information, see "Discovering Network Devices" on page 637.

The main objective of DCBx is to perform the following operations:

- **Discovery of DCB capability in a peer**: DCBx is used to learn about the capabilities of the peer device. It is a means to determine if the peer device supports a particular feature such as PFC.
- **DCB feature misconfiguration detection**: DCBx can be used to detect misconfiguration of a feature between the peers on a link. Misconfiguration detection is feature-specific because some features may allow asymmetric configuration.
- **Peer configuration of DCB features**: DCBx can be used by a device to perform configuration of DCB features in its peer device if the peer device is willing to accept configuration.

DCBx is expected to be deployed in Fibre Channel over Ethernet (FCoE) topologies in support of lossless operation for FCoE traffic. In these scenarios, all network elements are DCBx-enabled. In other words, DCBx is enabled end-to-end.

The DCBx protocol supports the propagation of configuration information for the following features:

- Enhanced Transmission Selection (ETS)
- Priority-based Flow Control (PFC)
- Application Priorities

These features use DCBx to send and receive device configuration and capability information to the peer DCBx device.

The Application Priorities information is simply captured from the peer and potentially propagated to other peers by the DCBx component, as well as being configured when iSCSI is enabled on an operationally active PFC port.

Interoperability with IEEE DCBx

To be interoperable with legacy industry implementations of the DCBx protocol, The PowerConnect 8024/8024F switches use a hybrid model to support both the IEEE version of DCBx (IEEE 802.1Qaz) and legacy DCBx versions.

The PowerConnect 8024/8024F switch automatically detects whether a peer is operating with either of the two CEE DCBx versions or the IEEE standard DCBx version (the default mode). You can also configure DCBx to manually select one of the legacy versions or IEEE standard mode. In auto-detect mode, the switch starts operating in IEEE DCBx mode on a port, and if it detects a legacy DCBx device based on the OUI of the organization TLV, then the switch changes its DCBx mode on that port to support the version detected. There is no timeout mechanism to move back to IEEE mode. Once the DCBx peer times out, multiple peers are detected, the link is reset (link down/up) or as commanded by the operator, DCBx resets its operational mode to IEEE.

The interaction between the DCBx component and other components remains the same irrespective of the operational mode it is executing. For instance, the DCBx component interacts with PFC to get needed information to pack the TLVs to be sent out on the interface. Based on the operational control mode of the port, DCBx packs it in the proper frame format.

DCBx and Port Roles

The behavior of each port is dependent on its operational mode and that of other ports in the stack. The port mode is a DCBx configuration item that is passed to the DCBx clients to control the processing of their configuration information. There are four port roles:

1. Manual
2. Auto-Upstream
3. Auto-Downstream
4. Configuration Source

Ports operating in the manual role do not have their configuration affected by peer devices or by internal propagation of configuration. These ports have their operational mode, traffic classes, and bandwidth information specified

explicitly by the operator. These ports advertise their configuration to their peer if DCBx is enabled on that port. Incompatible peer configurations are logged and counted with an error counter.

The default operating mode for each port is manual. A port that is set to manual mode sets the willing bit for DCBx client TLVs to false. Manually-configured ports never internally propagate or accept internal or external configuration from other ports; in other words, a manual configuration discards any automatic configuration. Manually-configured ports may notify the operator of incompatible configurations if client configuration exchange over DCBx is enabled. Manually configured ports are always operationally enabled for DCBx clients, regardless of whether DCBx is enabled. Operationally enabled means that the port reports that it is able to operate using the current configuration.

A port operating in the auto-upstream role advertises a configuration, but is also willing to accept a configuration from the link-partner and propagate it internally to the auto-downstream ports, as well as receive configuration propagated internally by other auto-upstream ports. Specifically, the willing parameter is enabled on the port and the recommendation TLV is sent to the peer and processed if received locally. The first auto-upstream port to successfully accept a compatible configuration becomes the configuration source. The configuration source propagates its configuration to other auto-upstream and auto-downstream ports. Only the configuration source may propagate configuration to other ports internally. Auto-upstream ports that receive internally propagated information ignore their local configuration and utilize the internally propagated information.

Peer configurations received on auto-upstream ports other than the configuration source result in one of two possibilities. If the configuration is compatible with the configuration source, then the DCBx client becomes operationally active on the upstream port. If the configuration is not compatible with the configuration source, then a message is logged indicating an incompatible configuration, an error counter is incremented, and the DCBx client is operationally disabled on the port. The expectation is that the network administrator configures the upstream devices appropriately so that all such devices advertise a compatible configuration.

A port operating in the auto-downstream role advertises a configuration but is not willing to accept one from the link partner. However, the port will accept a configuration propagated internally by the configuration source. Specifically,

the willing parameter is disabled on auto-downstream. By default, auto-downstream ports have the recommendation TLV parameter enabled. Auto-downstream ports that receive internally propagated information ignore their local configuration and utilize the internally propagated information. Auto-downstream ports propagate PFC, ETS, and application priority information received from the configuration source.

In the Configuration Source role, the port has been manually selected to be the configuration source. Configuration received over this port is propagated to the other auto configuration ports, however, no automatic election of a new configuration source port is allowed. Events that cause selection of a new configuration source are ignored. The configuration received over the configuration source port is maintained until cleared by the operator (set the port to the manual role).

Configuration Source Port Selection Process

When an auto-upstream or auto-downstream port receives a configuration from a peer, the DCBx client first checks if there is an active configuration source. If there is a configuration source already selected, the received configuration is checked against the local port operational values as received from the configuration source, and if compatible, the client marks the port as operationally enabled. If the configuration received from the peer is determined to not be compatible, a message is logged, an error counter is incremented and the DCBx clients become operationally disabled on the port. Operationally disabled means that FIP snooping and PFC will not operate over the port. The port continues to keep link up and exchanges DCBx packets. If a compatible configuration is later received, the DCBx clients will become operationally enabled.

If there is no configuration source, a port may elect itself as the configuration source on a first-come, first-serve basis from the set of eligible ports. A port is eligible to become the configuration source if the following conditions are true:

- No other port is the configuration source.
- The port role is auto-upstream.
- The port is enabled with link up and DCBx enabled.
- The port has negotiated a DCBx relationship with the partner.

- The switch is capable of supporting the received configuration values, either directly or by translating the values into an equivalent configuration.

Whether or not the peer configuration is compatible with the configured values is NOT considered.

The newly elected configuration source propagates DCBx client information to the other ports and is internally marked as being the port over which configuration has been received. Configuration changes received from the peer over the configuration source port are propagated to the other auto configuration ports. Ports receiving auto configuration information from the configuration source ignore their current settings and utilize the configuration source information.

When a configuration source is selected, all auto-upstream ports other than the configuration source are marked as willing disabled.

To reduce flapping of configuration information, if the configuration source port is disabled, disconnected, or loses LLDP connectivity, the system clears the selection of configuration source port (if not manually selected) and enables the willing bit on all auto-upstream ports. The configuration on the auto configuration ports is not cleared (configuration holdover). If the user wishes to clear the configuration on the system in this scenario, the user can put the configuration source port into manual mode.

When a new port is selected as the configuration source, it is marked as the configuration source, the DCBx configuration is refreshed on all auto configuration ports, and each port may begin configuration negotiation with their peer again (if any information has changed).

Disabling DCBX

If it is desired to disable DCBX, the network operator can use the following commands to eliminate the transmission of DCBX TLVs in the LLDP frames on an interface:

```
no lldp tlv-select dcbxp application-priority
no lldp tlv-select dcbxp congestion-notification
no lldp tlv-select dcbxp ets-config
no lldp tlv-select dcbxp ets-recommend
no lldp tlv-select dcbxp pfc
```

These commands eliminate only the DCBX TLVs from use by LLDP. They do not otherwise affect any manually configured DCBX capabilities or the normal operation of LLDP.

Configuring DCBx

You can use the CLI to configure DCBx.

Beginning in Privileged EXEC mode, use the following commands to configure DCBx.

Command	Purpose
configure	Enter global configuration mode.
lldp dcbx version {auto | cin | cee | ieee}	Optionally configure the administrative version for the DCBx protocol: • auto—Automatically select the version based on the peer response (default) • cin—Force the mode to Cisco-Intel-Nuova. (DCBx 1.0) • cee—Force the mode to CEE (DCBx 1.06) • ieee—Force the mode to IEEE 802.1Qaz
lldp tlv-select dcbxp [pfc | application-priority]	Enable LLDP to send specific DCBx TLVs if LLDP is enabled to transmit on the given interface. Entering the command with no parameters enables transmission of all TLVs. • pfc—Transmit the PFC configuration TLV • application-priority—Transmit the application priority TLV
interface *interface*	Enter interface configuration mode for the specified interface. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**. You can also specify a range of interfaces with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12.

Command	Purpose
lldp tlv-select dcbxp [pfc | application-priority]	Override the global configuration for the LLDP DCBx TLVs on this interface. Entering the command with no parameters enables transmission of all TLVs. • pfc—Transmit the PFC configuration TLV. • application-priority—Transmit the application priority TLV.
lldp dcbx port-role {auto-up |auto-down | manual | configuration-source}	Configure the DCBx port role on the interface: • **auto-up**—Advertises a configuration, but is also willing to accept a configuration from the link-partner and propagate it internally to the auto-downstream ports as well as receive configuration propagated internally by other auto-upstream ports. These ports have the willing bit enabled. These ports should be connected to FCFs. • **auto-down**—Advertises a configuration but is not willing to accept one from the link partner. However, the port will accept a configuration propagated internally by the configuration source. These ports have the willing bit set to disabled. Selection of a port based upon compatibility of the received configuration is suppressed. These ports should be connected to a trusted FCF. • **manual**—Ports operating in the Manual role do not have their configuration affected by peer devices or by internal propagation of configuration. These ports will advertise their configuration to their peer if DCBx is enabled on that port. The willing bit is set to disabled on manual role ports. • **configuration-source**—In this role, the port has been manually selected to be the configuration source. Configuration received over this port is propagated to the other auto configuration ports. Selection of a port based upon compatibility of the received configuration is suppressed. These ports should be connected to a trusted FCF. These ports have the willing bit enabled.
CTRL + Z	Exit to Privileged EXEC mode.

Command	Purpose
show lldp tlv-select interface {all | *interface*}	Display the interface TLV configuration for all interfaces or for the specified interface.
show lldp dcbx interface {all status | *interface* [detail]}	Display the interface TLV configuration for all interfaces or for the specified interface.

FIP Snooping

The FCoE Initialization Protocol (FIP) is used to perform the functions of FC_BB_E device discovery, initialization, and maintenance. FIP uses a separate EtherType from FCoE to distinguish discovery, initialization, and maintenance traffic from other FCoE traffic. FIP frames are standard Ethernet size (1518 Byte 802.1q frame), whereas FCoE frames are a maximum of 2240 bytes. FIP Snooping is supported on the PC80xx switches.

FIP snooping is a frame inspection method used by FIP Snooping Bridges to monitor FIP frames and apply policies based upon the L2 header information in those frames.

FIP snooping allows for:

- Auto configuration of Ethernet ACLs based on information in the Ethernet headers of FIP frames.
- Emulation of FC point-to-point links within the DCB Ethernet network.
- Enhanced FCoE security/robustness by preventing FCoE MAC spoofing.

The role of FIP snooping-enabled ports on the switch falls under one of the following types:

- Perimeter or Edge port (connected directly to a Fibre Channel end node or ENode).
- Fibre Channel forwarder (FCF) facing port (that receives traffic from FCFs targeted to the ENodes).

NOTE: The PowerConnect 8024/8024F FIP Snooping Bridge feature supports the configuration of the perimeter port role and FCF-facing port roles and is intended for use only at the edge of the switched network.

The default port role in an FCoE-enabled VLAN is as a perimeter port. FCF-facing ports are configured by the user.

Enabling and Disabling FIP Snooping

When FIP snooping is globally enabled on the switch, FC-BB-5 Annex D ACLs are installed on the switch, and FIP frames are snooped. FIP snooping will not allow FIP or FCoE frames to be forwarded over a port until the port is operationally enabled for PFC. VLAN tagging must be enabled on the interface in order to carry the dot1p values through the network.

When FIP snooping is disabled, received FIP frames are forwarded or flooded using the normal multicast rules.

NOTE: FIP Snooping will become operationally active on a port only when priority flow control is enabled and at least one lossless CoS queue is active.

Configuring the FC Map Value

When configured using fabric provided MAC addresses, FCoE devices transmit frames containing the FC map value in the upper 24 bits. Only frames that match the configured FC map value are passed across the VLAN. Frames with MAC addresses that do not match the FC map value are discarded.

Configuring Ports for FIP Snooping

To relay the FIP packets received from the hosts toward the Fibre Channel forwarders (FCFs), such as an FC router that has a direct FC link into storage, the switch needs to know the interfaces the FCFs are on. By default, an interface is configured to be a host facing interface, not an FCF facing interface.

Dell recommends that FCF-facing ports be placed into auto-upstream mode in order to receive DCBx information and propagate it to the Converged Network Adaptors (CNAs) on the downstream ports. Interfaces enabled for PFC should be configured in trunk or general mode and must be PFC operationally enabled before FCoE traffic can pass over the port.

Configuring FIP Snooping (CLI)

Beginning in Privileged EXEC mode, use the following commands to configure FIP snooping.

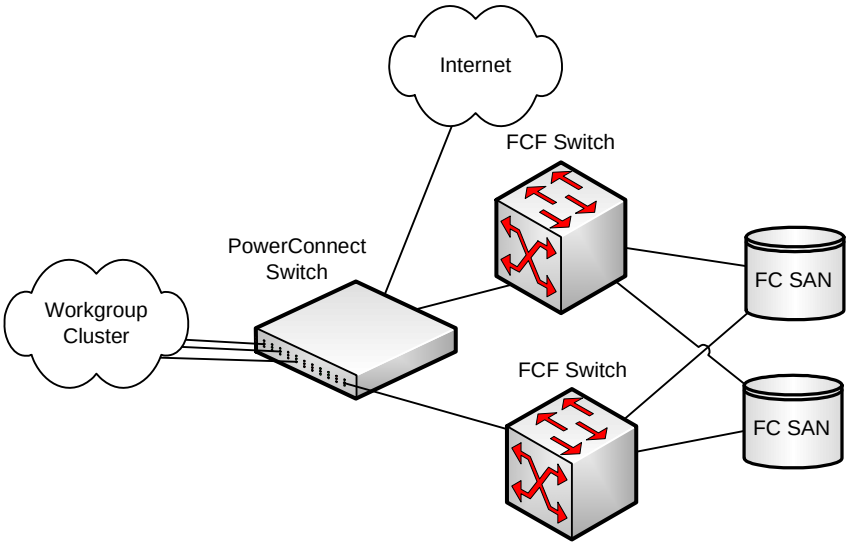
NOTE: FIP snooping will not allow FIP or FCoE frames to be forwarded over a port until the port is operationally enabled for PFC. VLAN tagging must be enabled on the interface in order to carry the dot1p values through the network. This section describes the FIP snooping commands only. For an example of configuring FIP snooping on the switch, see "FIP Snooping Configuration Example" on page 817.

Command	Purpose
configure	Enter global configuration mode.
feature fip-snooping	Globally enable FIP snooping on the switch.
vlan *vlan_id*	Enter VLAN configuration mode for a VLAN or range of VLANs.
fip-snooping enable	Enable the snooping of FIP packets on the specified VLAN or VLAN range.
fip-snooping fc-map *fc_map_value*	Optionally configure the FC-MAP value on a VLAN. The FC map value is used to help prevent the switch from being incorrectly configured. The range for *fc_map_value* is 0x0–0xffffff. The FC map value configured on the switch must match the FC map value configured on the FCF for the VLAN.
exit	Exit to global configuration mode.
interface *interface*	Enter interface configuration mode for the specified interface. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**. You can also specify a range of interfaces with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12.
fip-snooping port-mode fcf	Configure the interface to be connected to a FCF switch.
CTRL + Z	Exit to Privileged EXEC mode.
show fip-snooping sessions	Display information about the active FIP snooping sessions.

Command	Purpose
show fip-snooping fcf [*fcf-mac*]	Display information about the interfaces connected to Fibre Channel forwarder (FCF). Use the optional *fcf-mac* parameter to display additional information about the session with the specified FCF device.
show fip-snooping enode [*enode-mac*]	Display information about the interfaces connected to FCoE end nodes (ENodes). Use the optional *enode-mac* parameter to display FIP snooping sessions between the given FCF-MAC and ENode-MAC.
show fip-snooping statistics {**vlan** *vlan_id* | **interface** *interface*}	Display the statistics of the FIP packets snooped for all the FIP snooping-enabled VLANs or for the specified VLAN or interface.
clear fip-snooping statistics {**vlan** *vlan_id* | **interface** *interface*}	Clear the statistics of the FIP packets snooped for all the FIP snooping-enabled VLANs or for the specified VLAN or interface.

FIP Snooping Configuration Example

In this example, FIP snooping is deployed on a stack of two PowerConnect switches that are functioning as enterprise data center top of rack (ToR) switches. Two 10G ports on each stack member (ports 16 and 17 on each switch) connect to a pair FCoE-facing switches through redundant links. These ports are designated as default DCBx auto-upstream ports. The 10G ports that are directly connected to Converged Network Adapters (CNAs) on individual workstations (ports 1–3 on each switch) are designated as DCBx default auto-downstream ports.



To configure FIP snooping:

1. Enter global configuration mode and enable FIP snooping on the switch.

```
console#configure
console(config)#feature fip-snooping
```

2. Create VLAN 100. This command also enters the VLAN configuration mode for VLAN 100.

```
console(config)#vlan 100
console(config-vlan100)#fip-snooping enable
console(config-vlan100)#exit
```

3. Enter Interface Configuration mode for ports 1, 2, 3, 16, and 17 on both switches in the stack.

```
console(config)#interface range te1/0/1-
3,te1/0/16-17,te2/0/1-3,te2/0/16-17
```

4. Enable VLAN tagging to allow the ports to carry 802.1p priority values through the network.

```
console(config-if)#switchport mode general
console(config-if)#switchport general allowed vlan
add 100 tagged
```

5 Exit interface configuration mode for the range of interfaces.

6 Enter interface configuration mode for the CNA-facing ports and configure the DCBx port role as auto-downstream. This step automatically enables PFC on the ports.

```
console(config)#interface te1/0/1-3,te2/0/1-3
console(config-if)#lldp dcbx port-role auto-down
console(config-if#exit
```

7 Enter Interface Configuration mode for the ports connect to an FCF on both switches in the stack to configure the DCBx port role as auto-upstream.

```
console(config)#interface te1/0/16-17,te2/0/16-17
console(config-if)#lldp dcbx port-role auto-up
```

8 Set the FIP snooping port mode as *fcf* to indicate that these ports are connected to a Fibre Channel forwarder.

```
console(config-if)#fip-snooping port-mode fcf
console(config-if)#exit
```

9 Optionally, use the FIP snooping show commands to verify the configuration, view FIP snooping sessions, and view information about the ports that are connected to end nodes or FCFs.

Enhanced Transmission Selection

Networks classify and prioritize traffic to provide different service characteristics to end user traffic flows. Administrators may wish to guarantee or limit bandwidth for certain traffic, ensure lossless behavior for other traffic, and control the queue discipline/drop characteristics for best-effort traffic. Additionally, it is desirable for a switch to support sharing bandwidth among bursty sources while still enabling prioritization of time-sensitive or management traffic that requires minimum latency.

Enhanced transmission selection (ETS) provides uniform management for sharing bandwidth between congestion managed and traditional classes on a single bridged network. Using priority-based processing and weight allocations, Traffic Class Groups (TCGs) carrying different types of traffic such as LAN, SAN, and management traffic can be configured to provide minimum bandwidth guarantees, unused bandwidth sharing, and lossless or best-effort transmit characteristics. PowerConnect 81xx switches support

strict priority and Weighted Deficit Round Robin (WDRR) scheduling with up to two lossless traffic classes. WDRR schedules traffic based on average bandwidth consumed vs. frame counts.

ETS Operation

The normal (default) operation of PowerConnect switches, when uncongested, is that packets are scheduled for output in the order in which they are received, that is, using FIFO scheduling. The class of service (CoS) mechanism enables the administrator to schedule packets for output ahead of other packets when the switch is congested, choose which type of packets to drop when the switch is congested, and assign a minimum bandwidth guarantee to ensure scheduling fairness. These mechanisms operate at the CoS queue level; that is, the minimum bandwidth guarantees are made across all configured CoS queues.

NOTE: Minimum bandwidth guarantees and scheduling mechanisms apply only when the switch is congested. When the switch is not congested, packets egress the switch as soon as they are received.

ETS provides a second level of scheduling for packets selected for transmission by the CoS scheduler. ETS operates at the traffic class group (TCG) level and supports sharing of bandwidth across TCGs, bandwidth assignment for each TCG, and queue discipline (drop behavior) for each TCG. PC81xx switches support three TCGs internally, up to two of which may be configured as lossless.

When a packet arrives on an ingress port, it is forwarded to the appropriate CoS queue based upon the dot1p mapping. The dot1p mapping maps the user priority contained in the received VLAN Priority Tag to the traffic class or uses the default port assignment. A traffic class identifies a particular CoS queue. If the ingress port is configured to use 802.1p CoS mapping and the port is configured to trust the user priority value in the received frame, then the frame is forwarded to the appropriate CoS queue per its 802.1p user priority value.

At the first level of egress scheduling, each of the configured attributes of a CoS queue, namely scheduler algorithm, min-bandwidth and drop mechanism, are honored, and the packet is either dropped or forwarded to next level. Only frames selected by the first level scheduler are forwarded to the second level.

Strict priority traffic classes are serviced first in order of traffic class number. A strict priority traffic class is one that is configured as strict priority or has a traffic class group weight of 0 (unlimited bandwidth). Within a strict priority traffic class, frames are serviced (as determined by the traffic class number) in CoS queue priority order from the highest (6) to the lowest (0).

Each traffic class has an internal weight equal to the traffic class number plus one. After the strict priority traffic classes have been serviced, the remaining traffic classes are serviced according to their internal weight. For example, if CoS queues 0, 1, and 2 have an equal offered load toward a congested output port, CoS queue 2 will receive 3/6 of the bandwidth, CoS queue 1 will receive 2/6 of the bandwidth, and CoS queue 0 will receive 1/6 of the bandwidth.

The minimum bandwidth setting can be used to override the strict priority and weighted settings. The highest numbered strict priority queue will receive no more bandwidth than 100 percent minus the sum of the minimum bandwidth percentages assigned to the other queues. If used, it is recommended that minimum bandwidth percentages only be set high enough to ensure a minimum level of service for any queue; i.e., the sum of the minimum bandwidth percentages is a fraction of 100%. This ensures that the system can respond to bursts in traffic. Setting the minimum bandwidths such that they sum to 100% effectively sets the scheduler such that no queue can exceed its minimum bandwidth percentage when congestion occurs.



NOTE: CoS queue 7 is reserved for internal traffic. Non-strict priority CoS queues are serviced with WDRR scheduling using the bandwidth available after strict priority traffic is serviced.

Each CoS queue in the first level scheduler is mapped to one of the three traffic class groups in the second level scheduler. There, frames are serviced using the TCG configuration. The minimum bandwidth guarantee is first calculated across the TCGs. Strict priority TCGs are scheduled first but have their bandwidth reduced by the minimum bandwidth guarantees configured on other TCGs. Strict priority TCGs are scheduled from highest numbered TCG to lowest. When all TCGs have met their minimum bandwidth limits

(or the queues are empty), TCGs that have not met their maximum bandwidth limit are scheduled. Once the limits for a TCG are satisfied (maximum bandwidth, no frames available for transmission, etc.), the scheduler moves to the next TCG.

If no minimum or maximum bandwidth limits are configured, TCGs are serviced by the second-level scheduler using the configured TCG weights to define the relative bandwidth allocation among the TCGs. When an egress port is congested, packets are selected for discard using the configured tail-drop or WRED discipline. Minimum TCG bandwidth, maximum TCG bandwidth, and TCG weights are metered to within approximately 3% of the link bandwidth.

In the case that all TCGs are configured as strict priority, inter-TCG scheduling reverts to Weighted Deficit Round Robin scheduling based on the configured TCG weights.

Commands

This section provides information about the commands you use to configure and monitor ETS. For more information about the commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Command	Purpose
classofservice traffic-class-group	Maps the internal Traffic Class to an internal Traffic Class Group (TCG). The Traffic Class can range from 0-6, although the actual number of available traffic classes depends on the platform.
traffic-class-group max-bandwidth	Specifies the maximum transmission bandwidth limit for each TCG as a percentage of the interface rate. Also known as rate shaping, this has the effect of smoothing temporary traffic bursts over time so that the transmitted traffic rate is bound.
traffic-class-group min-bandwidth	Specifies the minimum transmission bandwidth guaranteed for each TCG before processing frames from lower numbered TCGs on an interface.
traffic-class-group strict	Activates the strict priority scheduler mode for each specified TCG.
traffic-class-group weight	Specifies the scheduling weight for each TCG. The scheduler attempts to balance the traffic selected for transmission from the TCGs such that, when the switch is congested, traffic is selected from the round robin configured TCGs in proportion to their weights.
show classofservice traffic-class-group	Displays the Traffic Class to Traffic Class Group mapping.
show interfaces traffic-class-group	**D**isplays the Traffic Class to Traffic Class Group mapping

ETS Configuration Example

This example configures four classes of traffic:

Best effort traffic	CoS Queue 0 for untagged and VLAN-tagged frames with VPTs 0, 1, and 2
Lossless FCoE/iSCSI traffic	CoS Queues 1 & 2 for VLAN tagged frames with VPTs 3 & 4 respectively
Expedited traffic	CoS Queue 3 on VLAN tagged frames with VPTs 5, 6, and 7

1. Enable Trust Mode on an Interface

The following command enables the use of the dot1p priority of the incoming packet. It may be configured on a single interface, a range of interfaces or all interfaces. By default, ports are configured to trust the incoming user priority.

```
console(config-if-Te1/0/1)#classofservice trust dot1p
```

To show the configured trust mode on an interface, use the following command:

```
console#show classofservice trust tengigabitethernet 1/0/1
Class of Service Trust Mode: Dot1P
```

2. Map Dot1p Priority to CoS Queues

This step maps dot1p priorities to the CoS queues. This mapping places the incoming packet in the selected CoS queue based on the dot1p priority. It may be configured on a single interface, a range of interfaces, or all interfaces.

To ensure lossless behavior, the dot1p priority must be mapped one to one to a CoS queue for the lossless priorities. Up to two lossless priorities may be configured on PC81xx switches. CoS queue 7 is reserved by the system and is not assignable. It is generally recommended that the administrator utilize CoS queues 0-3 as CoS queues 4-6 may be utilized by the system for other types of system traffic, e.g. routing protocol PDU handling. Frames with different user priorities assigned to a single CoS queue receive equal treatment.

This example maps user priorities 0, 1, and 2 to CoS queue 0 (background or best effort traffic), user priorities 3 and 4 to CoS queues 1 and 2 (FCoE and iSCSI traffic), and all other priorities to CoS queue 2 (low latency and network control traffic).

```
console(config-if-Te1/0/2)#classofservice dot1p-mapping 0 0
console(config-if-Te1/0/2)#classofservice dot1p-mapping 1 0
console(config-if-Te1/0/2)#classofservice dot1p-mapping 2 0
console(config-if-Te1/0/2)#classofservice dot1p-mapping 3 1
console(config-if-Te1/0/2)#classofservice dot1p-mapping 4 2
console(config-if-Te1/0/2)#classofservice dot1p-mapping 5 3
console(config-if-Te1/0/2)#classofservice dot1p-mapping 6 3
console(config-if-Te1/0/2)#classofservice dot1p-mapping 7 3
```

To show dot1p priority for an interface, use the following command:

```
console#show classofservice dot1p-mapping tengigabitethernet 1/0/1

 User Priority    Traffic Class
 -------------    -------------
       0                0
       1                0
       2                0
       3                1
       4                2
       5                3
       6                3
       7                3
```

3. Configure Minimum Bandwidth on the CoS Queues (optional)

This step configures the minimum bandwidth guarantee for each CoS queue when the switch is congested. If the switch is not congested, this setting has no effect. The min-bandwidth setting guarantees that any particular CoS queue is serviced often enough to ensure that the offered load can achieve the minimum transfer rate. The bandwidth is measured internally as bytes transferred per second. The minimum bandwidth setting is enforced on the egress queue; it does not rate-limit incoming frames. The minimum bandwidth setting is configured as a percentage of the total bandwidth and must be less than or equal to 100%. It may be configured on a single interface, a range of interfaces, or all interfaces.

The minimum bandwidth setting on the CoS queues comes in to play when there is congestion among the CoS queues belonging to a single TCG. This is an optional setting and is not generally required, as the secondary scheduler has the capability of guaranteeing minimum bandwidth for a TCG. If this value is set, the sum of the individual minimum bandwidths of the CoS queues belonging to any TCG should be set to the same value as the minimum bandwidth value of the TCG in the secondary scheduler, or no minimum bandwidth value should be configured for the TCG.

```
console(config-if-Te1/0/1)#cos-queue min-bandwidth 20 35 35 10 0 0 0
```

4. Configure the Scheduler Mode for the CoS Queues

This step enables strict priority scheduling on one or more CoS queues (traffic classes). Strict priority scheduling ensures that packets assigned to a higher CoS queue number are serviced before packets assigned to lower CoS queue numbers. The strict priority setting does not restrict frame ingress. Use the 'no' command to disable strict priority mode scheduling.

It is recommended that all queues enabled for strict priority scheduling be assigned to a single strict-priority-enabled TCG other than TCG0.

The following example sets CoS queue (traffic class) number 3 to be serviced with strict priority:

```
console(config-if-Te1/0/1)#cos-queue strict 3
```

To show the minimum bandwidth and scheduler modes for CoS queues, use the following command.

```
console#show interfaces cos-queue tengigabitethernet 1/0/1

Interface..................................... Te1/0/1
Interface Shaping Rate........................ 0 kbps
WRED Decay Exponent........................... 9

Queue Id   Min. Bandwidth   Scheduler Type   Queue Management Type
--------   --------------   --------------   ---------------------
0          20               Weighted         Tail Drop
1          35               Weighted         Tail Drop
2          35               Weighted         Tail Drop
3          10               Strict           Tail Drop
4          0                Weighted         Tail Drop
5          0                Weighted         Tail Drop
6          0                Weighted         Tail Drop
```

5. Map the CoS Queues to TCGs

In this step, CoS queues are mapped to Traffic Class Groups (TCGs). Since TCGs are serviced from highest numbered TCG to lowest, higher priority traffic should be assigned to higher numbered TCGs. In general, strict priority traffic (typically control plane or low bandwidth, low latency traffic) is assigned the highest numbered TCG. It is recommended that WDRR queues be assigned to TCG0.

The mapping may be configured on a single interface, a range of interfaces, or all the interfaces.

It is required that TCGs always be assigned in order from 0 to 2. It is further recommended that the operator always utilize consecutive TCGs starting with TCG 0; for example, use TCG 0 and 1 vs. TCG 0 and 2 only.

This example assigns best effort traffic to TCG 0, lossless traffic to TCG 1, and expedited service traffic to TCG 2.

```
console(config-if-Te1/0/1)# classofservice traffic-class-group 0 0
console(config-if-Te1/0/1)# classofservice traffic-class-group 1 1
console(config-if-Te1/0/1)# classofservice traffic-class-group 2 1
console(config-if-Te1/0/1)# classofservice traffic-class-group 3 2
```

To show the CoS queue to TCG mapping, use the following command:

```
console#show classofservice traffic-class-group tengigabitethernet
1/0/1
 Traffic Class    Traffic Class Group
 -------------    -------------------
       0                0
       1                1
       2                1
       3                2
       4                0
       5                0
       6                0
```

6. Set the Weight for Each TCG

This step configures the scheduling weight for each TCG. The weight percentage configures the scheduler to ensure that each TCG is nominally serviced in the ratio specified.

Each WDRR TCG should be assigned a nonzero weight. Weights may be configured on a single interface, a range of interfaces, or all interfaces, and must sum to 100%. It is recommended that strict priority TCGs be assigned a weight of 0%, since they are processed first and ignore the configured TCG weight.

It is recommended that the sum of minimum bandwidth percentages configured on the CoS queues mapped to any TCG be less than or equal to that of the weight percentage configured for the TCG, so that packets are not dropped due to the congestion in the TCG.

In example below, the TCG0 and TCG1 are allocated 30% and 70% of the bandwidth remaining after servicing TCG2 (strict priority) traffic. TCG2 traffic is handled with strict priority but can only consume up to 100% minus the sum of the minimum bandwidths of TCG0 and TCG1 (60%).

```
console(config-if-Te1/0/1)#traffic-class-group weight 30 70 0
```

7. Set the Minimum Bandwidth for Each TCG (optional)

In this step, the minimum bandwidth is configured for each TCG while in congestion. The minimum bandwidth setting ensures that a TCG is serviced during congestion even if a higher priority TCG has an offered load greater than 100% minus the minimum bandwidth assigned to the other TCGs.

The minimum bandwidth is enforced on the egress TCG regardless of the scheduling mode and does not directly affect incoming traffic. The minimum bandwidth for a TCG is configured as a percentage of the total bandwidth, and the configured minimum bandwidths may sum to less than 100%. The sum may not exceed 100%. Minimum bandwidth may be configured on a single interface, a range of interfaces, or all interfaces. It is recommended that the minimum bandwidth for a TCG be configured to be less than or equal to the weight of the TCG.

```
console(config-if-Te1/0/1)#traffic-class-group min-bandwidth 10 50 0
```

8. Set the Maximum Bandwidth for each TCG (optional)

In this step, the maximum bandwidth for each TCG is configured. The bandwidth allowed for the TCG will never exceed the maximum bandwidth configured, even if configured for strict priority. This limit is configured as a percentage of the total bandwidth and is used to shape egress traffic bursts to no greater than the configured value. The maximum bandwidth may be configured on a single interface, range of interfaces or all interfaces. When configured to be 0, unlimited bandwidth is allowed on the TCG.

It is recommended that the maximum bandwidth be configured to be greater than the minimum bandwidth or the weight or be configured to 0 (unlimited burst size).

```
console(config-if-Te1/0/1)#traffic-class-group max-bandwidth 50 90 20
```

9. Set the Scheduler Modes for the TCGs

This step enables strict priority scheduling on TCGs. Strict priority scheduling on multiple TCGs prioritizes traffic from the highest numbered TCG for transmission first. Strict priority scheduling on a single TCG selects that TCG for transmission before the WDRR TCGs. Use the 'no' command to disable strict priority scheduling. It is recommended that all CoS queues enabled for strict priority scheduling be assigned to a single TCG other than TCG0. This scheme allows a larger number of priorities to be configured as strict priorities.

```
console(config-if-Te1/0/1)#traffic-class-group strict 2
```

To show the weight, minimum bandwidth, maximum bandwidth, and scheduler modes of the TCGs, use the following command:

```
console#show interfaces traffic-class-group tengigabitethernet 1/0/1

Interface..................................... Te1/0/1

Traffic Class   Min.        Max
Group           Bandwidth   Bandwidth    Weight   Scheduler Type
------------    ---------   ------------ ------   --------------
0               10          50           30       Weighted Round Robin
1               50          90           70       Weighted Round Robin
2               0           20           0        Strict
```

ETS Theory of Operation

First Level of Scheduling

To understand the first level of scheduling, consider Table 29-1. Assume that we have eight ingress ports, each one receiving line rate traffic with one dot1p priority each. The table shows the mapping of dot1p priorities to the cos-queues, the min-bandwidth settings, and scheduler modes.

Table 29-3. First Level of Scheduling

Dot1p Priority	CoS Queue	CoS Min-bandwidth	Scheduler	TCG	Input to TCG
0	0	10	WRR	0	100% of pri0
1	1	10	WRR	1	10% each of
2	2	10	Strict	1	pri1, pri3 and
3	3	10	WRR	1	80% of pri2
4	4	15	WRR	2	25% each
5	5	15	WRR	2	of pri4, pri5
6	6	30	WRR	2	and 50% of
7	6	–	–	–	pri6/pri7

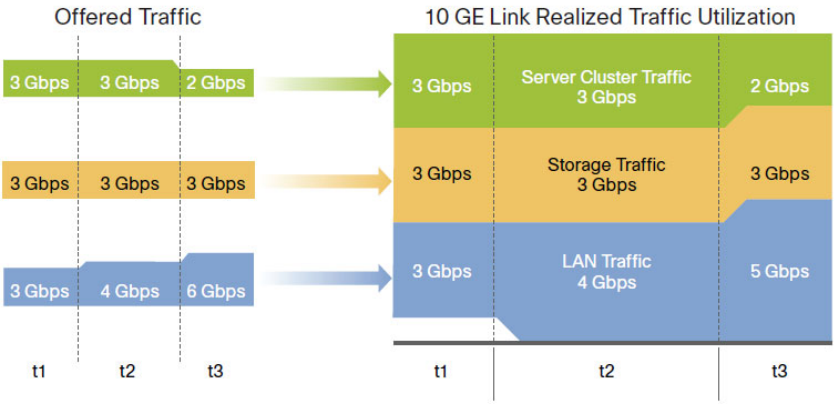
In this scenario, the input to the TCGs is as follows:

- All of the pri0 traffic reaches the TCG0. This is because there is no other priority contending for the same TCG. The minimum bandwidth setting on the CoS queue does not have any effect.
- TCG1 would receive 10% each of pri1 and pri3 and 80% of pri2. Even though strict mode is enabled for pri2, the minimum bandwidth of pri1 and pri3 is first honored before applying strict mode on pri2.
- TCG2 receives 25% each of pri4 and pri5 traffic and the other 50% can be of pri6 or pri7. This is based on the minimum bandwidth ratio, which is 15:15:30, converted to 1:1:2.

Second Level of Scheduling

To consolidate different traffic classes within different traffic types in a typical DCB environment, ETS provides an operational model for prioritization and bandwidth allocation for traffic. Figure 29-3 illustrates a typical example that consolidates three traffic types on a single 10GE link. For consolidation to be effective all traffic types must be serviced according to their requirements. For instance, SAN traffic must be guaranteed to be lossless, and the server cluster traffic (IPC) should be guaranteed a sufficiently high priority to meet the requirement of low latency.

Figure 29-3. Converged Link on the DCB Environment



In this example, to ensure that the server cluster traffic has low latency, it may be assigned to a TCG—say, TCG0—and a strict mode of scheduling is enabled on this group (weight set to 0%). SAN traffic can be assigned to TCG1 and LAN to TCG2. The TCG1 and TCG2 can be set to a weight of 50% each. With these configurations, Figure 29-3 illustrates how the load is managed at different times.

When the offered load is 3 Gbps for each traffic type at time t1, the allocated bandwidth for TCG1 and TCG2 is 3.5 Gbps each. All traffic types, including IPC, SAN, and LAN, are allowed to be transmitted, since the offered load is under the allocated bandwidth.

At time t2, a burst of LAN traffic is incoming at the rate of 4 Gbps, this burst is allowed to borrow the unused 0.5 Gbps bandwidth from SAN TCG and transmitted since the offered load of SAN is only 3 Gbps.

At time t3, when the offered load of IPC falls to 2 Gbps and the bursty LAN traffic is at 6 Gbps, the available bandwidth for SAN and LAN is 4 Gbps each according to the TCG weights, which are set as 50% each. However, ETS allows the LAN traffic to borrow unused bandwidth from SAN traffic, and then 5 Gbps of LAN traffic is transmitted.

Recommendations and Notes

Lossless traffic must map user priorities one-to-one onto CoS queues. Lossless traffic must always be carried within a VLAN-tagged frame to differentiate it from untagged best-effort LAN control traffic, such as ARP and LLDP.

Strict-priority CoS queues may starve other CoS queues of traffic if the offered load of the strict-priority CoS queues equals or exceeds the available capacity of the egress interfaces. It is recommended that either a CoS-queue-level min-bandwidth setting be utilized to ensure a minimum amount of bandwidth is processed on the non-strict-priority queues if there is a possibility that the strict priority traffic is not limited in bandwidth by some other means.

It is recommended that the sum of the minimum bandwidth percentages allocated to the group of CoS queues mapped to a single TCG be less than or equal to the weight percentage, so that packets are not dropped due to the congestion in the TCG.

The maximum bandwidth for a TCG should always be configured to be greater than the minimum bandwidth or the weight of the TCG.

In the case where all TCGs are configured as strict priority, inter-TCG scheduling is Weighted Deficit Round Robin, using the configured weights.

Configuring CoS queues with a total maximum bandwidth that is less than the corresponding TCG weight will result in bandwidth being wasted in the secondary scheduler.

Configuring CoS queues with a total minimum bandwidth that is greater than the corresponding TCG weight will flood the secondary scheduler and result in the minimum bandwidth parameter effectively overriding the TCG weight setting.

Traffic is passed across stacking links using WDRR for all CoS queues. This will affect the observed behavior of ETS on egress ports scheduling traffic from over-subscribed stacking links.

The three supported traffic class groups support an industry standard configuration such that one traffic class group offers lossless service (PFC enabled using WRED); one traffic class group contains priorities which are best effort (PFC disabled using WRED); and one traffic class group offers an expedited service (utilizes strict priority). When auto configuration is enabled on ports using ETS, the traffic classes are combined according to the above. Internally, the DCBX TCG mapping is restricted to using three TCGs regardless of the number of TCGs advertised by the configuration source, with the strict priority traffic mapped onto TCG 2, lossless traffic mapped onto TCG1, and all other traffic mapped onto TCG0.

For ETS to be operational in a manually configured environment, the following minimum steps must be performed:

1. Configure the CoS queue to Traffic Class Group mapping for the egress ports.
2. Enable the appropriate scheduling algorithm for each TCG.
3. Configure the weight percentage for each TCG.

Variation on the Example Configuration

This example configures three classes of traffic and utilizes the secondary (ETS) scheduler only:

- Best effort traffic: CoS Queue 0 for untagged and VLAN tagged frames with VPTs: 0, 1, and 2
- Lossless iSCSI: CoS Queue 1, VLAN tagged frames with VPT equal to 3
- Expedited Traffic: CoS Queue 2, VLAN tagged frames with VPTs: 4, 5, 6, 7

```
console(config-if-Te1/0/2)#classofservice dot1p-mapping 0 0
console(config-if-Te1/0/2)#classofservice dot1p-mapping 1 0
console(config-if-Te1/0/2)#classofservice dot1p-mapping 2 0
console(config-if-Te1/0/2)#classofservice dot1p-mapping 3 1
console(config-if-Te1/0/2)#classofservice dot1p-mapping 4 2
console(config-if-Te1/0/2)#classofservice dot1p-mapping 5 2
console(config-if-Te1/0/2)#classofservice dot1p-mapping 6 2
console(config-if-Te1/0/2)#classofservice dot1p-mapping 7 2
console(config-if-Te1/0/1)#classofservice traffic-class-group 0 0
console(config-if-Te1/0/1)#classofservice traffic-class-group 1 1
```

```
console(config-if-Te1/0/1)#classofservice traffic-class-group 2 2
console(config-if-Te1/0/1)#traffic-class-group weight 30 70 0
console(config-if-Te1/0/1)#traffic-class-group strict 2
```

PC81xx Operation

When DCBx is enabled on manually configured ports, it is not necessary for the ETS parameters to match, regardless of the version of DCBX negotiated or configured. Configuration mismatches are logged.

In auto configuration mode, ETS parameters from the configuration source are checked (Max TCs 3 and bandwidth equal to 100%) and if the system is capable of performing the configuration, it is accepted and propagated as received to the other auto-configuration ports. The ETS Recommendation TLVs are preferred over the ETS Configuration TLVs. Auto-configuration via DCBX overrides manually configured ETS parameters for auto-configured ports, however, manual configuration is restored should the port be placed back into the manual port role.

The ETS parameters received via DCBX are modified and applied to the system via the DCBX Mapping function as follows (references are to the 802.1Qaz parameters):

- Like traffic classes are combined up to the limits of the system; e.g., no more than 2 lossless CoS queues may be configured.
- The Priority Assignment Table (user priority to CoS queue mapping) is utilized by the system to map user priorities to the traffic classes (CoS queues).
- The TSA Assignment Table is converted to use 3 TCGs internally. Priorities with like characteristics are combined into TCGs, i.e. strict priority traffic is combined into a TCG, lossless traffic is combined into a TCG, etc. Generally, strict priority traffic is mapped onto TCG 2, lossless traffic is mapped onto TCG1 and best effort traffic is mapped onto TCG0.
- The bandwidths from the TC Bandwidth Table are summed based on the internal TCG mapping and are used to set the TCG weights. Other switches may assign bandwidth to strict priority queues. This bandwidth is counted in the sum to ETS 100% validation check, however, internally the weight for strict priority queues is ignored and they are configured for unlimited bandwidth. An implication of this assignment is that the percentage of bandwidth that may be consumed by a WDRR TCG after

processing strict priority traffic is skewed to be the bandwidth of the individual TCG divided by the sum of the weights of all WDRR configured TCGs.

The administrator may configure other parameters to work in conjunction with the received DCBX configuration, e.g. min-bandwidth per CoS queue and minimum or maximum bandwidth per TCG.

PC8024 Operation with DCBx

PowerConnect 8024 and 8024F switches can act as a proxy for ETS information via the auto configuration mechanism. ETS information received from the configuration source is transmitted via DCBX to the other auto configuration peers.

While 8024 and 8024F switches are not equipped with a hierarchical scheduler, they do support the following ETS capabilities:

- Both lossy and lossless service using PFC (may be configured via DCBX).
- Sharing of unused bandwidth with lower priority traffic classes (default behavior).
- WRR or strict-priority scheduling (configured manually).

30

Managing the MAC Address Table

This chapter describes the L2 MAC address table the switch uses to forward data between ports.

The topics covered in this chapter include:

- MAC Address Table Overview
- Default MAC Address Table Values
- Managing the MAC Address Table (Web)
- Managing the MAC Address Table (CLI)

MAC Address Table Overview

The MAC address table keeps track of the MAC addresses that are associated with each port to allow the switch to forward unicast traffic through the appropriate port. This table is sometimes called the bridge table or the forwarding database.

How Is the Address Table Populated?

The MAC address table can contain two types of addresses:

- Static: The address has been manually configured and does not age out.
- Dynamic: The address has been automatically learned by the switch and can age out when it is not in use.

Static addresses are configured by the administrator and added to the table. Dynamic addresses are learned by examining information in the Ethernet frame.

When a frame arrives on a port, the switch looks at the frame header to learn the source MAC address of the frame, then adds the address, VLAN ID, and the ingress port to the MAC address table. The address table is constantly updated as new addresses are learned, and unused addresses age out.

A frame that has a destination MAC address that matches an entry in the table is forwarded immediately to the associated port(s)/VLAN(s).

What Information Is in the MAC Address Table?

Each entry in the address table, whether it is static or dynamic, includes the MAC address, the VLAN ID associated with the MAC address, and the interface on which the address was learned or configured.

Each port can maintain multiple MAC addresses, and a MAC address can be associated with multiple VLANs.

How Is the MAC Address Table Maintained Across a Stack?

The MAC address table is synchronized across all stack members. When a member joins the stack, its previous MAC address table is overwritten by the table maintained by the stack.

Default MAC Address Table Values

Table 30-1 summarizes the default values for the MAC address table.

Table 30-1. MAC Address Table Defaults

Parameter	Default Value
Aging time	300 seconds
Dynamic addresses	Enabled (automatically learned)
Static addresses	None configured

Managing the MAC Address Table (Web)

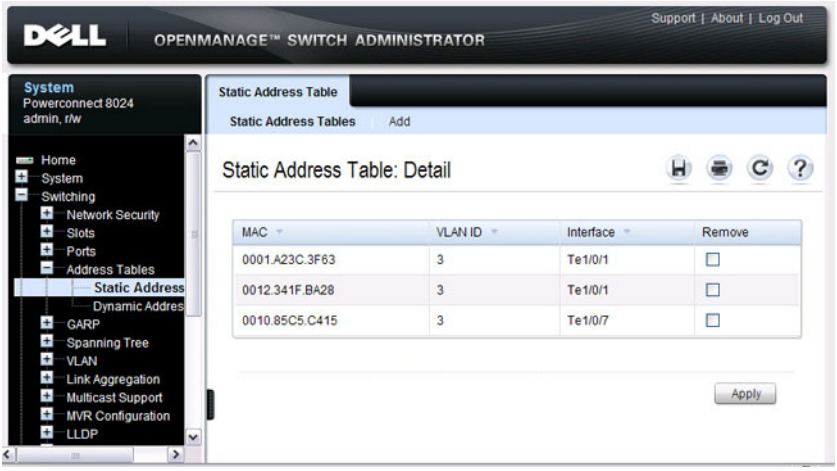
This section provides information about the OpenManage Switch Administrator pages to use to manage the MAC address table on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click at the top of the page.

Static Address Table

Use the **Static Address Table** page to view MAC addresses that have been manually added to the MAC address table and to configure static MAC addresses.

To display the **Static Address Table** page, click **Switching** → **Address Tables** → **Static Address Table** in the navigation panel.

Figure 30-1. Static MAC Address



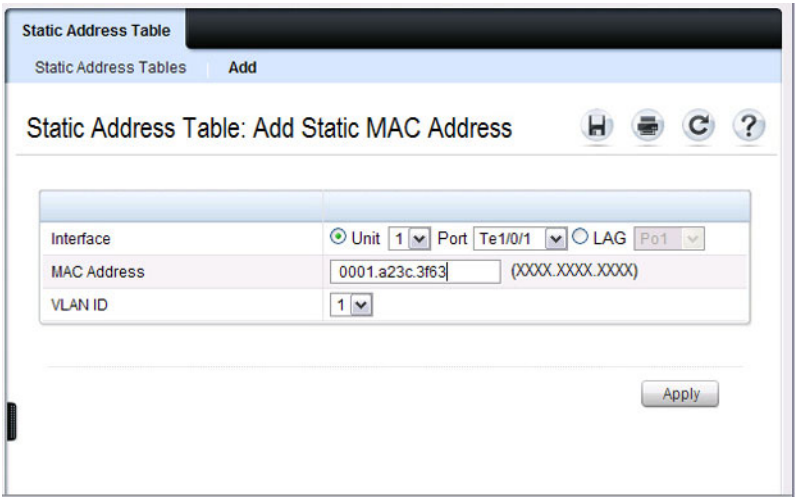
Adding a Static MAC Address

To add a static MAC address:

1. Open the **Static MAC Address** page.
2. Click **Add**.

   The **Add Static MAC Address** page displays.

Figure 30-2. Adding Static MAC Address



3 Select the interface to associate with the static address.

4 Specify the MAC address and an associated VLAN ID.

5 Click **Apply**.

The new static address is added to the **Static MAC Address Table**, and the device is updated.

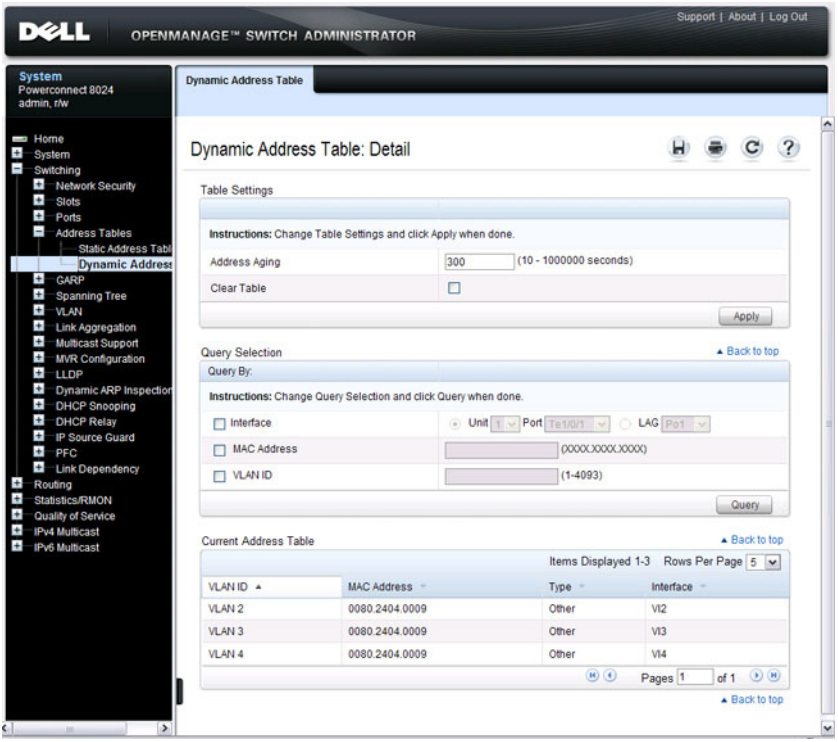
Dynamic Address Table

The **Dynamic Address Table** page contains fields for querying information in the dynamic address table, including the interface type, MAC addresses, VLAN, and table sorting key. Packets forwarded to an address stored in the address table are forwarded directly to those ports.

The **Dynamic Address Table** also contains information about the aging time before a dynamic MAC address is removed from the table.

To display the **Dynamic Address Table**, click **Switching** → **Address Tables** → **Dynamic Address Table** in the navigation panel.

Figure 30-3. Dynamic Address Table



Managing the MAC Address Table (CLI)

This section provides information about the commands you use to manage the MAC address table on the switch. For more information about the commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Managing the MAC Address Table

Beginning in Privileged EXEC mode, use the following commands to add a static MAC address to the table, control the aging time for dynamic addresses, and view entries in the MAC address table.

Command	Purpose
configure	Enter global configuration mode.
mac address-table static *mac-address* **vlan** *vlan-id* **interface** *interface*	Add a static MAC source address to the MAC address table. • *mac-address* — A valid MAC address in the format xxxx.xxxx.xxxx. • *vlan-id* —A valid VLAN. • *interface* — A valid port or LAG, including the interface type and number.
mac address-table aging-time {0 | 10-1000000}	Specify the number of seconds that must pass before an unused dynamically-learned MAC address is removed from the MAC address table. A value of 0 disables the aging time for the MAC address table.
exit	Exit to Privileged EXEC mode.
show mac address-table [**static** | **dynamic**]	View information about the entries in the MAC address table. Use the keywords **static** or **dynamic** to specify the address type to view. For dynamic entries, you can use the **clear mac address-table** command to remove entries from the table.
show mac address-table {**vlan** *vlan* | **interface** *interface* [**vlan** *vlan-id*]}	View information about the MAC addresses that have been configured or learned on the switch, a specific VLAN, or an interface (Ethernet port or LAG/port-channel).
show mac address-table count [{**vlan** *vlan-id* |**interface** *interface*}]	View information about the number of addresses that have been configured or learned on the switch, a specific VLAN, or an interface (Ethernet port or LAG/port-channel).

31

Configuring Routing Interfaces

This chapter describes the routing (layer 3) interfaces the PowerConnect 8000-series and 8100-series switches support, which includes VLAN routing interfaces, loopback interfaces, and tunnel interfaces.

The topics covered in this chapter are:

- Routing Interface Overview
- Default Routing Interface Values
- Configuring Routing Interfaces (Web)
- Configuring Routing Interfaces (CLI)

For information about configuring IPv6 characteristics on routing interfaces, see "Configuring IPv6 Routing" on page 1059.

For configuration examples that configure VLAN routing interfaces, see "IP Routing Configuration Example" on page 904 in the Configuring IP Routing chapter. For a configuration example that includes tunnel and loopback interface creation, see "Interconnecting an IPv4 Backbone and Local IPv6 Network" on page 1005.

Routing Interface Overview

Routing interfaces are logical interfaces that can be configured with an IP address. Routing interfaces provide a means of transmitting IP packets between subnets on the network.

What Are VLAN Routing Interfaces?

VLANs divide a single physical network (broadcast domain) into separate logical networks. To forward traffic across VLAN boundaries, a layer 3 device, such as router, is required. PowerConnect 8000-series and 8100-series switches can act as layer 3 devices when you configure VLAN routing interfaces. VLAN routing interfaces make it possible to transmit traffic

between VLANs while still containing broadcast traffic within VLAN boundaries. The configuration of VLAN routing interfaces makes inter-VLAN routing possible.

For each VLAN routing interface you can assign a static IP address, or you can allow a network DHCP server to assign a dynamic IP address.

When a port is enabled for bridging (L2 switching) rather than routing, which is the default, all normal bridge processing is performed for an inbound packet, which is then associated with a VLAN. Its MAC Destination Address (MAC DA) and VLAN ID are used to search the MAC address table. If routing is enabled for the VLAN, and the MAC DA of an inbound unicast packet is that of the internal router interface, the packet is routed. An inbound multicast packet is forwarded to all ports in the VLAN, plus the internal bridge-router interface, if it was received on a routed VLAN.

Since a port can be configured to belong to more than one VLAN, VLAN routing might be enabled for all of the VLANs on the port or for only some of the VLANs on the port. VLAN Routing can be used to allow more than one physical port to reside on the same subnet. It could also be used when a VLAN spans multiple physical networks, or when additional segmentation or security is required.

What Are Loopback Interfaces?

A loopback interface is a logical interface that is always up and, because it cannot go down, allows the switch to have a stable IP address that other network devices and protocols can use to reach the switch. The loopback can provide the source address for sent packets.

NOTE: In this context, loopback interfaces should not be confused with the loopback IP address, usually 127.0.0.1, assigned to a host for handling self-routed packets.

The loopback interface does not behave like a network switching port. Specifically, there are no neighbors on a loopback interface; it is a pseudo-device for assigning local addresses so that the other layer 3 devices can communicate with the switch by using the loopback IP address. The loopback interface is always up and can receive traffic from any of the existing active interfaces. Thus, given reachability from a remote client, the address of the loopback can be used to communicate with the switch through various

services such as Telnet and SSH. In this way, the IP address on a loopback behaves identically to any of the local addresses of the VLAN routing interfaces in terms of the processing of incoming packets.

What Are Tunnel Interfaces?

Tunnels are a mechanism for transporting a packet across a network so that it can be evaluated at a remote location or *tunnel endpoint*. The tunnel, effectively, hides the packet from the network used to transport the packet to the endpoint. This allows for the transmission of packets that the transport network cannot process directly, such as in one of the following cases:

- The packet protocol is not supported.
- The packet is in an incompatible addressing space.
- The packet is encrypted.

PowerConnect 8000-series and 8100-series switches support tunnels to encapsulate IPv6 traffic in IPv4 tunnels to provide functionality to facilitate the transition of IPv4 networks to IPv6 networks.

The switch supports two types of tunnels: configured (6-in-4) and automatic (6-to-4). Configured tunnels have an explicit configured endpoint and are considered to be point-to-point interfaces. Automatic tunnels determine the endpoint of the tunnel from the destination address of packets routed into the tunnel. These tunnels correspond to Non-Broadcast Multi-Access (NBMA) interfaces. A configured tunnel interface has a single tunnel associated with it, while an automatic tunnel interface has an infinite number of tunnels (limited only by the address encoding scheme).

Because tunnels are used as logical interfaces, you can define static routes that reference the tunnels. Additionally, dynamic routing can be configured to use the tunnels.

Why Are Routing Interfaces Needed?

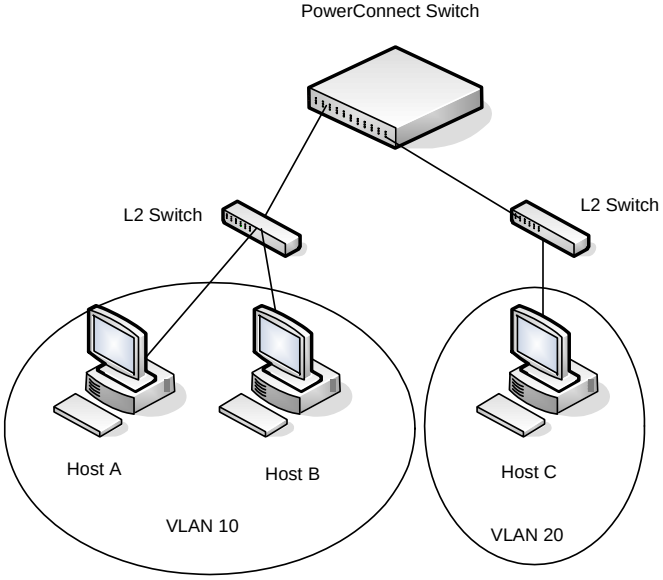
The routing interfaces this chapter describes have very different applications and uses, as this section describes. If you use the switch as a layer 2 device that handles switching only, routing interface configuration is not required. When the switch is used as a layer 2 device, it typically connects to an external layer 3 device that handles the routing functions.

VLAN Routing

VLAN routing is required when the switch is used as a layer 3 device. VLAN routing must be configured to allow the switch to forward IP traffic between subnets and allow hosts in different networks to communicate.

In Figure 31-1 the PowerConnect switch is configured as an L3 device and performs the routing functions for hosts connected to the L2 switches. For Host A to communicate with Host B, no routing is necessary. These hosts are in the same VLAN. However, for Host A in VLAN 10 to communicate with Host C in VLAN 20, the PowerConnect switch must perform inter-VLAN routing.

Figure 31-1. Inter-VLAN Routing



Loopback Interfaces

When packets are sent to the loopback IP address, the network should be able to deliver the packets as long as any physical interface on the switch is up. There are many cases where you need to send traffic to a switch, such as in switch management. The loopback interface IP address is a good choice for communicating with the switch in these cases because the loopback interface cannot go down when the switch is powered on and operational.

Tunnel Interface

Tunnels can be used in networks that support both IPv6 and IPv4. The tunnel allows non-contiguous IPv6 networks to be connected over an IPv4 infrastructure.

Default Routing Interface Values

By default, no routing interfaces are configured.

When you create a VLAN, no IP address is configured, and DHCP is disabled. After you configure an IP address on a VLAN or loopback interface, routing is automatically enabled on the VLAN interface, and the interface has the default configuration shown in Table 31-1.

Most interface configuration parameters are not applicable to loopback interfaces, so you cannot change the default values. However, when you create a loopback interface, the default values are similar to those of VLAN routing interfaces, as Table 31-1 shows.

Table 31-1. VLAN Routing Interface and Loopback Interface Defaults

Parameter	Default Value
Forward Net Directed Broadcasts	Disabled
Encapsulation Type	Ethernet (N/A for loopbacks)
Proxy Arp	Enabled
Local Proxy Arp	Disabled
IP MTU	1500
Bandwidth	Not configured.
Destination Unreachables	Enabled
ICMP Redirects	Enabled

When you create a tunnel, it has the default values shown in Table 31-2

Table 31-2. Tunnel Interface Defaults

Parameter	Default Value
Tunnel mode	6-in-4 configured
Link Local Only Mode	Disabled
Source address	None
Destination address	0.0.0.0

Configuring Routing Interfaces (Web)

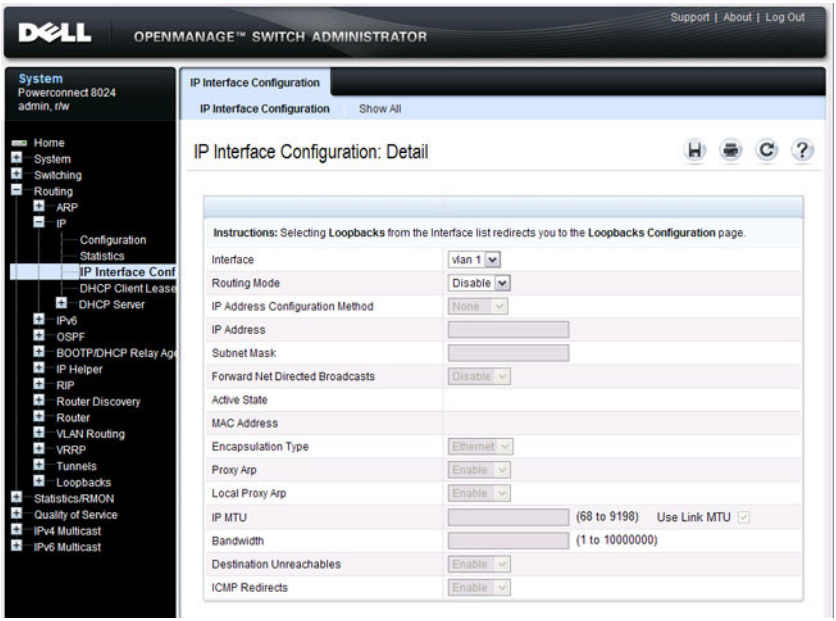
This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring VLAN routing interfaces, loopback interfaces, and tunnels on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

IP Interface Configuration

Use the **IP Interface Configuration** page to update IP interface data for this switch. The IP interface configuration includes the ability to configure the bandwidth, Destination Unreachable messages, and ICMP Redirect messages.

To display the page, click **Routing → IP → IP Interface Configuration** in the navigation panel.

Figure 31-2. IP Interface Configuration

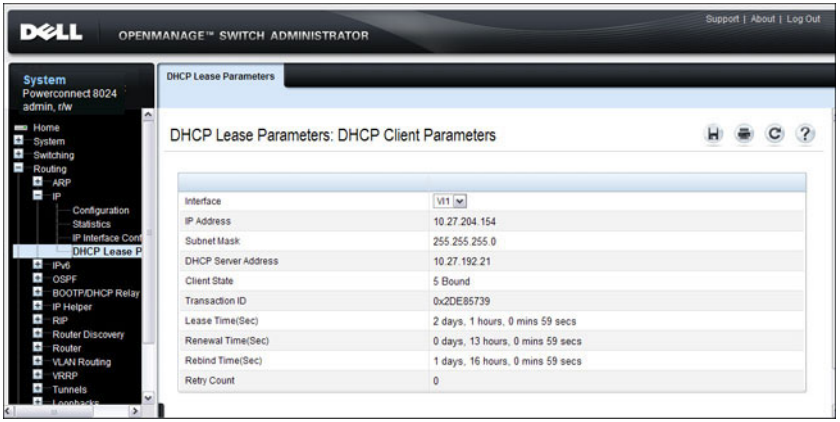


DHCP Lease Parameters

Use the **DHCP Lease Parameters** page to view information about the network information automatically assigned to an interface by the DHCP server.

To display the page, click **Routing** → **IP** → **DHCP Lease Parameters** in the navigation panel.

Figure 31-3. DHCP Lease Parameters

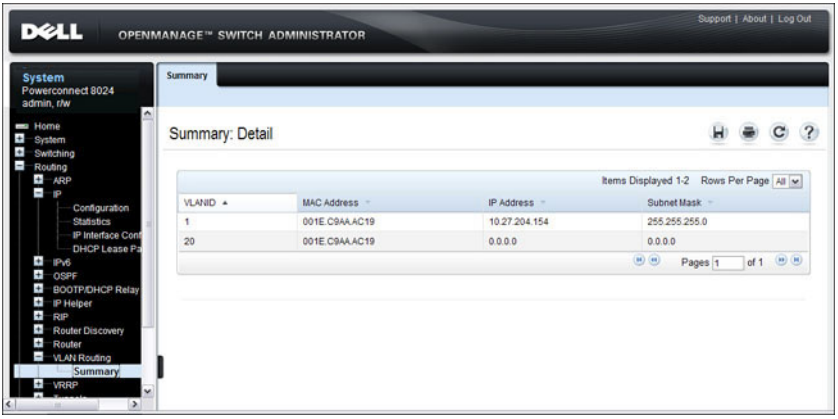


VLAN Routing Summary

Use the **VLAN Routing Summary** page to view summary information about VLAN routing interfaces configured on the switch.

To display the page, click **Routing** → **VLAN Routing** → **Summary** in the navigation panel.

Figure 31-4. VLAN Routing Summary

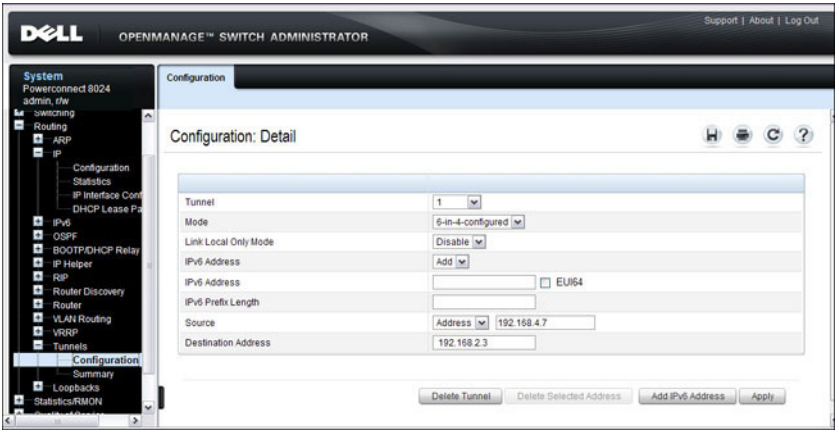


Tunnel Configuration

Use the **Tunnels Configuration** page to create, configure, or delete a tunnel.

To display the page, click **Routing** → **Tunnels** → **Configuration** in the navigation panel.

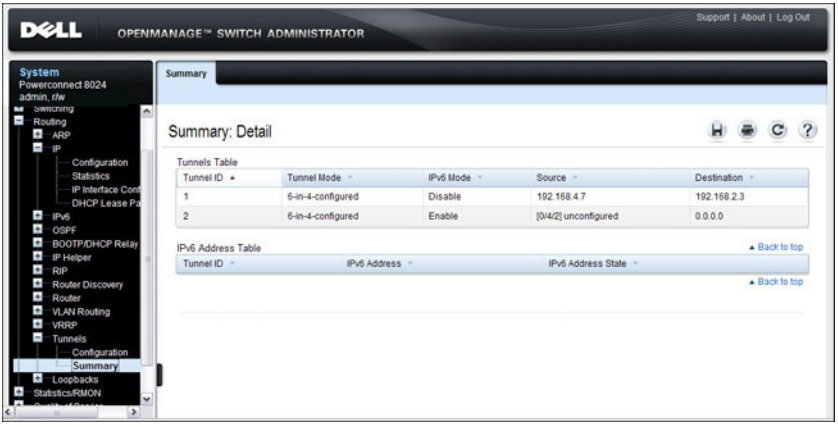
Figure 31-5. Tunnel Configuration



Tunnels Summary

Use the **Tunnels Summary** page to display a summary of configured tunnels. To display the page, click **Routing** → **Tunnels** → **Summary** in the navigation panel.

Figure 31-6. Tunnels Summary

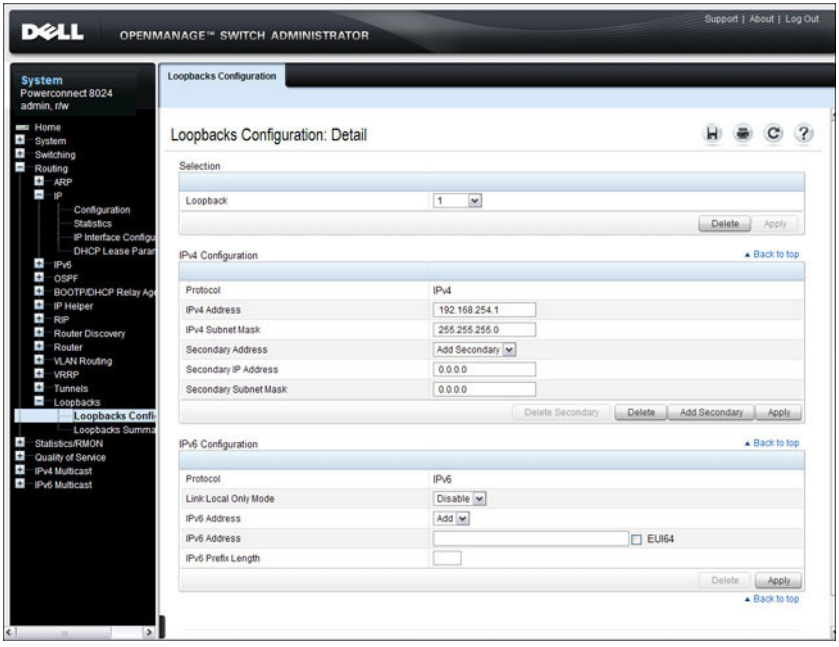


Loopbacks Configuration

Use the **Loopbacks Configuration** page to create, configure, or remove loopback interfaces. You can also set up or delete a secondary address for a loopback.

To display the page, click **Routing → Loopbacks → Loopbacks Configuration** in the navigation panel.

Figure 31-7. Loopback Configuration

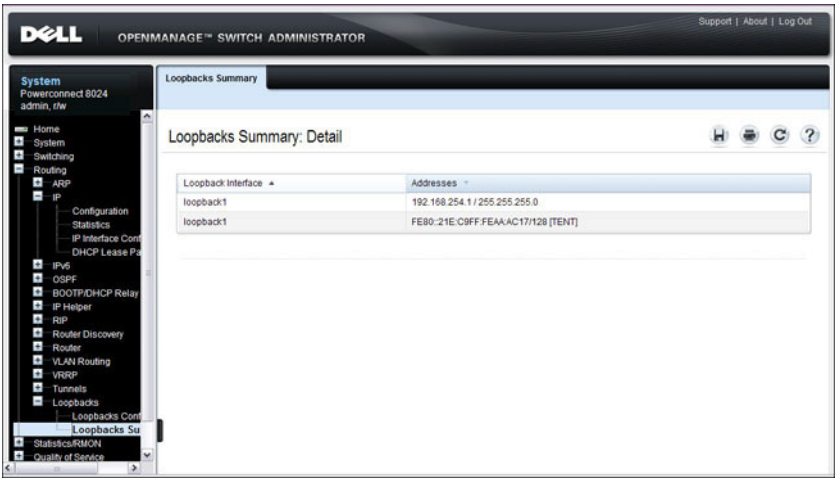


Loopbacks Summary

Use the **Loopbacks Summary** page to display a summary of configured loopback interfaces on the switch.

To display the page, click **Routing** → **Loopbacks** → **Loopbacks Summary** in the navigation panel.

Figure 31-8. Loopbacks Summary



Configuring Routing Interfaces (CLI)

This section provides information about the commands you use to configure VLAN routing interfaces, loopbacks, and tunnels on the switch. For more information about the commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Configuring VLAN Routing Interfaces (IPv4)

Beginning in Privileged EXEC mode, use the following commands to configure a VLAN as a routing interface and set the IP configuration parameters.

Command	Purpose
configure	Enter Global Configuration mode.
interface vlan *vlan-id*	Enter Interface Configuration mode for the specified VLAN.
ip address {dhcp | none | *ip_address subnet_mask* [**secondary**]}	Configure the IP address. Use the **dhcp** keyword to enable the DHCP client and obtain an IP address from a network DHCP server. Use **none** to release the address obtained from the DHCP server. Use *ip_address* and *subnet_mask* to assign a static IP address. If you configure a static address, you can use the **secondary** keyword to specify that the address is a secondary IP address.
ip netdirbcast	Enable the forwarding of network-directed broadcasts.
encapsulation {ethernet | snap}	Configure the link-layer encapsulation type for the packet. Routed frames are always ethernet encapsulated when a frame is routed to a VLAN.
ip proxy-arp	Enable proxy ARP on the interface. Without proxy ARP, the switch responds to an ARP request only if the target IP address is an address configured on the interface where the ARP request arrived. This command is not available in interface range mode.

Command	Purpose
ip local-proxy-arp	Enable local proxy ARP on the interface to allow the switch to respond to ARP requests for hosts on the same subnet as the ARP source.
ip mtu *size*	Set the IP Maximum Transmission Unit (MTU) on a routing interface. The IP MTU is the size of the largest IP packet that can be transmitted on the interface without fragmentation. The range is 68–9198 bytes.
bandwidth *size*	Set the configured bandwidth on this interface to communicate the speed of the interface to higher level protocols. OSPF uses the bandwidth value to compute link cost. The range is 1–10000000.
ip unreachables	Allow the switch to send ICMP Destination Unreachable messages in response to packets received on the interface.
ip redirects	Allow the switch to send ICMP Redirect messages in response to packets received on the interface.
exit	Exit to Global Config mode.
ip default-gateway *ip_address*	Configure the default gateway. All switch interfaces use the same default gateway.
exit	Exit to Privileged EXEC mode.
show dhcp lease [**interface** *interface*]	View information about the DHCP leases acquired for all interfaces or for the specified interface. For a VLAN, the *interface_string* parameter is **vlan** followed by the VLAN ID, with no space, for example vlan10.
show ip interface vlan *vlan-id*	View the IP interface configuration information for the specified routing VLAN.

Configuring Loopback Interfaces

Beginning in Privileged EXEC mode, use the following commands to configure a loopback interface.

Command	Purpose
configure	Enter Global Configuration mode.
interface loopback *loopback-id*	Create the loopback interface and enter Interface Configuration mode for the specified loopback interface.
ip address *ip_address subnet_mask* [**secondary**]	Configure a static IP address and subnet mask. Use the **secondary** keyword to specify that the address is a secondary IP address.
CTRL + Z	Exit to Privileged EXEC mode.
show ip interface loopback *loopback-id*	View interface configuration information for the specified loopback interface.

Configuring Tunnels

Beginning in Privileged EXEC mode, use the following commands to configure a loopback interface.

NOTE: For information about configuring the IPv6 interface characteristics for a tunnel, see "Configuring IPv6 Routing" on page 1059.

Command	Purpose
configure	Enter Global Configuration mode.
interface tunnel *tunnel-id*	Create the tunnel interface and enter Interface Configuration mode for the specified tunnel.
tunnel mode ipv6ip [**6to4**]	Specify the mode of the tunnel. If you use the **6to4** keyword, the tunnel is an automatic tunnel. If you omit the keyword, the tunnel is a point-to-point (configured) tunnel.
ipv6 enable	Enable IPv6 on this interface using the Link Local address.
tunnel source {*ipv4addr* | **vlan** *vlan-id*}	Specify the source transport address of the tunnel, either, which can be an IPv4 address or a VLAN routing interface.
tunnel destination *ipv4addr*	Specify the destination transport IPv4 address of the tunnel.
CTRL + Z	Exit to Privileged EXEC mode.
show interfaces tunnel [*tunnel-id*]	View configuration information for all tunnels or for the specified tunnel.

32

Configuring DHCP Server Settings

This chapter describes how to configure the switch to dynamically assign network information to hosts by using the Dynamic Host Configuration Protocol (DHCP).

The topics covered in this chapter include:

- DHCP Overview
- Default DHCP Server Values
- Configuring the DHCP Server (Web)
- Configuring the DHCP Server (CLI)
- DHCP Server Configuration Examples

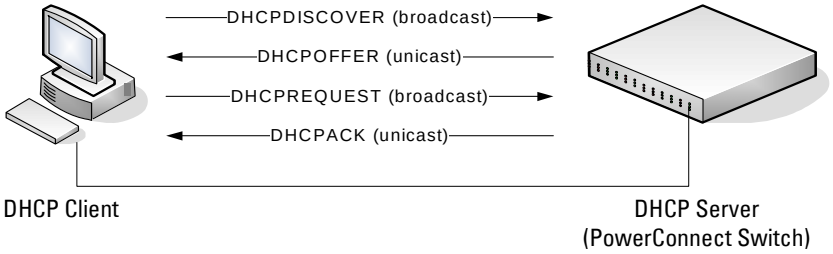
DHCP Overview

DHCP is generally used between clients and servers for the purpose of assigning IP addresses, gateways, and other network settings such as DNS and SNTP server information.

How Does DHCP Work?

When a host connects to the network, the host's DHCP client broadcasts a message requesting information from any DHCP server that receives the broadcast. One or more DHCP servers respond to the request. The response includes the requested information, such as the IP address, subnet mask, and default gateway IP address. The client accepts an offer from one of the servers, and the server sends an acknowledgment to the client to confirm the transaction.

Figure 32-1. Message Exchange Between DHCP Client and Server



The DHCP server maintains one or more set of IP addresses the and other configuration information available, by request, to DHCP clients. Each set of information is known as an address pool.

After a client leases an IP address from the DHCP server, the server adds an entry to its database. The entry is called a binding.

What are DHCP Options?

DHCP options are collections of data with type codes that indicate how the options should be used. Options can specify information that is required for the DHCP protocol, IP stack configuration parameters for the client, information allowing the client to rendezvous with DHCP servers, and so on.

When a client broadcasts a request for information, the request includes the option codes that correspond to the information the client wants the DHCP server to supply. The Web pages and CLI commands to configure DHCP server settings include many predefined options for the information that is most commonly requested by DHCP clients. For example, DHCP client discover requests typically include options for the IP address (option 50), subnet mask (option 1), default gateway (option 3), and DNS server (option 6). These options are predefined.

For options that are not predefined, you can enter the option code and specify the data type along with the data that the switch should include in DHCP offers. RFC2132 specifies many of the DHCP options. Additional options are described in later RFCs.

What Additional DHCP Features Does the Switch Support?

The switch software includes a DHCP client that can request network information from a DHCP server on the network during the initial system configuration process. For information about enabling the DHCP client, see "Setting the IP Address and Other Basic Network Information" on page 119.

If the switch is functioning as a Layer 3 device, the Layer 3 DHCP Relay Agent can relay DHCP messages between DHCP clients and DHCP servers that are located in different IP subnets.

The DHCP Layer 2 Relay feature permits Layer 3 Relay agent functionality in Layer 2 switched networks. The switch supports L2 DHCP relay configuration on individual ports, link aggregation groups (LAGs) and VLANs. For information about Layer 2 and Layer 3 DHCP Relay, see "Configuring L2 and L3 Relay Features" on page 907.

DHCP Snooping is a security feature that monitors DHCP messages between a DHCP client and DHCP server. It filters harmful DHCP messages and builds a bindings database of (MAC address, IP address, VLAN ID, port) tuples that are specified as authorized. DHCP snooping can be enabled globally and on specific VLANs. For information about DHCP Snooping, see "Snooping and Inspecting Traffic" on page 743.

Default DHCP Server Values

By default, the DHCP server is disabled, and no address pools are configured. You must create at least one address pool and enable the DHCP server to allow the switch to dynamically assign network information to hosts with DHCP clients that broadcast requests.

The DHCP server can lease a maximum of 256 addresses.

Configuring the DHCP Server (Web)

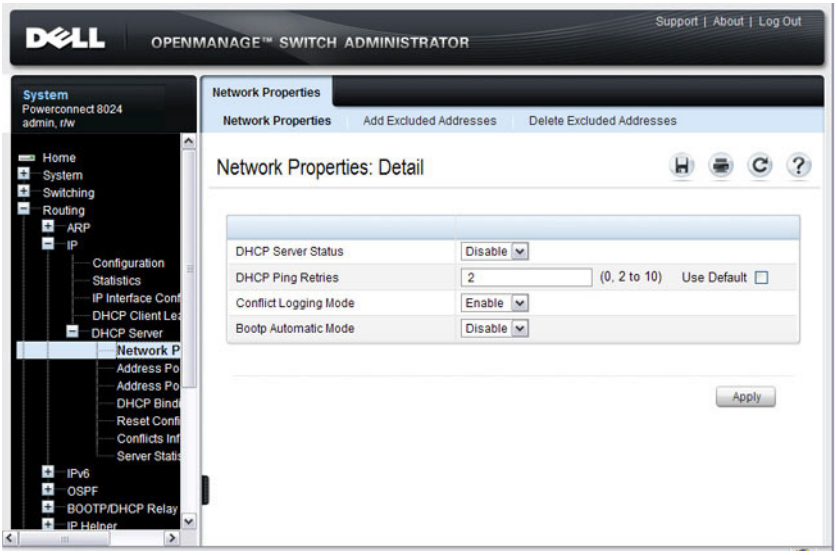
This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring the DHCP server on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

DHCP Server Network Properties

Use the **Network Properties** page to define global DHCP server settings and to configure addresses that are not included in any address pools.

To display the **Network Properties** page, click **Routing** → **IP** → **DHCP Server** → **Network Properties** in the navigation panel.

Figure 32-2. DHCP Server Network Properties



Adding Excluded Addresses

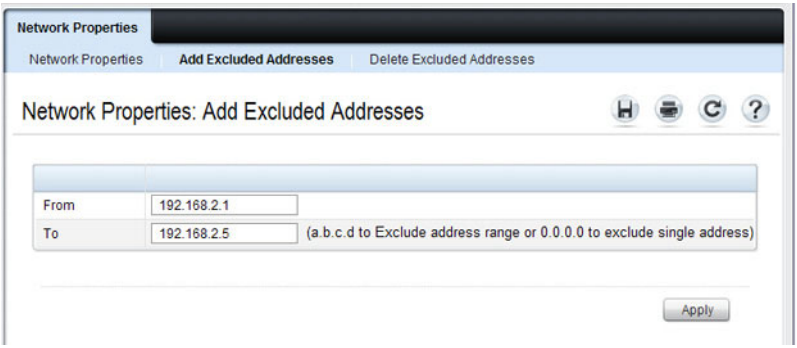
To exclude an address:

1. Open the **Network Properties** page.
2. Click **Add Excluded Addresses** to display the **Add Excluded Addresses** page.
3. In the From field, enter the first IP address to exclude from any configured address pool.
4. If the address in the From field is the only address to exclude, or if the excluded addresses are non-contiguous, leave the To field as the default value of 0.0.0.0. Otherwise, enter the last IP address to excluded from a contiguous range of IP addresses.

   In Figure 32-3, the From field contains the IP address 192.168.2.1, and the To field contains the IP address 192.168.2.5. This means that the following IP addresses are not available for lease:

   - 192.168.2.1
   - 192.168.2.2
   - 192.168.2.3
   - 192.168.2.4
   - 192.168.2.5

Figure 32-3. Add Excluded Addresses



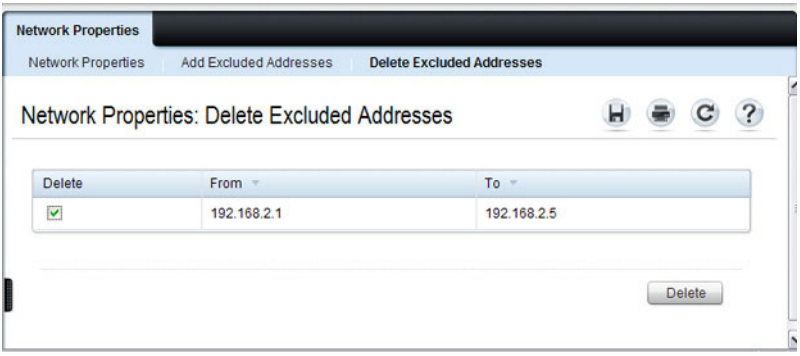
5. Click **Apply**.

Deleting Excluded Addresses

To remove an excluded address:

1. Open the **Network Properties** page.
2. Click **Delete Excluded Addresses** to display the **Delete Excluded Addresses** page.
3. Select the check box next to the address or address range to delete.

Figure 32-4. Delete Excluded Addresses



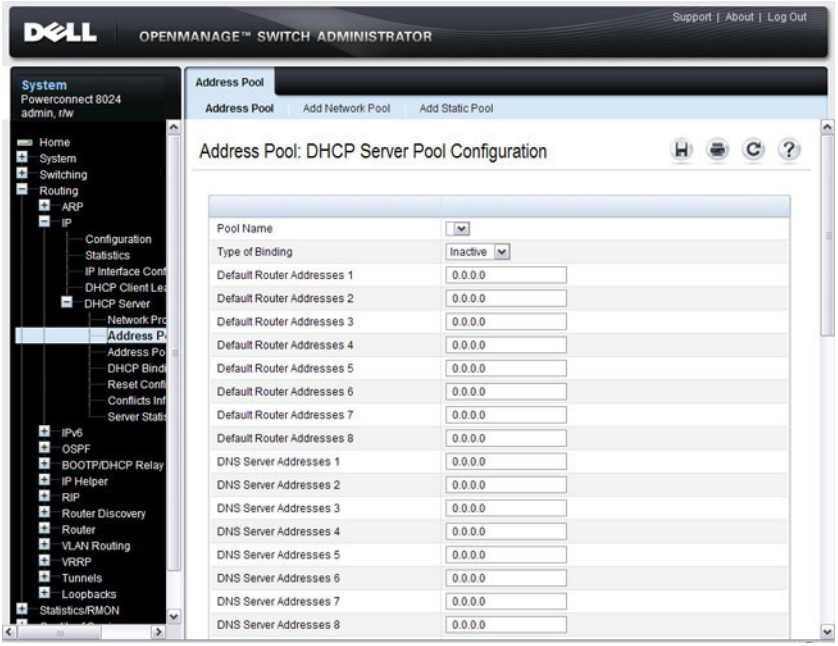
4. Click **Apply**.

Address Pool

Use the **Address Pool** page to create the pools of IP addresses and other network information that can be assigned by the server.

To display the **Address Pool** page, click **Routing** → **IP** → **DHCP Server** → **Address Pool** in the navigation panel.

Figure 32-5. Address Pool



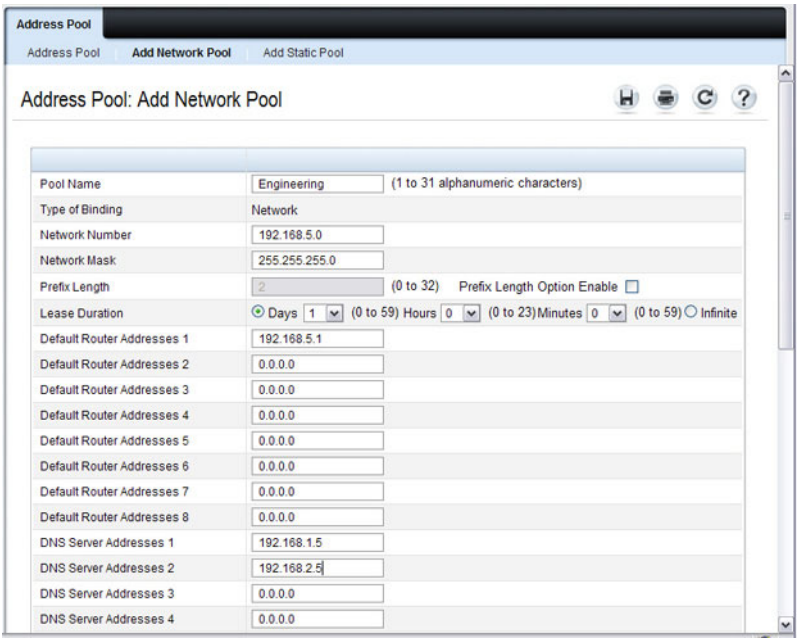
Adding a Network Pool

To create and configure a network pool:

1. Open the **Address Pool** page.
2. Click **Add Network Pool** to display the **Add Network Pool** page.
3. Assign a name to the pool and complete the desired fields.

   In Figure 32-6, the network pool name is Engineering, and the address pool contains all IP addresses in the 192.168.5.0 subnet, which means a client that receives an address from the DHCP server might lease an address in the range of 192.168.5.1 to 192.168.5.254.

Figure 32-6. Add Network Pool



The Engineering pool also configures clients to use 192.168.5.1 as the default gateway IP address and 192.168.1.5 and 192.168.2.5 as the primary and secondary DNS servers.

NOTE: The IP address 192.168.5.1 should be added to the global list of excluded addresses so that it is not leased to a client.

4 Click **Apply**.

Adding a Static Pool

To create and configure a static pool of IP addresses:

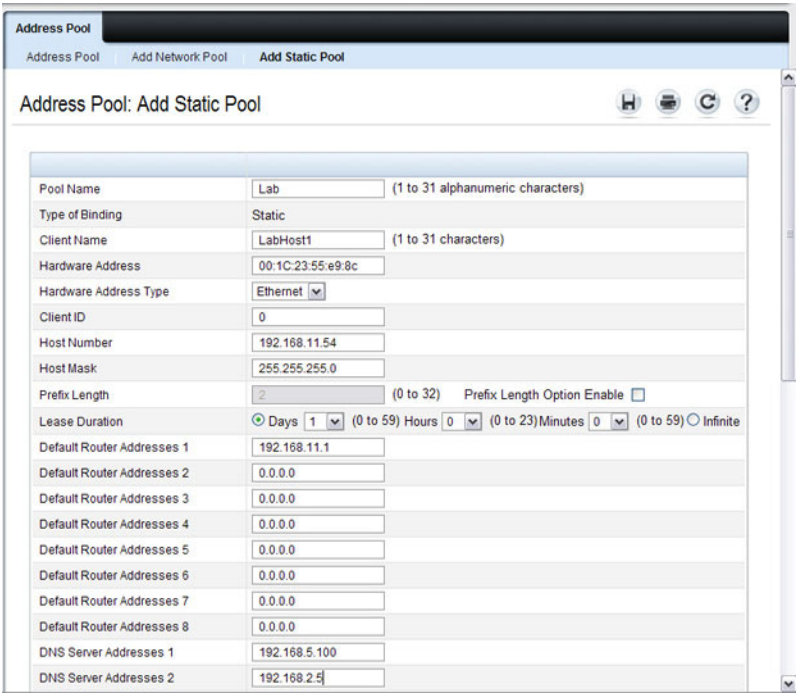
1 Open the **Address Pool** page.

2 Click **Add Static Pool** to display the **Add Static Pool** page.

3 Assign a name to the pool and complete the desired fields.

In Figure 32-7, the Static pool name is Lab, and the name of the client in the pool is LabHost1. The client's MAC address is mapped to the IP address 192.168.11.54, the default gateway is 192.168.11.1, and the DNS servers the client will use have IP addresses of 192.168.5.100 and 192.168.2.5.

Figure 32-7. Add Static Pool



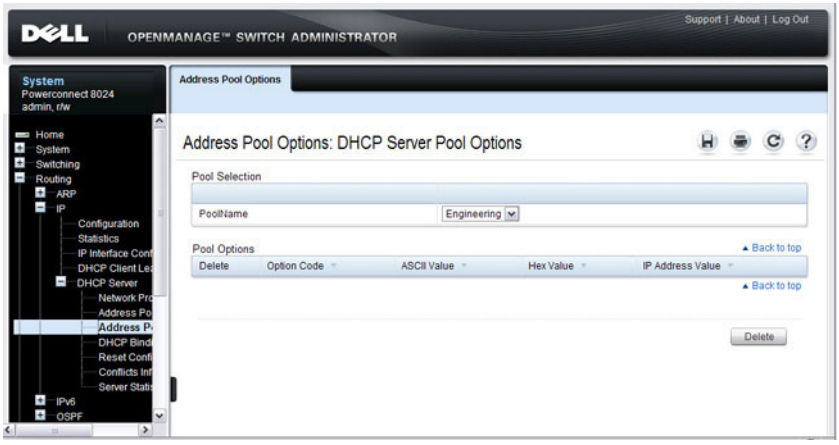
4 Click **Apply**.

Address Pool Options

Use the **Address Pool Options** page to view manually configured options. You can define options when you create an address pool, or you can add options to an existing address pool.

To display the **Address Pool Options** page, click **Routing** → **IP** → **DHCP Server** → **Address Pool Options** in the navigation panel.

Figure 32-8. Address Pool Options



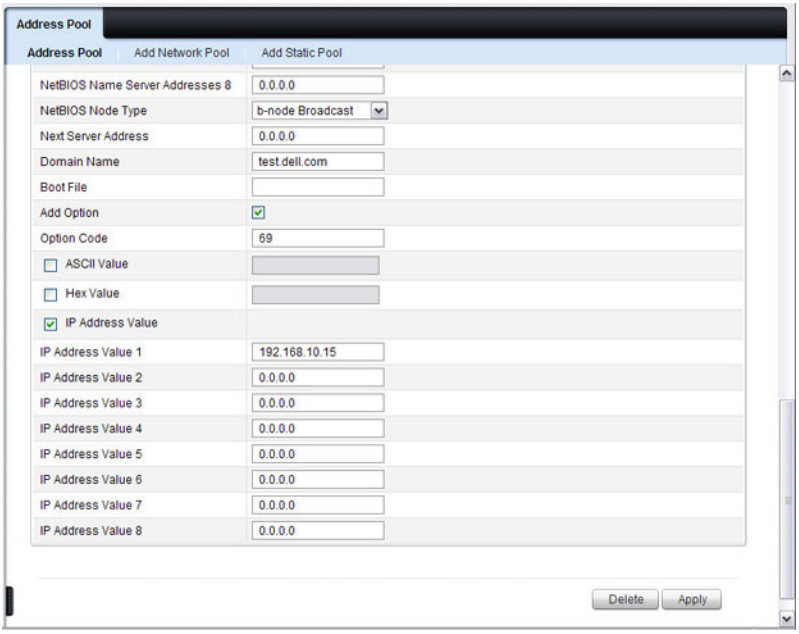
Defining DHCP Options

To configure DHCP options:

1. Open the **Address Pool** page.
2. Select the **Add Options** check box.
3. Select the check box that corresponds to the value type (ASCII, Hexadecimal, or IP address).
4. Specify the value(s) in the corresponding field.

   Figure 32-9 shows an example of adding the SMTP server IP address. The option code for the SMTP server is 69, and the IP address of the SMTP server is 192.168.10.15.

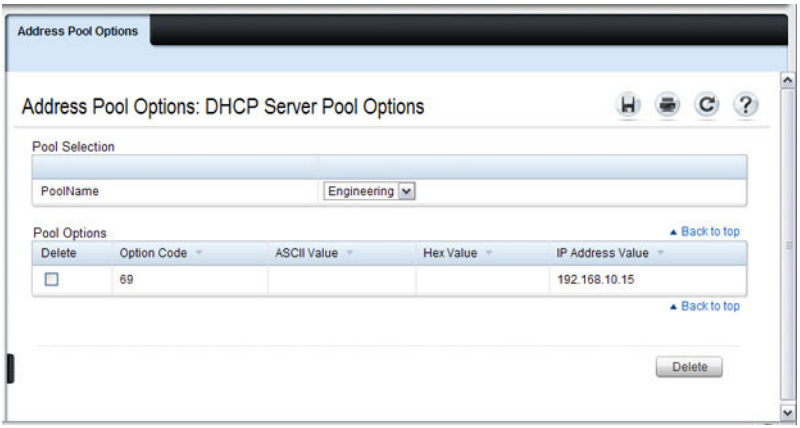
Figure 32-9. Add DHCP Option



5 Click **Apply**.

6 To verify that the option has been added to the address pool, open the **Address Pool Options** page.

Figure 32-10. View Address Pool Options

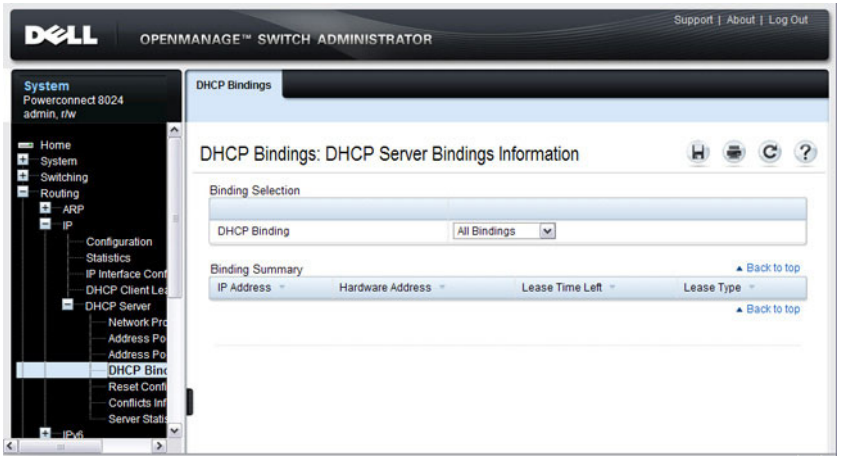


DHCP Bindings

Use the **DHCP Bindings** page to view information about the clients that have leased IP addresses from the DHCP server.

To display the **DHCP Bindings** page, click **Routing** → **IP** → **DHCP Server** → **DHCP Bindings** in the navigation panel.

Figure 32-11. DHCP Bindings

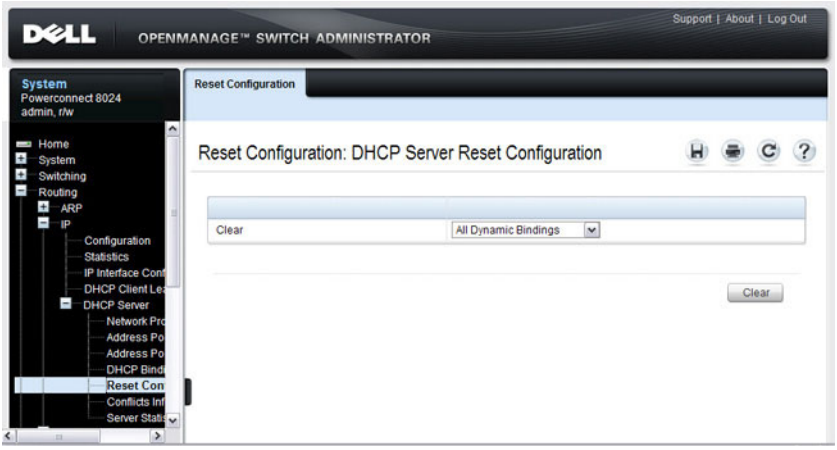


DHCP Server Reset Configuration

Use the **Reset Configuration** page to clear the client bindings for one or more clients. You can also reset bindings for clients that have leased an IP address that is already in use on the network.

To display the **Reset Configuration** page, click **Routing** → **IP** → **DHCP Server** → **Reset Configuration** in the navigation panel.

Figure 32-12. Reset DHCP Bindings

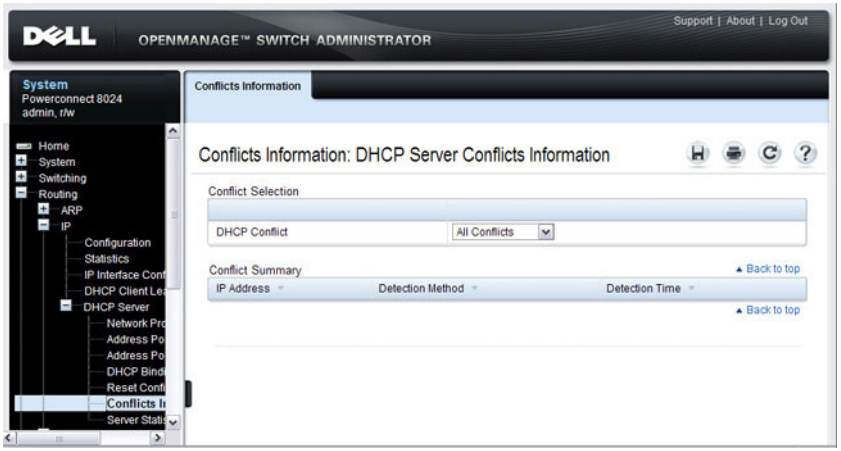


DHCP Server Conflicts Information

Use the **Conflicts Information** page to view information about clients that have leased an IP address that is already in use on the network.

To display the **Conflicts Information** page, click **Routing** → **IP** → **DHCP Server** → **Conflicts Information** in the navigation panel.

Figure 32-13. DHCP Server Conflicts Information

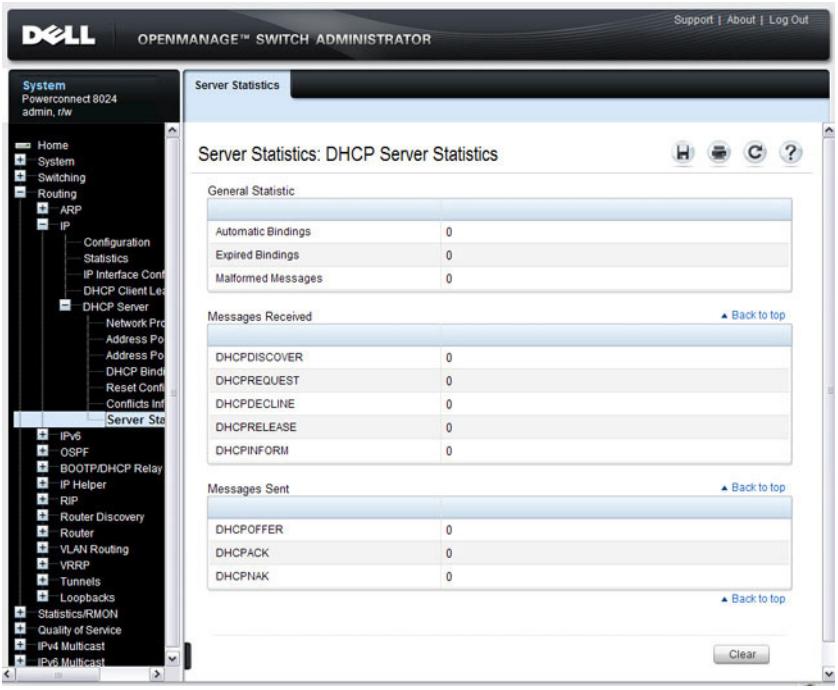


DHCP Server Statistics

Use the **Server Statistics** page to view general DHCP server statistics, messages received from DHCP clients, and messages sent to DHCP clients.

To display the **Server Statistics** page, click **Routing** → **IP** → **DHCP Server** → **Server Statistics** in the navigation panel.

Figure 32-14. DHCP Server Statistics



Configuring the DHCP Server (CLI)

This section provides information about the commands you use to configure and monitor the DHCP server and address pools. For more information about the commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Configuring Global DHCP Server Settings

Beginning in Privileged EXEC mode, use the following commands to configure settings for the DHCP server.

Command	Purpose
configure	Enter Global Configuration mode.
service dhcp	Enable the DHCP server.
ip dhcp ping packets	Specify the number, in a range from 2-10, of packets a DHCP server sends to a pool address as part of a ping operation.
ip dhcp conflict logging	Enable conflict logging on DHCP server
ip dhcp bootp automatic	Enable the allocation of the addresses to the BootP client.
ip dhcp excluded-address *lowaddress* [*highaddress*]	Specify the IP addresses that a DHCP server should not assign to DHCP clients. You can specify a single IP address, or you can specify a contiguous range by using both the low-address and high-address variables.
exit	Exit to Privileged EXEC mode.
show ip dhcp global configuration	Verify the global DHCP server configuration.

Configuring a Dynamic Address Pool

Beginning in Privileged EXEC mode, use the following commands to create an address pool with network information that is dynamically assigned to hosts with DHCP clients that request the information.

Command	Purpose
configure	Enter Global Configuration mode.
ip dhcp pool *name*	Create a DHCP address pool and enters DHCP pool configuration mode.
network *network-ip* [*mask* | *prefixlength*]	Configure the subnet number and mask for a DHCP address pool. Clients requesting an IP address can be assigned any non-excluded IP address within this network.
lease [*duration*] | **infinite**}]	Specify the duration of the lease for an IP address that is assigned from a DHCP server to a DHCP client. • *duration*— Days the lease is valid. You can optionally specify the hours and minutes after specifying the days. • **infinite** — 60 day lease
default-router *address1* [*address2....address8*]	Specify the list of default gateway IP addresses to be assigned to the DHCP client.
dns-server *address1* [*address2....address8*]	Specify the list of DNS server IP addresses to be assigned to the DHCP client.
domain-name *domain*	Specify the domain name for a DHCP client.
option *code* {**ascii** *string* | **hex** *string1* [*string2...string8*] | **ip** *address1* [*address2...address8*]}	Manually configure DHCP options.
CTRL + Z	Exit to Privileged EXEC mode.
show ip dhcp pool configuration {*name* | **all**}	View the settings for the specified address pool or for all configured address pools.

Configuring a Static Address Pool

Beginning in Privileged EXEC mode, use the following commands to create a static address pool and specify the network information for the pool. The network information configured in the static address pool is assigned only to the host with the hardware address or client identifier that matches the information configured in the static pool.

Command	Purpose
configure	Enter Global Configuration mode.
ip dhcp pool *name*	Create a DHCP address pool and enters DHCP pool configuration mode.
client-name *name*	Specify the DHCP client name.
hardware-address *mac* [*type*]	Specify the hardware address of the client in the static pool. • *mac*—MAC address of the hardware platform of the client consisting of 6 bytes in dotted hexadecimal format. • *type* — Indicates the protocol of the hardware platform. It is 1 for Ethernet and 6 for IEEE 802.
client-identifier *uniqueidentifier*	Specify the unique identifier for a DHCP client. The unique-identifier is a valid notation in hexadecimal format. In some systems, such as Microsoft DHCP clients, the client identifier is required instead of hardware addresses. The unique-identifier is a concatenation of the media type and the MAC address. For example, the Microsoft client identifier for Ethernet address c819.2488.f177 is 01c8.1924.88f1.77 where 01 represents the Ethernet media type.
host *address* [*mask* | *prefix-length*]	Specify the IP address and (optionally) network mask for a manual binding to a DHCP client.
lease [*duration*] | **infinite**}]	Specify the duration of the lease for an IP address that is assigned from a DHCP server to a DHCP client. • *duration*— Days the lease is valid. You can optionally specify the hours and minutes after specifying the days. • **infinite** — 60 day lease

Command	Purpose
default-router *address1* [*address2....address8*]	Specify the list of default gateway IP addresses to be assigned to the DHCP client.
dns-server *address1* [*address2....address8*]	Specify the list of DNS server IP addresses to be assigned to the DHCP client.
domain-name *domain*	Specify the domain name for a DHCP client.
option *code* {**ascii** *string* | **hex** *string1* [*string2...string8*] | **ip** *address1* [*address2...address8*]}	Manually configure DHCP options.
CTRL + Z	Exit to Privileged EXEC mode.
show ip dhcp pool configuration {*name* | **all**}	View the settings for the specified address pool or for all configured address pools.

Monitoring DHCP Server Information

Beginning in Privileged EXEC mode, use the following commands to view bindings, conflicts, and statistics, and to clear the information.

Command	Purpose
show ip dhcp binding [*address*]	View the current binding information in the DHCP server database. Specify the IP address to view a specific binding.
clear ip dhcp binding {*address* | *}	Delete an automatic address binding from the DHCP server database. Use * to clear all bindings.
show ip dhcp conflict [*address*]	View the current binding conflicts in the DHCP server database. Specify the IP address to view a specific conflict.
clear ip dhcp conflict {*address* | *}	Clear an address conflict from the DHCP Server database. Use * to clear all conflicts.
show ip dhcp server statistics	View DHCP server statistics.
clear ip dhcp server statistics	Reset all DHCP server statistics to zero.

DHCP Server Configuration Examples

This section contains the following examples:

- Configuring a Dynamic Address Pool
- Configuring a Static Address Pool

Configuring a Dynamic Address Pool

The commands in this example create an address pool that dynamically assigns network information to hosts with DHCP clients that broadcast DHCP messages. The hosts are assigned an IP address from the 192.168.5.0 network. The IP addresses 192.168.5.1–192.168.5.20, and 192.168.5.100 are excluded from the address pool.

To configure the switch:

1. Create an address pool named "Engineering" and enter into DHCP pool configuration mode for the pool.

   ```
   console#configure
   console(config)#ip dhcp pool Engineering
   ```

2. Specify the IP addresses that are available in the pool.

   ```
   console(config-dhcp-pool)#network 192.168.5.0
   255.255.255.0
   ```

3. Specify the IP address to use as the default gateway.

   ```
   console(config-dhcp-pool)#default-router
   192.168.5.1
   ```

4. Specify the primary and secondary DNS servers the hosts will use.

   ```
   console(config-dhcp-pool)#dns-server 192.168.5.10
   console(config-dhcp-pool)#dns-server 192.168.5.11
   ```

5. Specify the domain name to be assigned to clients that lease an address from this pool.

   ```
   console(config-dhcp-pool)#domain-name
   engineering.dell.com
   console(config-dhcp-pool)#exit
   ```

6 In Global Configuration mode, add the addresses to exclude from the pool. Clients will not be assigned these IP addresses.

```
console(config)#ip dhcp excluded-address
192.168.5.1 192.168.5.20
console(config)#ip dhcp excluded-address
192.168.5.100
```

7 Enable the DHCP server on the switch.

```
console(config)#service dhcp
console(config)#exit
```

8 View DHCP server settings.

```
console#show ip dhcp global configuration

Service DHCP...................Enable
Number of Ping Packets.........2
Excluded Address...............192.168.2.1 to 192.168.2.20
                               1.2.2.2 to 1.5.5.5
                               192.168.5.1 to 192.168.5.20
                               192.168.5.100 to 192.168.5.100
Conflict Logging...............Enable
Bootp Automatic................Disable
```

9 View information about all configured address pools.

```
console#show ip dhcp pool configuration all

Pool: Engineering
Pool Type..........................  Network
Network............................  192.168.5.0 255.255.255.0
Lease Time.........................  1 days 0 hrs 0 mins
DNS Servers........................  192.168.5.11
Default Routers....................  192.168.5.1
Domain Name........................  engineering.dell.com
```

Configuring a Static Address Pool

The commands in this example create an address pool that assigns the address 192.168.2.10 to the host with a MAC address of 00:1C:23:55:E9:F3. When this hosts sends a DHCP message requesting network information, the switch will offer the information configured in this example, which includes a custom DHCP option to assign the SMTP server IP address.

To configure the switch:

1. Create an address pool named "Tyler PC" and enter into DHCP pool configuration mode for the pool.

   ```
   console#configure
   console(config)#ip dhcp pool "Tyler PC"
   ```

2. Specify the IP addresses that are available in the pool.

   ```
   console(config-dhcp-pool)#hardware-address
   00:1C:23:55:E9:F3
   ```

3. Specify the IP address and subnet mask to assign to the client.

   ```
   console(config-dhcp-pool)#host 192.168.2.10
   255.255.255.0
   ```

4. Specify the IP address to use as the default gateway.

   ```
   console(config-dhcp-pool)#default-router
   192.168.2.1
   ```

5. Specify the primary and secondary DNS servers the hosts will use.

   ```
   console(config-dhcp-pool)#dns-server
   192.168.2.100
   console(config-dhcp-pool)#dns-server
   192.168.5.101
   ```

6. Specify the domain name to be assigned to clients that lease an address from this pool.

   ```
   console(config-dhcp-pool)#domain-name
   executive.dell.com
   ```

7. Specify the option that configures the SMTP server IP address to the host.

   ```
   console(config-dhcp-pool)#option 69 ip
   192.168.1.33
   console(config-dhcp-pool)#exit
   ```

8 View information about the static address pool.

```
console#show ip dhcp pool configuration "Tyler PC"

Pool: Tyler PC
Pool Type.....................Static
Client Name.....................TylerPC
Hardware Address.................. 00:1c:23:55:e9:f3
Hardware Address Type.............ethernet
Host.............................. 192.168.2.10 255.255.255.0
Lease Time........................ 1 days 0 hrs 0 mins
DNS Servers....................... 192.168.2.101
Default Routers................... 192.168.2.1
Domain Name....................... executive.dell.com
Option............................ 69 ip 192.168.1.33
```

33

Configuring IP Routing

This chapter describes how to configure routing on the switch, including global routing settings, Address Resolution Protocol (ARP), router discovery, and static routes.

The topics covered in this chapter include:

- IP Routing Overview
- Default IP Routing Values
- Configuring IP Routing Features (Web)
- Configuring IP Routing Features (CLI)
- IP Routing Configuration Example

IP Routing Overview

The PowerConnect 8000-series and 8100-series switches are multilayer switches that support static and dynamic routing. Table 33-1 describes some of the general routing features that you can configure on the switch.

Table 33-1. IP Routing Features

Feature	Description
ICMP message control	You can configure the type of ICMP messages that the switch responds to as well as the rate limit and burst size.
Default gateway	The switch supports a single default gateway. A manually configured default gateway is more preferable than a default gateway learned from a DHCP server.
ARP table	The switch maintains an ARP table that maps an IP address to a MAC address. You can create static ARP entries in the table and manage various ARP table settings such as the aging time of dynamically-learned entries.

Table 33-1. IP Routing Features (Continued)

Feature	Description
ICMP Router Discovery Protocol (IRDP)	Hosts can use IRDP to identify operational routers on the subnet. Routers periodically advertise their IP addresses. Hosts listen for these advertisements and discover the IP addresses of neighboring routers.
Routing table entries	You can configure the following route types in the routing table: • Default: The default route is the route the switch will use to send a packet if the routing table does not contain a longer matching prefix for the packet's destination. • Static: A static route is a route that you manually add to the routing table. • Static Reject: Packets that match a reject route are discarded instead of forwarded. The router may send an ICMP Destination Unreachable message.
Route preferences	The common routing table collects static, local, and dynamic (routing protocol) routes. When there is more than one route to the same destination prefix, the routing table selects the route with the best (lowest) route preference.

Default IP Routing Values

Table 33-2 shows the default values for the IP routing features this chapter describes.

Table 33-2. IP Routing Defaults

Parameter	Default Value
Default Time to Live	64
Routing Mode	Disabled globally and on each interface
ICMP Echo Replies	Enabled
ICMP Redirects	Enabled
ICMP Rate Limit Interval	1000 milliseconds
ICMP Rate Limit Burst Size	100
Maximum Next Hops	4
Global Default Gateway	None
Dynamic ARP Entry Age Time	1200 seconds
Automatic Renewal of Dynamic ARP Entries	Disabled
ARP Response Timeout	1 second
ARP Retries	4
Maximum Static ARP Entries	128
IRDP Advertise Mode	Disabled
IRDP Advertise Address	224.0.0.1
IRDP Maximum Advertise Interval	600 seconds
IRDP Minimum Advertise Interval	450 seconds
IRDP Advertise Lifetime	1800 seconds
IRDP Preference Level	0

Table 33-2. IP Routing Defaults (Continued)

Parameter	Default Value
Route Preference Values	Preference values are as follows: • Local—0 • Static—1 • OSPF Intra—110 • OSPF Inter—110 • OSPF External—110 • RIP—120

Configuring IP Routing Features (Web)

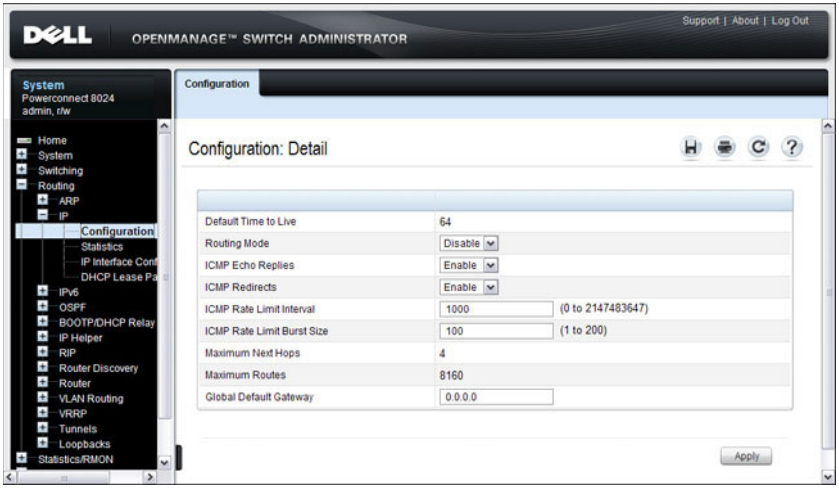
This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring IPv4 routing features on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

IP Configuration

Use the **Configuration** page to configure routing parameters for the switch as opposed to an interface. The IP configuration settings allow you to enable or disable the generation of various types of ICMP messages.

To display the page, click **Routing → IP → Configuration** in the navigation panel.

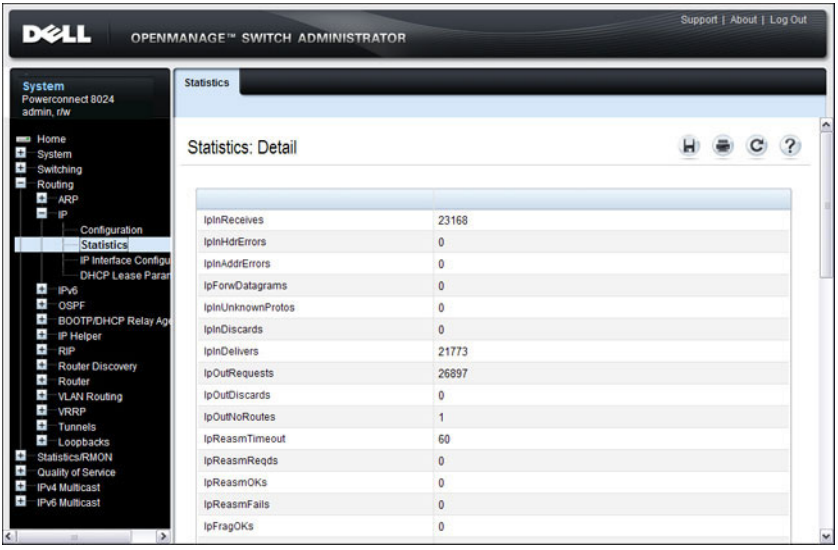
Figure 33-1. IP Configuration



IP Statistics

The IP statistics reported on the **Statistics** page are as specified in RFC 1213. To display the page, click **Routing** → **IP** → **Statistics** in the navigation panel.

Figure 33-2. IP Statistics

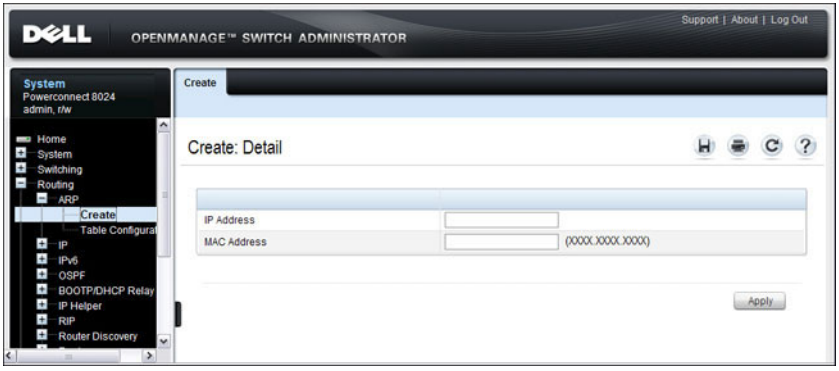


ARP Create

Use the **Create** page to add a static ARP entry to the Address Resolution Protocol table.

To display the page, click **Routing** → **ARP** → **Create** in the navigation panel.

Figure 33-3. ARP Create

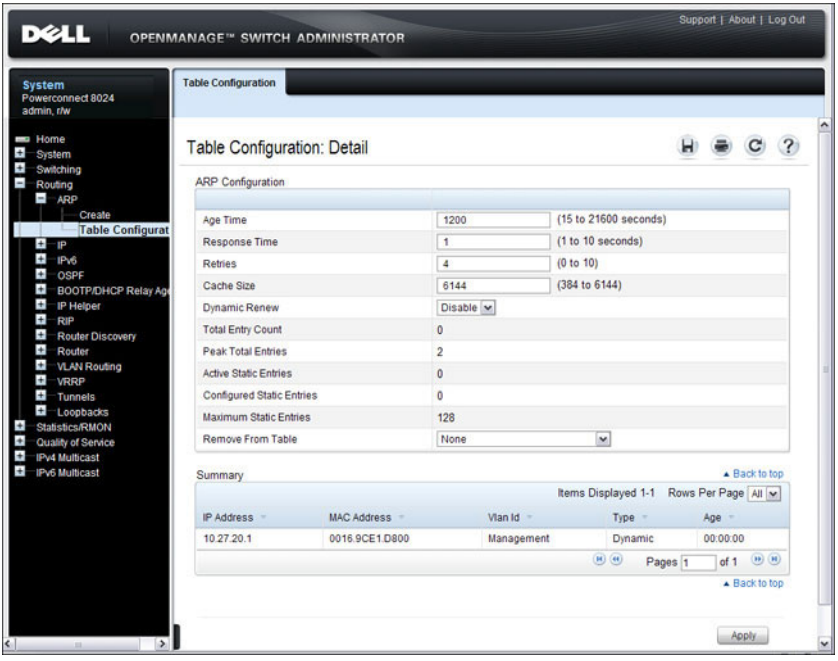


ARP Table Configuration

Use the **Table Configuration** page to change the configuration parameters for the Address Resolution Protocol Table. You can also use this screen to display the contents of the table.

To display the page, click **Routing** → **ARP** → **Table Configuration** in the navigation panel.

Figure 33-4. ARP Table Configuration

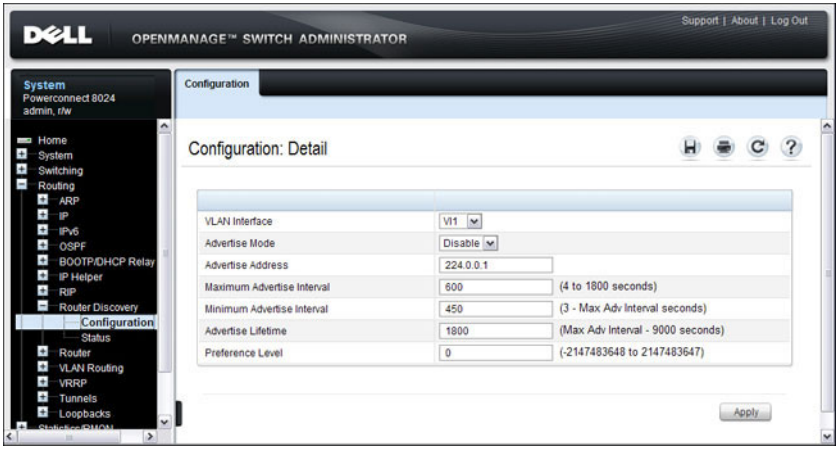


Router Discovery Configuration

Use the **Configuration** page to enter or change router discovery parameters.

To display the page, click **Routing → Router Discovery → Configuration** in the navigation panel.

Figure 33-5. Router Discovery Configuration

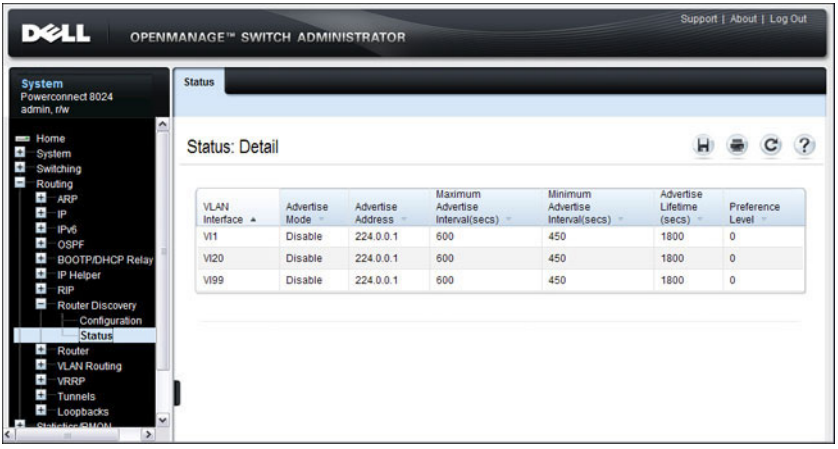


Router Discovery Status

Use the **Status** page to display router discovery data for each interface.

To display the page, click **Routing** → **Router Discovery** → **Status** in the navigation panel.

Figure 33-6. Router Discovery Status

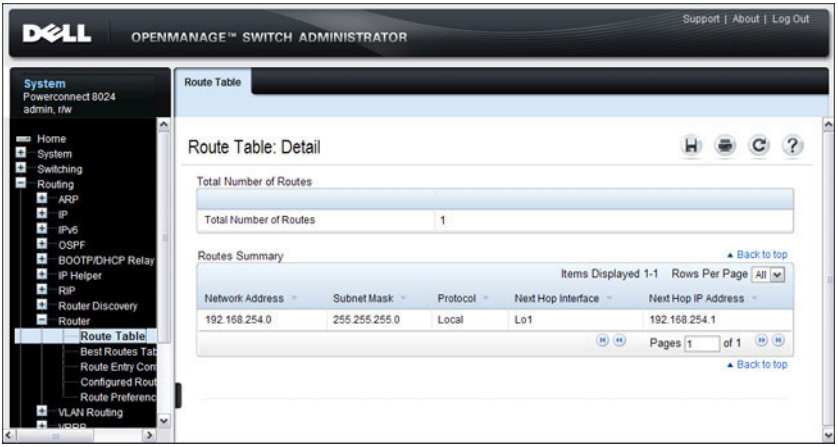


Route Table

Use the **Route Table** page to display the contents of the routing table.

To display the page, click **Routing** → **Router** → **Route Table** in the navigation panel.

Figure 33-7. Route Table

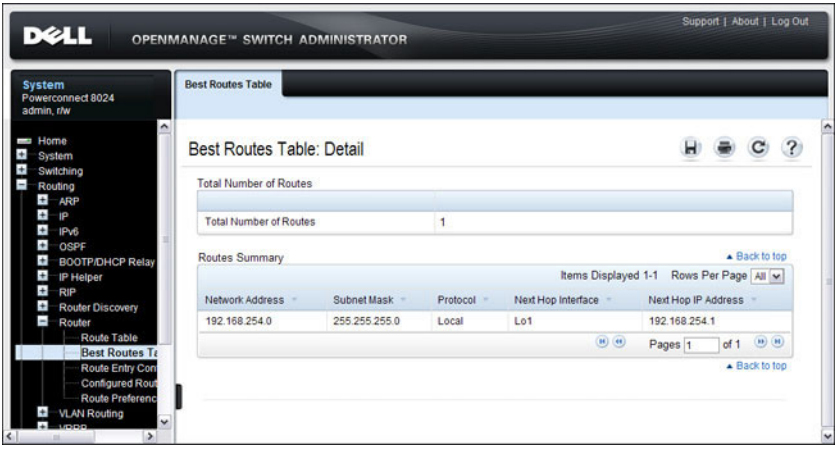


Best Routes Table

Use the **Best Routes Table** page to display the best routes from the routing table.

To display the page, click **Routing** → **Router** → **Best Routes Table** in the navigation panel.

Figure 33-8. Best Routes Table

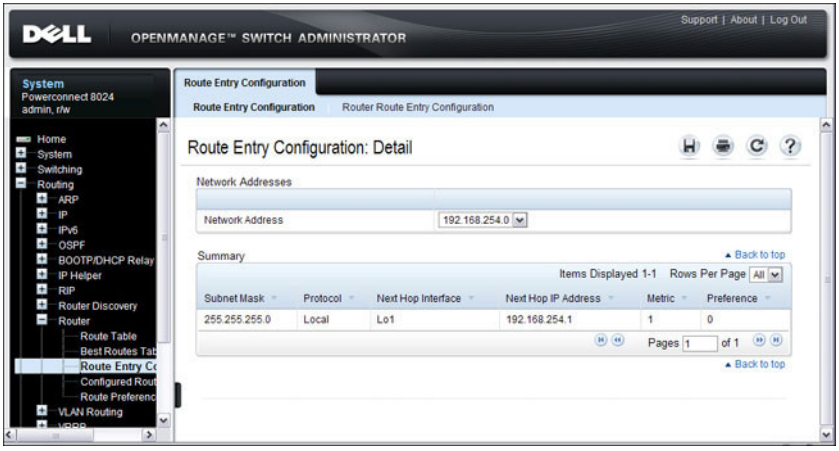


Route Entry Configuration

Use the **Route Entry Configuration** page to add new and configure router routes.

To display the page, click **Routing** → **Router** → **Route Entry Configuration** in the navigation panel.

Figure 33-9. Route Entry Configuration



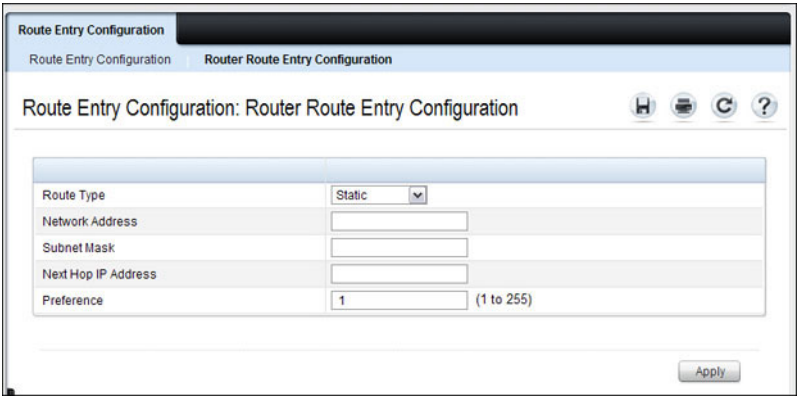
Adding a Route and Configuring Route Preference

To configure routing table entries:

1. Open the **Route Entry Configuration** page.
2. Click **Router Route Entry Configuration**.

   The screen refreshes and the **Router Route Entry Configuration** page displays.

Figure 33-10. Router Route Entry and Preference Configuration



3 Next to **Route Type**, use the drop-down box to add a **Default, Static, or Static Reject** route.

The fields to configure are different for each route type.

- **Default** — Enter the default gateway address in the **Next Hop IP Address** field.
- **Static** — Enter values for **Network Address**, **Subnet Mask**, **Next Hop IP Address**, and **Preference**.
- **Static Reject** — Enter values for **Network Address**, **Subnet Mask**, and **Preference**.

4 Click **Apply**.

The new route is added to the routing table.

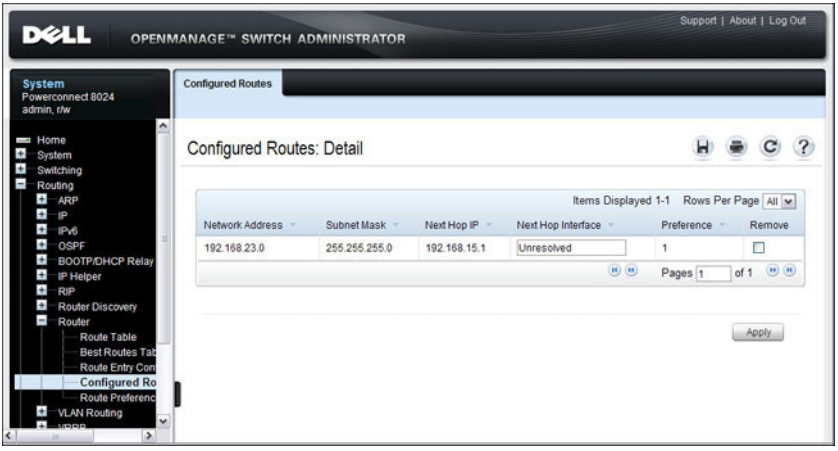
Configured Routes

Use the **Configured Routes** page to display the routes that have been manually configured.

NOTE: For a static reject route, the next hop interface value is Null0. Packets to the network address specified in static reject routes are intentionally dropped.

To display the page, click **Routing** → **Router** → **Configured Routes** in the navigation panel.

Figure 33-11. Configured Routes



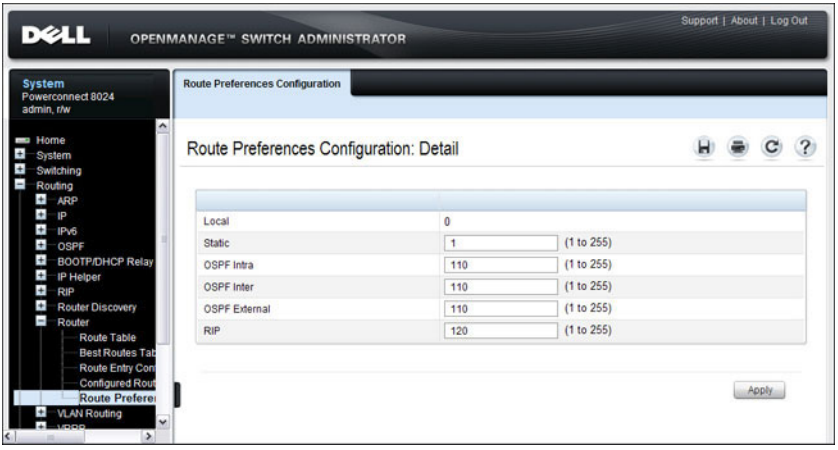
To remove a configured route, select the check box in the **Remove** column of the route to delete, and click **Apply**.

Route Preferences Configuration

Use the **Route Preferences Configuration** page to configure the default preference for each protocol (for example 60 for static routes). These values are arbitrary values that range from 1 to 255, and are independent of route metrics. Most routing protocols use a route metric to determine the shortest path known to the protocol, independent of any other protocol.

To display the page, click **Routing → Router → Route Preferences Configuration** in the navigation panel.

Figure 33-12. Router Route Preferences Configuration



Configuring IP Routing Features (CLI)

This section provides information about the commands you use to configure IPv4 routing on the switch. For more information about the commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Configuring Global IP Routing Settings

Beginning in Privileged EXEC mode, use the following commands to configure various global IP routing settings for the switch.

Command	Purpose
configure	Enter global configuration mode.
ip routing	Globally enable IPv4 routing on the switch.
ip icmp echo-reply	Allow the switch to generate ICMP Echo Reply messages.
ip icmp error-interval *burst-interval* [*burst-size*]	Limit the rate at which IPv4 ICMP error messages are sent. • *burst-interval* — How often the token bucket is initialized (Range: 0–2147483647 milliseconds). • *burst-size* — The maximum number of messages that can be sent during a burst interval (Range: 1–200).
ip redirects	Allow the switch to generate ICMP Redirect messages.
ip default-gateway *ip-address*	Configure the global default gateway for the switch. The gateway configured here takes precedence over a default gateway assigned by a network DHCP server.
exit	Exit to Privileged EXEC mode.
show ip brief	View the global IP settings for the switch.

Adding Static ARP Entries and Configuring ARP Table Settings

Beginning in Privileged EXEC mode, use the following commands to configure static ARP entries in the ARP cache and to specify the settings for the ARP cache.

Command	Purpose
configure	Enter global configuration mode.
arp *ip-address hardware-address*	Create a static ARP entry in the ARP table. • *ip-address* — IP address of a device on a subnet attached to an existing routing interface. • *hardware-address* — A unicast MAC address for that device.
arp timeout *seconds*	Configure the ARP entry ageout time.
arp resptime *seconds*	Configure the ARP request response timeout.
arp retries *integer*	Configure the ARP count of maximum requests for retries. The range is 1–10.
arp cachesize *integer*	Configure the maximum number of entries in the ARP cache.
arp dynamicrenew	Allow the ARP component to automatically renew dynamic ARP entries when they age out.
exit	Exit to Privileged EXEC mode.
show arp [**brief**]	View the user-configured (static) ARP entries. The static entries display regardless of whether they are reachable over an interface. Use the **brief** keyword to view only the ARP table settings.
clear arp-cache [**gateway**]	Remove all dynamic ARP entries from the ARP cache. Include the keyword **gateway** to remove gateway entries as well.
clear arp-cache management	Remove all dynamic ARP entries from the ARP cache that were learned on the management interface.
arp purge *ip-address*	Remove the specified IP address from the ARP cache. This command removes dynamic and gateway ARP entries only.

Configuring Router Discovery (IRDP)

Beginning in Privileged EXEC mode, use the following commands to configure IRDP settings.

Command	Purpose
configure	Enter global configuration mode.
interface *interface*	Enter interface configuration mode for the specified VLAN routing interface. The *interface* variable includes the interface type (**vlan**) and number, for example **vlan 100**.
ip irdp	Enable IRDP on the interface.
ip irdp address *ip-address*	Configure the address that the interface uses to send the router discovery advertisements. The allowed addresses are 224.0.0.1 (all-hosts IP multicast address) or 255.255.255.255 (limited broadcast address)
ip irdp holdtime *seconds*	Configure the value of the holdtime field of the router advertisement sent from this interface.
ip irdp maxadvertinterval *seconds*	Configure the maximum time allowed between sending router advertisements from the interface.
ip irdp minadvertinterval *seconds*	Configure the minimum time allowed between sending router advertisements from the interface.
ip irdp preference *integer*	Configure the preference of the address as a default router address relative to other router addresses on the same subnet.
exit	Exit to Global Config mode.
exit	Exit to Privileged EXEC mode.
show ip irdp [**vlan** *vlan-id*]	View the router discovery information for all interfaces, or for a specified interface.

Configuring Route Table Entries and Route Preferences

Beginning in Privileged EXEC mode, use the following commands to configure IRDP settings.

Command	Purpose
configure	Enter global configuration mode.
ip route default *nextHopRtr* [*preference*]	Configure the default route. • *nextHopRtr* — IP address of the next hop router. • *preference* — Specifies the preference value (administrative distance) of an individual static route. (Range: 1-255)
ip route *ip-addr* {*subnetmask* | *prefix length* } {*nextHopRtr* / **null***}* [*preference*]	Configure a static route. Use the keyword **null** instead of the next hop router IP address to configure a static reject route. • *ip-address* — IP address of destination interface. • *subnet-mask* — Subnet mask of destination interface. • *prefix-length* — Length of prefix. Must be preceded with a forward slash (/). (Range: 0-32 bits) • *nextHopRtr* — IP address of the next hop router. • **null** — Specifies that the route is a static reject route. • *preference* — Specifies the preference value (administrative distance) of an individual static route. (Range: 1-255)
ip route distance *integer*	Set the default distance (preference) for static routes. Lower route preference values are preferred when determining the best route.
exit	Exit to Privileged EXEC mode.

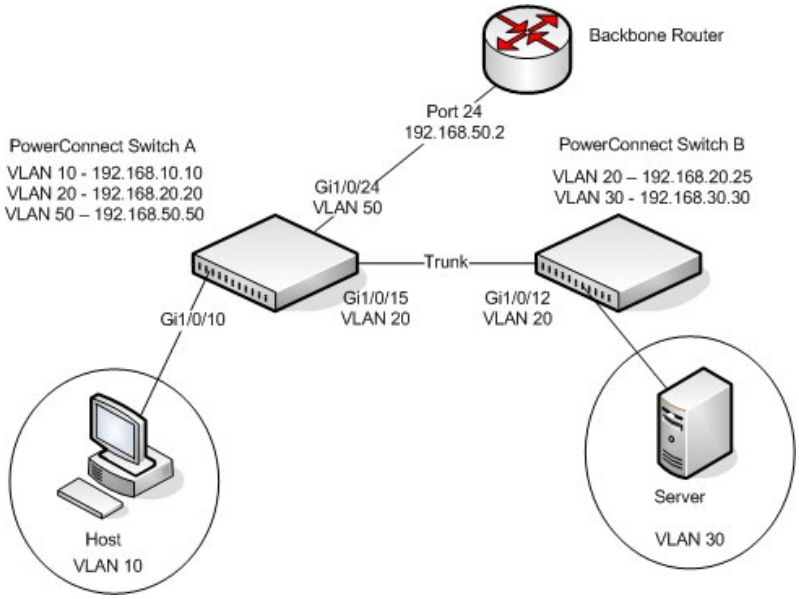
Command	Purpose
show ip route [*ip-address* [*mask* | *prefix-length*] [**longer-prefixes**] | *protocol*]	View the routing table. • *ip-address* — Specifies the network for which the route is to be displayed and displays the best matching best-route for the address. • *mask* — Subnet mask of the IP address. • *prefix-length* — Length of prefix, in bits. Must be preceded with a forward slash ('/'). (Range: 0-32 bits) • **longer-prefixes** — Indicates that the *ip-address* and *subnet-mask* pair becomes the prefix, and the command displays the routes to the addresses that match that prefix. • *protocol* — Specifies the protocol that installed the routes. (Range: **connected, ospf, rip, static**)
show ip route configured	View the configured routes, whether they are reachable or not.
show ip route summary	View summary information about the routing table.
show ip protocols	View the parameters and current state of the active routing protocols.
show ip route preferences	View detailed information about the route preferences.

IP Routing Configuration Example

In this example, the PowerConnect switches are L3 switches with VLAN routing interfaces. VLAN routing is configured on PowerConnect Switch A and PowerConnect Switch B. This allows the host in VLAN 10 to communicate with the server in VLAN 30. A static route to the VLAN 30 subnet is configured on Switch A. Additionally, a default route is configured on Switch A so that all traffic with an unknown destination is sent to the backbone router through port 24, which is a member of VLAN 50. A default route is configured on PowerConnect Switch B to use Switch A as the default gateway. The hosts use the IP address of the VLAN routing interface as their default gateway.

This example assumes that all L2 VLAN information, such as VLAN creation and port membership, has been configured.

Figure 33-13. IP Routing Example Topology



Configuring PowerConnect Switch A

To configure Switch A.

1. Enable routing on the switch.

```
console#configure
console(config)#ip routing
```

2. Assign an IP address to VLAN 10. This command also enables IP routing on the VLAN.

```
console(config)#interface vlan 10
console(config-if-vlan10)#ip address 192.168.10.10
255.255.255.0
console(config-if-vlan10)#exit
```

3. Assign an IP address to VLAN 20.

```
console#configure
console(config)#interface vlan 20
console(config-if-vlan20)#ip address 192.168.20.20
255.255.255.0
console(config-if-vlan20)#exit
```

4. Assign an IP address to VLAN 50.

```
console#configure
console(config)#interface vlan 50
console(config-if-vlan50)#ip address 192.168.50.50
255.255.255.0
console(config-if-vlan50)#exit
```

5. Configure a static route to the network that VLAN 30 is in, using the IP address of the VLAN 20 interface on Switch B as the next hop address.

```
console(config)#ip route 192.168.30.0
255.255.255.0 192.168.20.25
```

6. Configure the backbone router interface as the default gateway.

```
console(config)#ip route default 192.168.50.2
```

Configuring PowerConnect Switch B

To configure Switch B:

1. Enable routing on the switch.

```
console#configure
console(config)#ip routing
```

2. Assign an IP address to VLAN 20. This command also enables IP routing on the VLAN.

```
console#configure
console(config)#interface vlan 20
console(config-if-vlan20)#ip address 192.168.20.25
255.255.255.0
console(config-if-vlan20)#exit
```

3. Assign an IP address to VLAN 30. This command also enables IP routing on the VLAN.

```
console#configure
console(config)#interface vlan 30
console(config-if-vlan30)#ip address 192.168.30.30
255.255.255.0
console(config-if-vlan30)#exit
```

4. Configure the VLAN 20 routing interface on Switch A as the default gateway so that any traffic with an unknown destination is sent to Switch A for forwarding.

```
console(config)#ip route default 192.168.20.20
```

34

Configuring L2 and L3 Relay Features

This chapter describes how to configure the L2 DHCP Relay, L3 DHCP Relay, and IP Helper features on PowerConnect 8000-series and 8100-series switches.

The topics covered in this chapter include:

- L2 and L3 Relay Overview
- Default L2/L3 Relay Values
- Configuring L2 and L3 Relay Features (Web)
- Configuring L2 and L3 Relay Features (CLI)
- Relay Agent Configuration Example

L2 and L3 Relay Overview

When a DHCP client and server are in the same IP subnet, they can directly connect to exchange IP address requests and replies. However, buying and maintaining a DHCP server on each subnet can be expensive and is often impractical. The relay features on the PowerConnect 8000-series and 8100-series switches can help enable communication between DHCP clients and DHCP servers that reside in different subnets. Configuring L3 DHCP relay also enables the bootstrap protocol (BOOTP) relay.

What Is L3 DHCP Relay?

Network infrastructure devices can be used to relay packets between a DHCP client and server on different subnets. Such a device, an Layer 3 Relay agent, is often a router or L3 switch. The L3 relay agent must have an IP interface on the client subnets and, if it does not have an IP interface on the server's subnet, it should be able to route traffic toward the server's subnet.

The PowerConnect DHCP Relay Agent enables DHCP clients and servers to exchange DHCP messages across different subnets. The relay agent receives the requests from the clients, and checks the valid hops and *giaddr* fields in the DHCP request. If the number of hops is greater than the configured number, the agent discards the packet. If the *giaddr* field is zero, the agent must fill in this field with the IP address of the interface on which the request was received. The agent unicasts the valid packets to all configured DHCP servers. Each server responds with a unicast BOOTREPLY addressed to the relay agent closest to the client as indicated by *giaddr* field. Upon reception of the BOOTREPLY from the server, the agent forwards this reply as broadcast or unicast on the interface where the BOOTREQUEST arrived. This interface can be identified by the *giaddr* field or option 82.

The PowerConnect 8000-series and 8100-series switches DHCP component also supports DHCP relay agent options to identify the client interface. If configured, the relay agent inserts these options when forwarding the request to the server and removes them when sending the reply to the clients.

If an interface has more than one IP address, the relay agent uses the primary IP address configured as its relay agent IP address.

What Is L2 DHCP Relay?

In Layer 2 switched networks, there may be one or more infrastructure devices (for example, a switch) between the client and the L3 Relay agent/DHCP server. In this instance, some of the client device information required by the L3 Relay agent may not be visible to it. In this case, an L2 Relay agent can be used to add the information that the L3 Relay Agent and DHCP server need to perform their roles in address and configuration and assignment.

Before it relays DHCP requests from clients, the switch can add a Circuit ID and a Remote ID. These IDs provide information about the circuit and port number connected to the client. This information is added as suboptions in the DHCP Option 82 packets as defined in sections 3.1 and 3.2 of RFC3046. The switch removes this option from packets that it relays from L3 Relay agents/DHCP servers to clients.

These sub-options may be used by the DHCP server to affect how it treats the client and also may be used by the relay agent to limit broadcast replies to the specific circuit or attachment point of the client.

Enabling L2 Relay on VLANs

You can enable L2 DHCP relay on a particular VLAN. The VLAN is identified by a service VLAN ID (S-VID), which a service provider uses to identify a customer's traffic while traversing the provider network to multiple remote sites. The switch uses the VLAN membership of the switch port client (the customer VLAN ID, or C-VID) to perform a lookup a corresponding S-VID.

If the S-VID is enabled for DHCP Relay, then the packet can be forwarded. If the C-VID does not correspond to an S-VID that is enabled for DHCP Relay, then the switch will not relay the DHCP request packet.

What Is the IP Helper Feature?

The IP Helper feature provides the ability for a router to forward configured UDP broadcast packets to a particular IP address. This allows applications to reach servers on non-local subnets. This is possible even when the application is designed to assume a server is always on a local subnet or when the application uses broadcast packets to reach the server (with the limited broadcast address 255.255.255.255, or a network directed broadcast address).

You can configure relay entries globally and on routing interfaces. Each relay entry maps an ingress interface and destination UDP port number to a single IPv4 address (the helper address). Multiple relay entries may be configured for the same interface and UDP port, in which case the relay agent relays matching packets to each server address. Interface configuration takes priority over global configuration. If the destination UDP port for a packet matches any entry on the ingress interface, the packet is handled according to the interface configuration. If the packet does not match any entry on the ingress interface, the packet is handled according to the global IP helper configuration.

You can configure discard relay entries. Discard entries are used to discard packets received on a specific interface when those packets would otherwise be relayed according to a global relay entry. Discard relay entries may be configured on interfaces, but are not configured globally.

Additionally, you can configure which UDP ports are forwarded. Certain UDP port numbers can be selected from the web interface or specified by name in the CLI, but you can also configure a relay entry with any UDP port number. You may configure relay entries that do not specify a destination UDP port. The relay agent assumes that these entries match packets with the UDP destination ports listed in Table 34-1 (the list of default ports).

Table 34-1. Default Ports - UDP Port Numbers Implied By Wildcard

Protocol	UDP Port Number
IEN-116 Name Service	42
DNS	53
NetBIOS Name Server	137
NetBIOS Datagram Server	138
TACACS Server	49
Time Service	37
DHCP	67
Trivial File Transfer Protocol	69

The system limits the number of relay entries to four times the maximum number of routing interfaces (512 relay entries). There is no limit to the number of relay entries on an individual interface, and no limit to the number of servers for a given {interface, UDP port} pair.

Certain configurable DHCP relay options do not apply to relay of other protocols. You may optionally set a maximum hop count or minimum wait time using the **bootpdhcprelay maxhopcount** and **bootpdhcprelay minwaittime** commands.

The relay agent relays DHCP packets in both directions. It relays broadcast packets from the client to one or more DHCP servers, and relays packets to the client that the DHCP server unicasts back to the relay agent. For other protocols, the relay agent only relays broadcast packets from the client to the server. Packets from the server back to the client are assumed to be unicast directly to the client. Because there is no relay in the return direction for protocols other than DHCP, the relay agent retains the source IP address from the original client packet. The relay agent uses a local IP address as the source IP address of relayed DHCP client packets.

When a switch receives a broadcast UDP packet on a routing interface, the relay agent verifies that the interface is configured to relay to the destination UDP port. If so, the relay agent unicasts the packet to the configured server IP addresses. Otherwise, the relay agent verifies that there is a global

configuration for the destination UDP port. If so, the relay agent unicasts the packet to the configured server IP addresses. Otherwise the packet is not relayed.

NOTE: If the packet matches a discard relay entry on the ingress interface, the packet is not forwarded, regardless of the global configuration.

The relay agent relays packets that meet only the following conditions:

- The destination MAC address must be the all-ones broadcast address (FF:FF:FF:FF:FF:FF).
- The destination IP address must be the limited broadcast address (255.255.255.255) or a directed broadcast address for the receive interface.
- The IP time-to-live (TTL) must be greater than 1.
- The protocol field in the IP header must be UDP (17).
- The destination UDP port must match a configured relay entry.

Table 34-2 shows the most common protocols and their UDP port numbers and names that are relayed.

Table 34-2. UDP Port Allocations

UDP Port Number	Acronym	Application
7	Echo	Echo
11	SysStat	Active User
15	NetStat	NetStat
17	Quote	Quote of the day
19	CHARGEN	Character Generator
20	FTP-data	FTP Data
21	FTP	FTP
37	Time	Time
42	NAMESERVER	Host Name Server
43	NICNAME	Who is
53	DOMAIN	Domain Name Server
69	TFTP	Trivial File Transfer
111	SUNRPC	Sun Microsystems Rpc
123	NTP	Network Time
137	NetBiosNameService	NT Server to Station Connections
138	NetBiosDatagramService	NT Server to Station Connections
139	NetBios	SessionServiceNT Server to Station Connections
161	SNMP	Simple Network Management
162	SNMP-trap	Simple Network Management Traps
513	who	Unix Rwho Daemon
514	syslog	System Log
525	timed	Time Daemon

Default L2/L3 Relay Values

By default L2 DHCP relay is disabled. L3 relay (UDP) is enabled, but no UDP destination ports or server addresses are defined on the switch or on any interfaces.

Table 34-3. L2/L3 Relay Defaults

Parameter	Default Value
L2 DHCP Relay	
Admin Mode	Disabled globally and on all interfaces and VLANs
Trust Mode	Disabled on all interfaces
Circuit ID	Disabled on all VLANs
Remote ID	None configured
L3 DHCP Relay	
UDP Relay Mode (IP Helper)	Enabled
Hop Count	4
Minimum Wait Time	0 seconds
Circuit ID Option Mode	Disabled
Circuit ID Check Mode	Enabled
Information Option-Insert	Disabled on all VLAN interfaces
Information Check-Reply	Enabled on all VLAN interfaces

Configuring L2 and L3 Relay Features (Web)

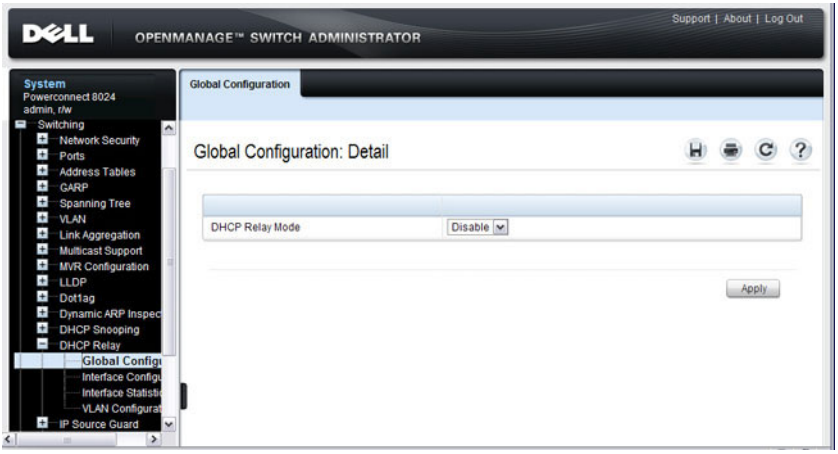
This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring L2 and L3 relay features on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

DHCP Relay Global Configuration

Use this page to enable or disable the switch to act as a DHCP Relay agent. This functionality must also be enabled on each port you want this service to operate on (see "DHCP Relay Interface Configuration" on page 915). The switch can also be configured to relay requests only when the VLAN of the requesting client corresponds to a service provider's VLAN ID that has been enabled with the L2 DHCP relay functionality (see "DHCP Relay VLAN Configuration" on page 918).

To access this page, click **Switching → DHCP Relay → Global Configuration** in the navigation panel.

Figure 34-1. DHCP Relay Global Configuration



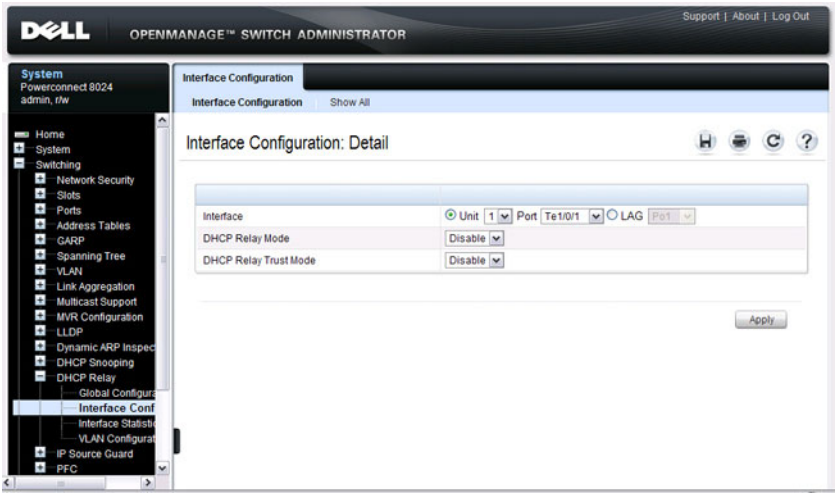
DHCP Relay Interface Configuration

Use this page to enable L2 DHCP relay on individual ports.

NOTE: L2 DHCP relay must also be enabled globally on the switch.

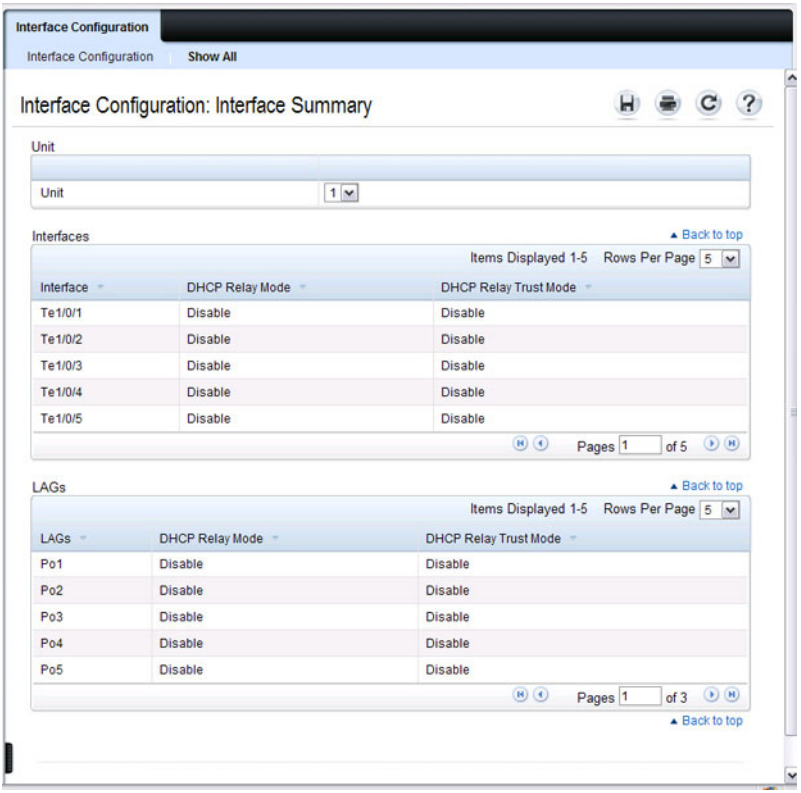
To access this page, click **Switching** → **DHCP Relay** → **Interface Configuration** in the navigation panel.

Figure 34-2. DHCP Relay Interface Configuration



To view a summary of the L2 DHCP relay configuration on all ports and LAGS, click **Show All**.

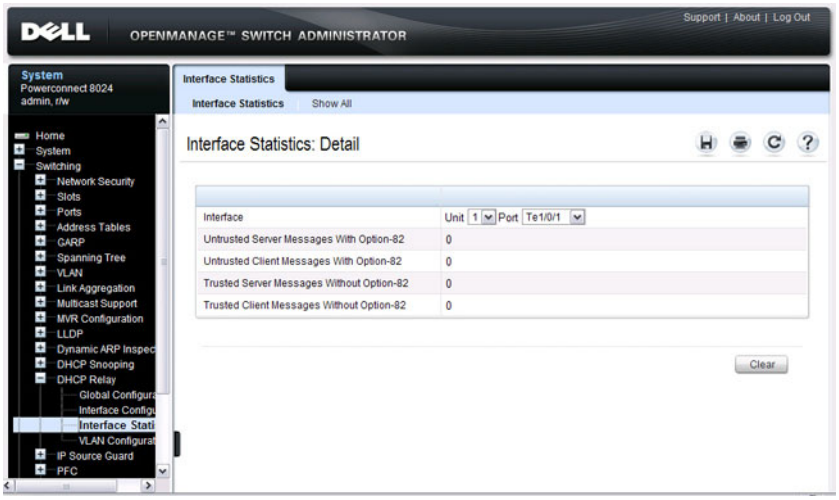
Figure 34-3. DHCP Relay Interface Summary



DHCP Relay Interface Statistics

Use this page to display statistics on DHCP Relay requests received on a selected port. To access this page, click **Switching → DHCP Relay → Interface Statistics** in the navigation panel.

Figure 34-4. DHCP Relay Interface Statistics

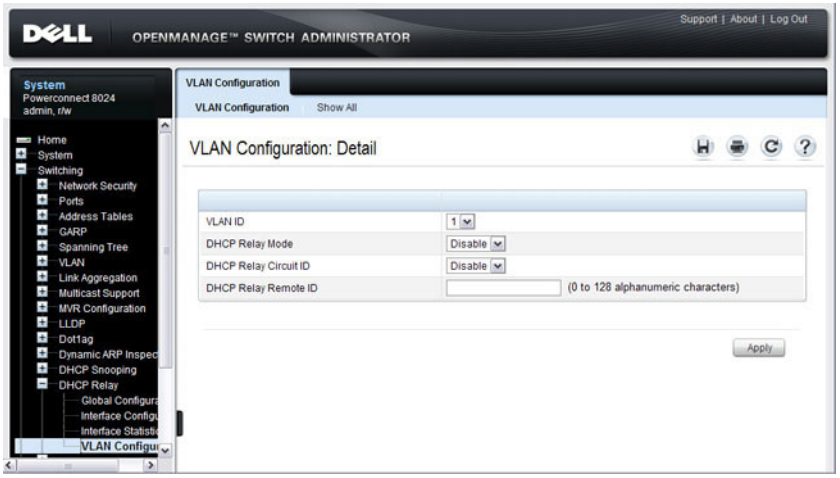


DHCP Relay VLAN Configuration

Use this page to enable and configure DHCP Relay on specific VLANs.

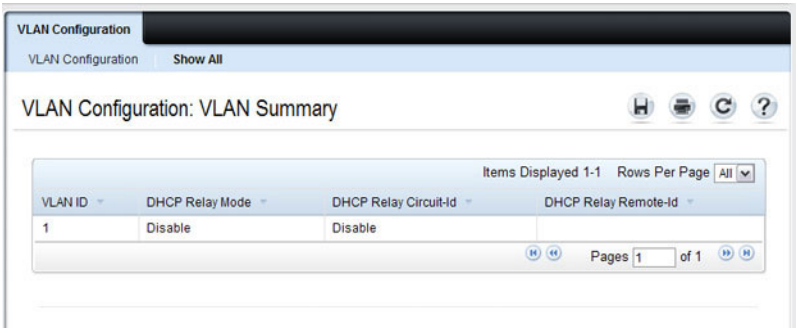
To access this page, click **Switching → DHCP Relay → VLAN Configuration** in the navigation panel.

Figure 34-5. DHCP Relay VLAN Configuration



To view a summary of the L2 DHCP relay configuration on all VLANs, click **Show All**.

Figure 34-6. DHCP Relay VLAN Summary

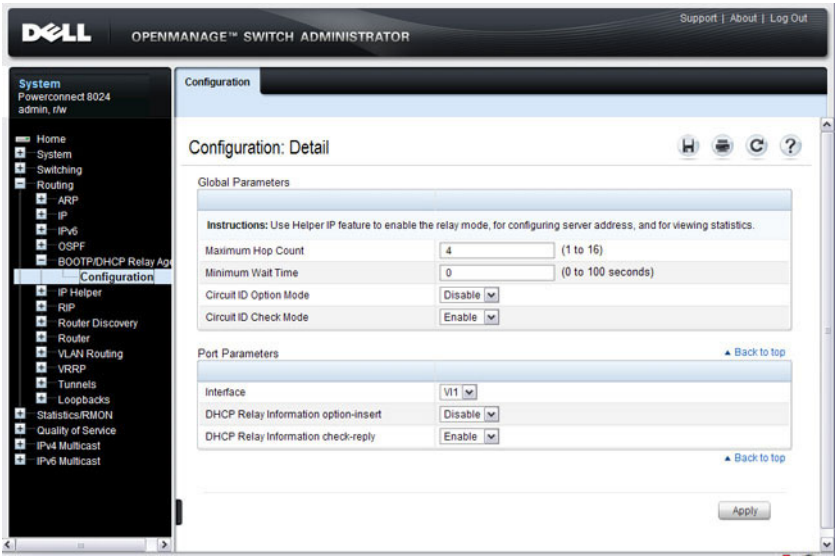


DHCP Relay Agent Configuration

Use the **Configuration** page to configure and display a DHCP relay agent.

To display the page, click **Routing → DHCP Relay Agent → Configuration** in the navigation panel.

Figure 34-7. DHCP Relay Agent Configuration

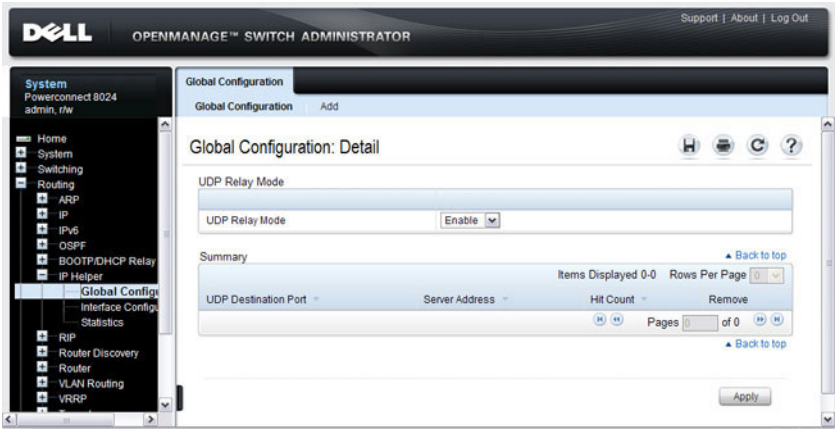


IP Helper Global Configuration

Use the **Global Configuration** page to add, show, or delete UDP Relay and Helper IP configuration

To display the page, click **Routing → IP Helper → Global Configuration** in the navigation panel.

Figure 34-8. IP Helper Global Configuration

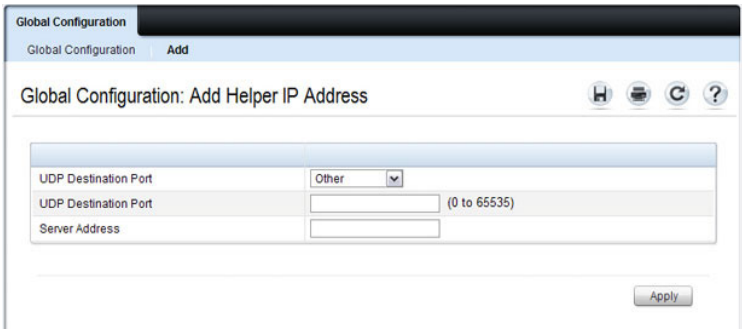


Adding an IP Helper Entry

To configure an IP helper entry:

1. Open the IP Helper **Global Configuration** page.
2. Click **Add** to display the **Add Helper IP Address** page:

Figure 34-9. Add Helper IP Address



3. Select a UDP Destination port name from the menu or enter the UDP Destination Port ID. Select the Default Set to configure for the relay entry for the default set of protocols.

   NOTE: If the DefaultSet option is specified, the device by default forwards UDP Broadcast packets for the following services: IEN-116 Name Service (port 42), DNS (port 53), NetBIOS Name Server (port 137), NetBIOS Datagram Server (port 138), TACACS Server (Port 49), and Time Service (port 37).

4. Enter the IP address of the server to which the packets with the given UDP Destination Port will be relayed.
5. Click **Apply**.

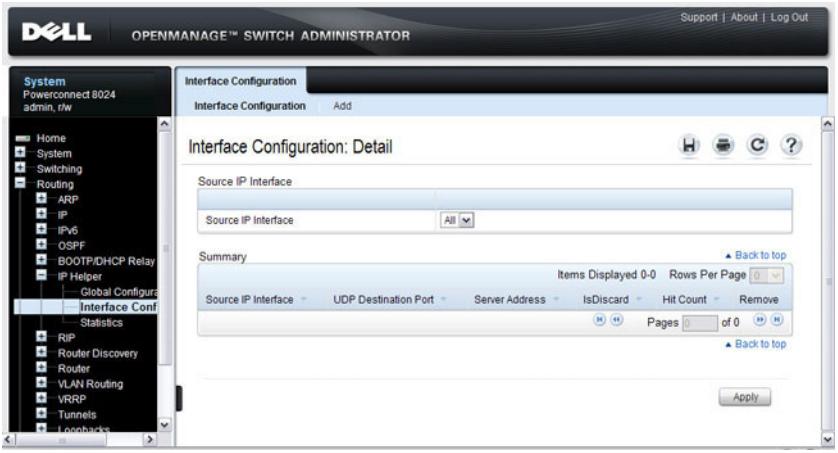
   The UDP/Helper Relay is added and the device is updated.

IP Helper Interface Configuration

Use the **Interface Configuration** page to add, show, or delete UDP Relay and Helper IP configuration for a specific interface.

To display the page, click **Routing → IP Helper → Interface Configuration** in the navigation panel.

Figure 34-10. IP Helper Interface Configuration

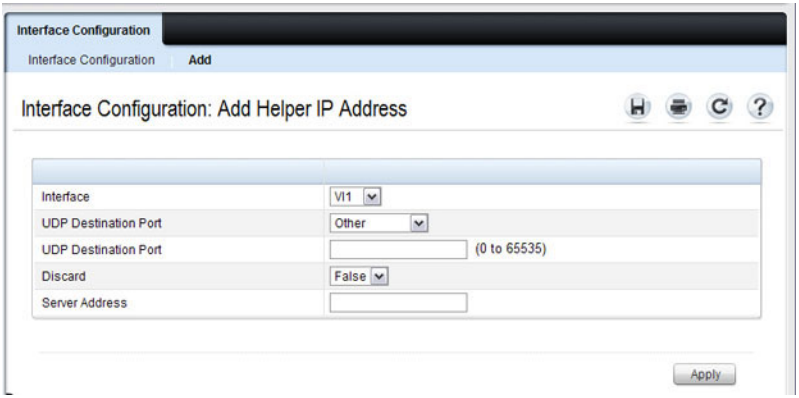


Adding an IP Helper Entry to an Interface

To add an IP helper entry to an interface:

1. Open the IP Helper **Interface Configuration** page.
2. Click **Add** to display the **Add IP Helper Address** page:

Figure 34-11. Add Helper IP Address



3. Select the interface to use for the relay.
4. Select a UDP Destination port name from the menu or enter the UDP Destination Port ID. Select the Default Set to configure for the relay entry for the default set of protocols.

   NOTE: If the DefaultSet option is specified, the device by default forwards UDP Broadcast packets for the following services: IEN-116 Name Service (port 42), DNS (port 53), NetBIOS Name Server (port 137), NetBIOS Datagram Server (port 138), TACACS Server (Port 49), and Time Service (port 37).

5. Choose whether to discard (True) or keep (False) packets arriving on the given interface with the given destination UDP port.
6. Enter the IP address of the server to which the packets with the given UDP Destination Port will be relayed.
7. Click **Apply**.

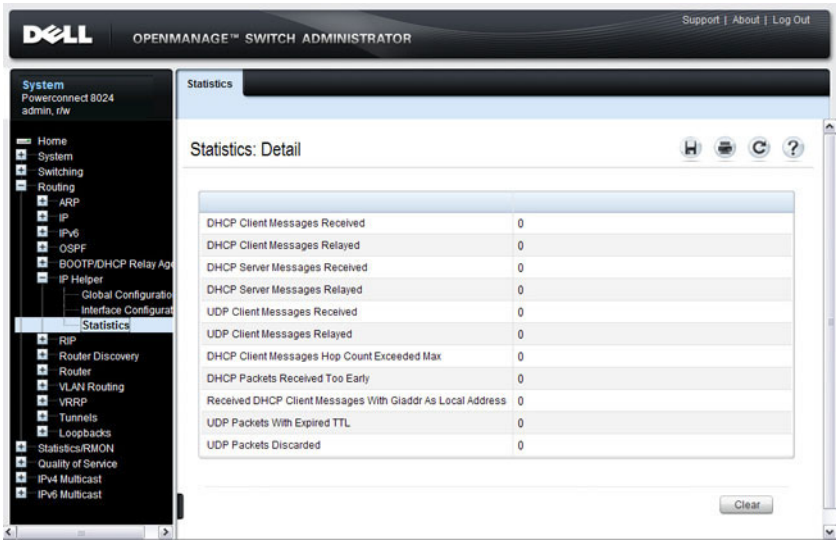
   The UDP/Helper Relay is added to the interface and the device is updated.

IP Helper Statistics

Use the **Statistics** page to view UDP Relay Statistics for the switch.

To display the page, click **Routing** → **IP Helper** → **Statistics** in the navigation panel.

Figure 34-12. IP Helper Statistics



Configuring L2 and L3 Relay Features (CLI)

This section provides information about the commands you use to configure L2 and L3 relay features on the switch. For more information about the commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Configuring L2 DHCP Relay

Beginning in Privileged EXEC mode, use the following commands to configure switch and interface L2 DHCP relay settings.

Command	Purpose
configure	Enter global configuration mode.
dhcp l2relay	Globally enable L2 DHCP relay on the switch
interface *interface*	Enter interface configuration mode for the specified port or LAG. The *interface* variable includes the interface type and number, for example **tengigabitethernet 1/0/3**. For a LAG, the interface type is **port-channel**.
	You can also specify a range of ports with the **interface range** command, for example, **interface range tengigabitethernet 1/0/8-12** configures interfaces 8, 9, 10, 11, and 12.
dhcp l2relay	Enable L2 DHCP relay on the port(s) or LAG(s).
dhcp l2relay trust	Configure the interface(s) to mandate Option-82 on receiving DHCP packets.
exit	Exit to Global Configuration mode.
dhcp l2relay vlan *vlan-range*	Enable the L2 DHCP Relay agent for a set of VLANs. All DHCP packets which arrive on interfaces in the configured VLAN are subject to L2 Relay processing.
dhcp l2relay circuit-id vlan *vlan-range*	Enable setting the DHCP Option 82 Circuit ID for a VLAN. When enabled, the interface number is added as the Circuit ID in DHCP option 82.

Command	Purpose
dhcp l2relay remote-id *remoteId* **vlan** *vlan-range*	Enable setting the DHCP Option 82 Remote ID for a VLAN. When enabled, the supplied string is used for the Remote ID in DHCP Option 82. The *remoteId* variable is a string to be used as the remote ID in the Option 82 (Range: 1 - 128 characters).
exit	Exit to Privileged EXEC mode.
show dhcp l2relay all	View L2 DHCP relay settings on the switch.
show dhcp l2relay interface [**all** | *interface*]	View L2 DHCP relay settings for all interfaces or for the specified interface.
show dhcp l2relay vlan *vlan-range*	View L2 DHCP relay settings for the specified VLAN
show dhcp l2relay stats interface [**all** | *interface*]	View the number of DHCP packets processed and relayed by the L2 relay agent. To reset the statistics to 0, use the **clear dhcp l2relay statistics interface** [**all** | *interface*] **command.**
show dhcp l2relay agent-option vlan *vlan-id*	View the DHCP L2 Relay Option-82 configuration for the specified VLAN.
show dhcp l2relay circuit-id vlan *vlan-id*	View the DHCP L2 Relay circuit ID configuration for the specified VLAN.
show dhcp l2relay remote-id vlan *vlan-id*	View the DHCP L2 Relay remote ID configuration for the specified VLAN.

Configuring L3 Relay (IP Helper) Settings

Beginning in Privileged EXEC mode, use the following commands to configure switch and interface L3 DHCP relay and IP helper settings.

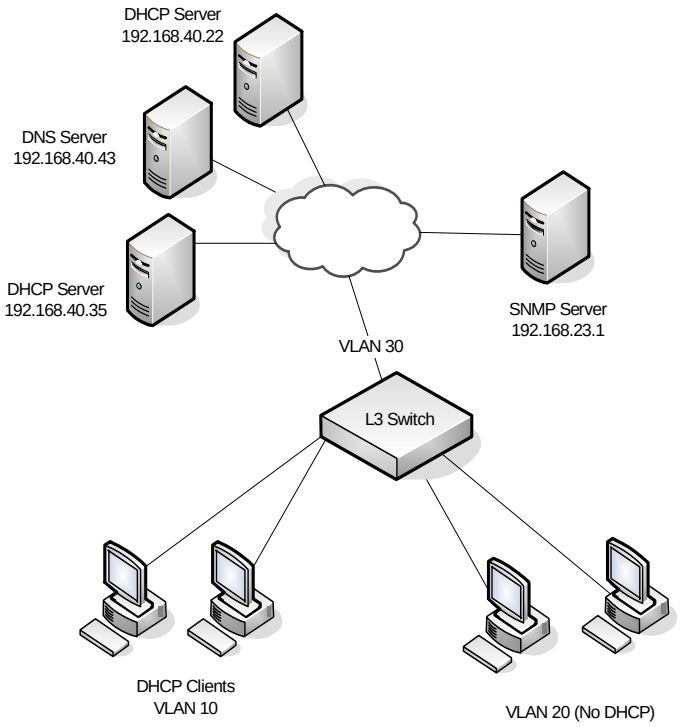
Command	Purpose
configure	Enter global configuration mode.
ip helper enable	Use this command to enable the IP helper feature. It is enabled by default.
ip helper-address *server-address* [*dest-udp-port* | **dhcp** | **domain** | **isakmp** | **mobile-ip** | **nameserver** | **netbios-dgm** | **netbios-ns** | **ntp** | **pim-auto-rp** | **rip** | **tacacs** | **tftp** | **time**]	Configure the relay of certain UDP broadcast packets received on any interface. Specify the one of the protocols defined in the command or the UDP port number. • *server-address* — The IPv4 unicast or directed broadcast address to which relayed UDP broadcast packets are sent. The server address cannot be an IP address configured on any interface of the local router. • *dest-udp-port* — A destination UDP port number from 0 to 65535.
interface vlan *vlan-id*	Enter interface configuration mode for the specified VLAN routing interface. You can also specify a range of VLAN routing interfaces with the **interface range vlan** command, for example, **interface range vlan 10,20,30** configures VLAN interfaces 10, 20, and 30. **NOTE:** All VLANs must be configured as VLAN routing interfaces.

Command	Purpose
ip helper-address {*server-address* | **discard**} [*dest-udp-port* | **dhcp** | **domain** | **isakmp** | **mobile-ip** | **nameserver** | **netbios-dgm** | **netbios-ns** | **ntp** | **pim-auto-rp** | **rip** | **tacacs** | **tftp** | **time**]	Configure the relay of certain UDP broadcast packets received on the VLAN routing interface(s). This command takes precedence over an ip helper-address command given in global configuration mode. Specify the one of the protocols defined in the command or the UDP port number. • *server-address* — The IPv4 unicast or directed broadcast address to which relayed UDP broadcast packets are sent. The server address cannot be an IP address configured on any interface of the local router. • *dest-udp-port* — A destination UDP port number from 0 to 65535.
exit	Exit to Global Config mode.
exit	Exit to Privileged EXEC mode.
show ip helper-address [**vlan** *vlan-id*]	View IP helper (L3 relay) settings for all interfaces or for the specified VLAN routing interface.
show ip helper statistics	View the number of DHCP and other UDP packets processed and relayed by the UDP relay agent. To reset the statistics to 0, use the **clear ip helper statistics** command.

Relay Agent Configuration Example

The example in this section shows how to configure the L3 relay agent (IP helper) to relay and discard various protocols.

Figure 34-13. L3 Relay Network Diagram



This example assumes that multiple VLAN routing interfaces have been created, and configured with IP addresses.

To configure the switch:

1 Relay DHCP packets received on VLAN 10 to 192.168.40.35

```
console#config
console(config)#interface vlan 10
console(config-if-vlan10)#ip helper-address
192.168.40.35 dhcp
```

2 Relay DNS packets received on VLAN 10 to 192.168.40.43

```
console(config-if-vlan10)#ip helper-address
192.168.40.35 domain
console(config-if-vlan10)#exit
```

3 Relay SNMP traps (port 162) received on VLAN 20 to 192.168.23.1

```
console(config)#interface vlan 20
console(config-if-vlan20)#ip helper-address
192.168.23.1 162
```

4 The clients on VLAN 20 have statically-configured network information, so the switch is configured to drop DHCP packets received on VLAN 20

```
console(config-if-vlan20)#ip helper-address
discard dhcp
console(config-if-vlan20)#exit
```

5 DHCP packets received from clients in any VLAN other than VLAN 10 and VLAN 20 are relayed to 192.168.40.22.



NOTE: The following command is issued in Global Configuration mode, so it applies to all interfaces except VLAN 10 and VLAN 20. IP helper commands issued in Interface Configuration mode override the commands issued in Global Configuration Mode.

```
console(config)#ip helper-address 192.168.40.22
dhcp
```

6 Verify the configuration.

```
console#show ip helper-address

IP helper is enabled

I/F      UDP Port     Discard  Hit Count    Server Address
----     --------     -----    ---------    ---------------
Vl10     domain       No       0            192.168.40.43
Vl10     dhcp         No       0            192.168.40.35
Vl20     dhcp         Yes      0
Vl20     162          No       0            192.168.23.1
Any      dhcp         No       0            192.168.40.22
```

35

Configuring OSPF and OSPFv3

This chapter describes how to configure Open Shortest Path First (OSPF) and OSPFv3. OSPF is a dynamic routing protocol for IPv4 networks, and OSPFv3 is used to route traffic in IPv6 networks. The protocols are configured separately within the software, but their functionality is largely similar for IPv4 and IPv6 networks.

NOTE: In this chapter references to OSPF apply to OSPFv2 and OSPFv3 unless otherwise noted.

The topics covered in this chapter include:

- OSPF Overview
- OSPF Feature Details
- Default OSPF Values
- Configuring OSPF Features (Web)
- Configuring OSPFv3 Features (Web)
- Configuring OSPF Features (CLI)
- Configuring OSPFv3 Features (CLI)
- OSPF Configuration Examples

OSPF Overview

OSPF is an Interior Gateway Protocol (IGP) that performs dynamic routing within a network. PowerConnect 8000-series and 8100-series switches support two dynamic routing protocols: OSPF and Routing Information Protocol (RIP).

Unlike RIP, OSPF is a link-state protocol. Larger networks typically use the OSPF protocol instead of RIP.

What Are OSPF Areas and Other OSPF Topology Features?

The top level of the hierarchy of an OSPF network is known as an OSPF domain. The domain can be divided into areas. Routers within an area must share detailed information on the topology of their area, but require less detailed information about the topology of other areas. Segregating a network into areas enables limiting the amount of route information communicated throughout the network.

Areas are identified by a numeric ID in IP address format n.n.n.n (note, however, that these are not used as actual IP addresses). For simplicity, the area can be configured and referred to in normal integer notation. For example, Area 20 is identified as 0.0.0.20 and Area 256 as 0.0.1.0. The area identified as 0.0.0.0 is referred to as Area 0 and is considered the OSPF backbone. All other OSPF areas in the network must connect to Area 0 directly or through a virtual link. The backbone area is responsible for distributing routing information between non-backbone areas.

A virtual link can be used to connect an area to Area 0 when a direct link is not possible. A virtual link traverses an area between the remote area and Area 0.

A stub area is an area that does not accept external LSAs (LSAs generated by redistributing routes) that were learned from a protocol other than OSPF or were statically configured. These routes typically send traffic outside the AS. Therefore, routes from a stub area to locations outside the AS use the default gateway. A virtual link cannot be configured across a stub area. A Not So Stubby Area can import limited external routes only from a connected ASBR.

What Are OSPF Routers and LSAs?

When a PowerConnect switch is configured to use OSPF for dynamic routing, it is considered to be an OSPF router. OSPF routers keep track of the state of the various links they send data to. Routers exchange OSPF link state advertisements (LSAs) with other routers. External LSAs provide information on static routes or routes learned from other routing protocols.

OSPF defines various router types:

- Backbone routers have an interface in Area 0.
- Area border routers (ABRs) have interfaces in multiple areas.
- Internal routers have all their interfaces in a single OSPF area.
- Autonomous system boundary routers (ASBRs) redistribute routes from other protocols and originate external LSAs.

How Are Routes Selected?

OSPF determines the best route using the route metric and the type of the OSPF route. The following order is used for choosing a route if more than one type of route exists:

1. Intra-area (the destination prefix is in the same area as the router computing the route)
2. Inter-area (the destination is not in the same area as the router computing the route
3. External Type 1
4. External Type 2

How Are OSPF and OSPFv3 Different?

OSPFv3 is the Open Shortest Path First routing protocol for IPv6. It is similar to OSPFv2 in its concept of a link state database, intra/inter area, and AS external routes and virtual links. It differs from its IPv4 counterpart in a number of respects. Peering is done through link-local addresses, and the protocol is link rather than network centric; and addressing semantics have been moved to leaf LSAs.

OSPF Feature Details

This section provides details on the following OSPF features:

- Max Metric
- Static Area Range Cost
- LSA Pacing
- LSA Pacing

Max Metric

RFC 3137 introduced stub router behavior to OSPFv2. As a stub, a router can inform other routers that it is not available to forward data packets. This can be useful if OSPF has run out of resources (for example, memory) to compute a complete routing table, or to avoid routing transients as OSPF learns its neighbors and a complete set of routes at startup. Thus, OSPF can enter stub router mode either automatically (as a result of a resource condition) or by configuration.

When OSPF enters stub router mode, it re-originates its router LSAs and sets the metric on each of its non-stub links to the maximum value, 0xFFFF. Whenever OSPF originates a router LSA while in stub router mode, it sets the metrics in this way. Stub router mode is global and applies to router LSAs for all areas. Other routers prefer alternate paths that avoid the stub router; however, if no alternate path is available, another router may compute a transit route through a stub router. Because the stub router does not adjust the metric for stub links in its router LSA, routes to destinations on these networks are unaffected. Thus, stub router mode does not affect management connections to the router, even if the router and management station depend on OSPF routes to communicate with each other.

The feature supports two modes of operation. The network administrator can put OSPF in stub router mode. OSPF remains in stub router mode until the network administrator takes OSPF out of stub router mode. Alternatively, the network administrator can configure OSPF to start in stub router mode for a configurable period of time after the router boots up. On a stack, the startup period also applies when a unit takes over as the management unit. The **`clear configuration`** command also restarts OSPF in stub router

mode. OSPF does not begin in stub router mode when OSPF is globally enabled. If the operator wants to avoid routing transients when he enables or configures OSPF, he can manually set OSPF in stub router mode.

If OSPF is in startup stub router mode and encounters a resource limitation that would normally cause OSPF to become a stub router, OSPF cancels the timer to exit startup stub router and remains in stub router mode until the network administrator takes action.

The network administrator can optionally configure OSPF to override the metric in summary LSAs while in stub router mode. The option applies to both type 3 and type 4 summary LSAs.

When a router is in stub router mode, all its virtual links are down. This is because the cost to the virtual neighbor is guaranteed to be greater than or equal to 0xFFFF. RFC 2328 section 15 states that:

> "...a virtual link whose underlying path has cost greater than hexadecimal 0xffff (the maximum size of an interface cost in a router-LSA) should be considered non-operational."

To configure a router for stub router mode, use the **max-metric router-lsa** command in Global Router Configuration mode. The following example sets the router to start in stub router mode on a restart and remain in stub router mode for 5 minutes:

```
ABR-R0(config)#router ospf
ABR-R0(config-router)#max-metric router-lsa on-startup 300
```

The following example sets the router to advertise the metric in type 3 and type 4 summary LSAs as 32768 for 5 minutes after a restart, after which time the router will exit stub router mode and advertise the full set of LSAs:

```
ABR-R0(config)#router ospf
ABR-R0(config-router)#max-metric router-lsa on-startup 300 summary-
lsa 32768
```

The following example causes the router to exit stub router mode, whether entered automatically due to resource constraints or due to configuration by the operator. Virtual links are enabled when the router exits stub router mode.

```
ABR-R0(config)#router ospf
ABR-R0(config-router)#no max-metric router-lsa
```

Static Area Range Cost

This feature allows a network operator to configure a fixed OSPF cost that is always advertised when an area range is active. This feature applies to both OSPFv2 and OSPFv3.

An OSPF domain can be divided into areas to limit the processing required on each router. Area Border Routers (ABRs) advertise reachability across area boundaries. It is common to summarize the set of prefixes that an ABR advertises across an area boundary. RFC 2328 specifies that when an ABR originates a type 3 LSA for an active area range, the cost in the LSA is set to "the largest cost of any of the component networks." Thus, when an area's topology changes in a way that increases the largest cost, the type 3 LSA must be re-originated. In some cases, advertising the change in cost may be less important than preventing the topology change from propagating outside the area (thus causing routers in other areas to process and flood a changed LSA and rerun their routing table calculations). For this reason, it is common to give the network administrator the option of configuring the cost for an area range. When a static cost is configured, the cost advertised in the type 3 LSA does not depend on the cost of the component networks. Thus, topology changes within an area do not propagate outside the area, resulting in greater stability within the OSPF domain.

PowerConnect switches also use area ranges to summarize type 7 LSAs when they are translated to type 5 LSAs. The cost option may be configured on area ranges used for type 7 to type 5 translation.

If an area range is configured for type 3 summarization and the static cost is set to the maximum value, 16,777,215, the range is not advertised. Setting this static cost is equivalent to configuring a range with the not-advertise option. A summary LSA with this metric (LSInfinity) cannot be advertised, according to RFC 2328 section 12.4.3. This behavior is consistent with the industry standard.

If an area range is configured for type 7 to type 5 translation, a type 5 LSA is sent if the metric is set to 16,777,215; however, other routers will not compute a route from a type 5 LSA with this metric.

See "Configuring the Static Area Range Cost" on page 1008 for a configuration example.

LSA Pacing

OSPF refreshes each self-originated LSA every 30 minutes. Because a router tends to originate many LSAs at the same time, either at startup or when adjacencies are formed or when routes are first learned, LSA refreshes tend to be grouped. Further, Area Border Routers (ABRs) attached to the same area tend to originate summary LSAs into the area at the same time. This behavior leads to periodic bursts of LS Update packets. Update bursts can lead to high CPU utilization, packet loss, and retransmission, if a receiver cannot absorb all packets in a burst. These losses occur primarily in two places: 1) at the Class of Service (CoS) queue where the hardware queues packets to the CPU, and 2) when a message buffer is allocated for an incoming packet.

This feature makes changes to OSPFv2 to improve the efficiency of LSA flooding, with the expectation that the improvements will greatly reduce or eliminate the packet drops caused by bursts in OSPF control packets. The changes are as follows:

- Introduce LSA transmit pacing, limiting the rate of LS Update packets that OSPF can send
- Introduce LSA refresh groups, so that OSPF efficiently bundles LSAs into LS Update packets when periodically refreshing self-originated LSAs

To configure LSA transmit pacing, use the timers pacing flood command in router config mode:

```
ABR-R0(config)#router ospf
ABR-R0(config-router)#timers pacing flood 50
```

This will cause LSA Update packets to be sent at no less than a 50 millisecond interval.

When OSPF refreshes LSAs, it considers all self-originated LSAs whose age is from 1800 to 1800 plus the pacing group size. Grouping LSAs for refresh allows OSPF to combine refreshed LSAs into a minimal number of LS Update packets. Minimizing the number of Update packets makes LSA distribution more efficient. To configure an LSA Refresh window, use the timers pacing lsa-group command in router-config mode:

```
ABR-R0(config)#router ospf
ABR-R0(config-router)#timers pacing lsa-group 300
```

This sets the LSA Refresh window to 2100 seconds or about 35 minutes.

Flood Blocking

OSPF is a link state routing protocol. Routers describe their local environment in Link State Advertisements (LSAs), which are distributed throughout an area or OSPF domain. Through this process, each router learns enough information to compute a set of routes consistent with the routes computed by all other routers.

Normally, OSPF floods an LSA on all interfaces within the LSA's flooding scope. Flooding ensures that all routers receive all LSAs. A router normally receives a duplicate copy of each LSA once on each interface in the LSA's flooding scope. The duplicate deliveries make OSPF LSA distribution robust, but in highly interconnected networks, can cause a lot of buffer and CPU usage. Buffer and CPU use can be reduced by selectively blocking LSA flooding on some interfaces, while ensuring that LSAs are flooded on enough interfaces to guarantee delivery of all LSAs to all routers. When enabling flood blocking, the network administrator must ensure there is sufficient LSA flooding even when there are router and link failures.

This feature enables a network administrator to disable LSA flooding on an interface. Flood blocking only affects flooding of LSAs with area or AS (i.e., domain-wide) scope. Such LSAs are expected to be flooded to neighbors on other, unblocked interfaces, and eventually reach neighbors on blocked interfaces. An LSA with interface flooding scope cannot be blocked; there is no other way for interface-scope LSAs to reach neighbors on the blocked interface. Allowing interface-scope LSAs on blocked interfaces allows graceful restart to work, even if the restarting router has neighbors on flood blocked interfaces.

When an interface is blocked, LSAs with area or AS scope are not sent to any neighbor on that interface. When flood blocking is enabled, OSPF does not advertise any LSAs with area or AS scope in its database description packets sent to neighbors on a blocked interface. When OSPF receives an LSA from a neighbor and the local database copy is newer than the received LSA, OSPF normally sends the newer LSA directly to the neighbor. If the neighbor is on a blocked interface, OSPF neither acknowledges the LSA nor sends the newer LSA. Instead, OSPF expects that the neighbor will receive the newer LSA indirectly.

Flooding is enabled by default.

Flood blocking cannot be enabled on virtual interfaces. While the feature could be allowed on virtual interfaces, it is less likely to be used on a virtual interface, since virtual interfaces are created specifically to allow flooding between two backbone routers. So the option of flood blocking on virtual interfaces is not supported.

See "Configuring Flood Blocking" on page 1013 for a configuration example.

Default OSPF Values

OSPF is globally enabled by default. To make it operational on the router, you must configure a router ID and enable OSPF on at least one interface.

Table 35-1 shows the global default values for OSPF and OSPFv3.

Table 35-1. OSPF/OSPFv3 Global Defaults

Parameter	Default Value
Router ID	None
Admin Mode	Enabled
RFC 1583 Compatibility	Enabled (OSPFv2 only)
ABR Status	Enabled
Opaque LSA Status	Enabled (OSPFv2 only)
Exit Overflow Interval	Not configured
SPF Delay Time	5 (OSPFv2 only)
SPF Hold Time	10 (OSPFv2 only)
External LSDB Limit	None
Default Metric	Not configured
Maximum Paths	4
AutoCost Reference Bandwidth	100 Mbps
Default Passive Setting	Disabled
Default Information Originate	Disabled

Table 35-2 shows the per-interface default values for OSPF and OSPFv3.

Table 35-2. OSPF Per-Interface Defaults

Parameter	Default Value
Admin Mode	Disabled
Advertise Secondaries	Enabled (OSPFv2 only)
Router Priority	1
Retransmit Interval	5 seconds
Hello Interval	10 seconds
Dead Interval	40 seconds
LSA Ack Interval	1 second
Interface Delay Interval	1 second
MTU Ignore	Disabled
Passive Mode	Disabled
Network Type	Broadcast
Authentication Type	None (OSPFv2 only)
Metric Cost	Not configured

Configuring OSPF Features (Web)

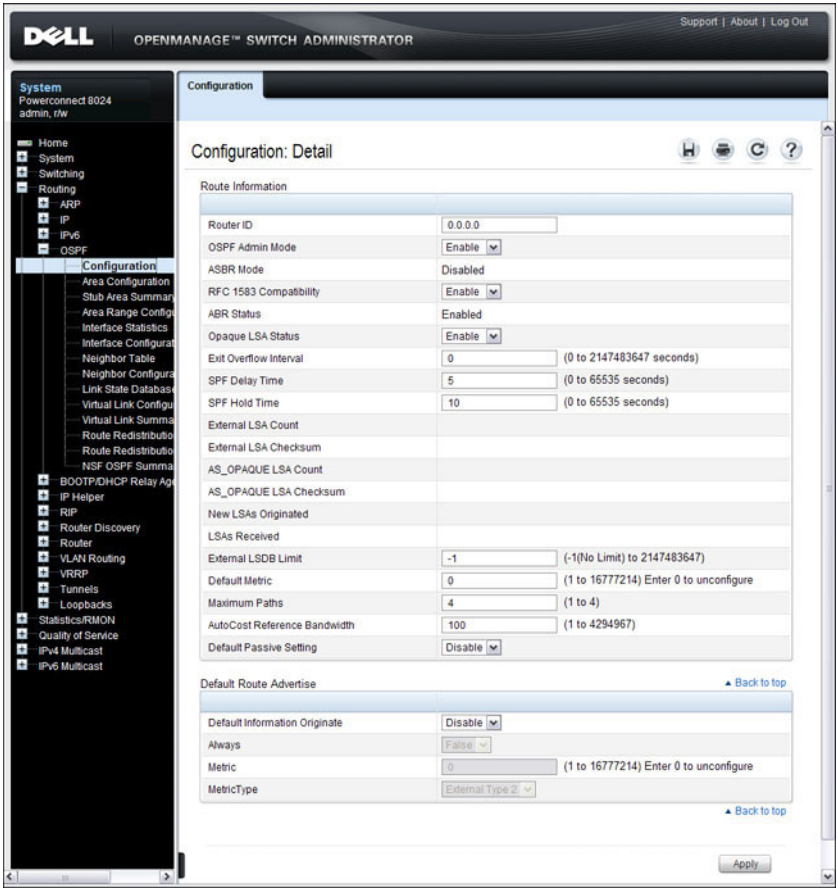
This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring OSPF features on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click  at the top of the page.

OSPF Configuration

Use the **Configuration** page to enable OSPF on a router and to configure the related OSPF settings.

To display the page, click **Routing → OSPF → Configuration** in the navigation panel.

Figure 35-1. OSPF Configuration

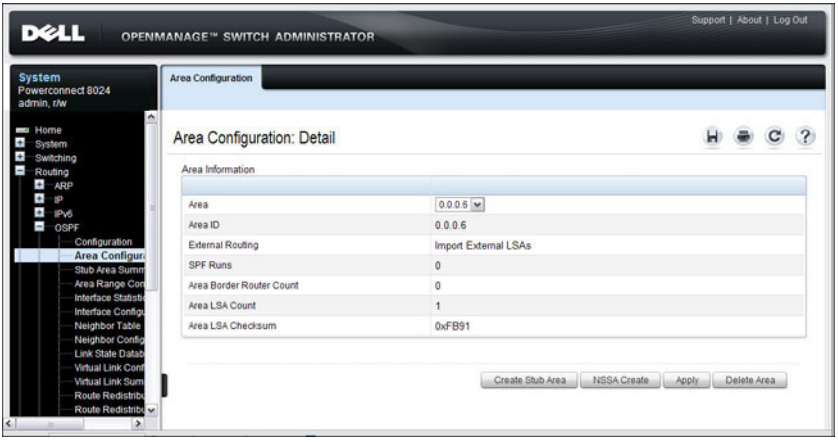


OSPF Area Configuration

The **Area Configuration** page lets you create a Stub area configuration and NSSA once you've enabled OSPF on an interface through **Routing → OSPF → Interface Configuration**. At least one router must have OSPF enabled for this web page to display.

To display the page, click **Routing → OSPF → Area Configuration** in the navigation panel. If a Stub Area has been created, the fields in the Stub Area Information are available. If a NSSA has been created, the fields in the NSSA Area Information are available.

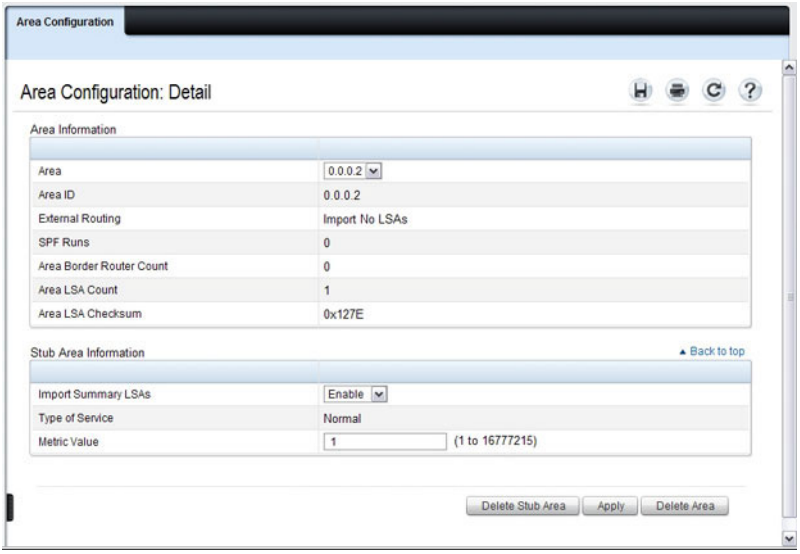
Figure 35-2. OSPF Area Configuration



Configuring an OSPF Stub Area

To configure the area as an OSPF stub area, click **Create Stub Area**. The pages refreshes, and displays additional fields that are specific to the stub area.

Figure 35-3. OSPF Stub Area Configuration

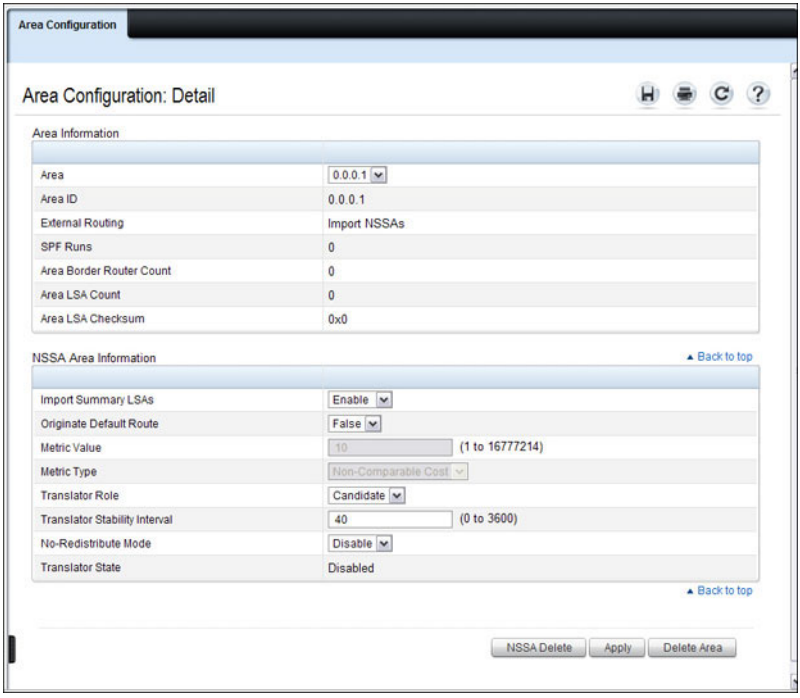


Use the **Delete Stub Area** button to remove the stub area.

Configuring an OSPF Not-So-Stubby Area

To configure the area as an OSPF not-so-stubby area (NSSA), click **NSSA Create**. The pages refreshes, and displays additional fields that are specific to the NSSA.

Figure 35-4. OSPF NSSA Configuration



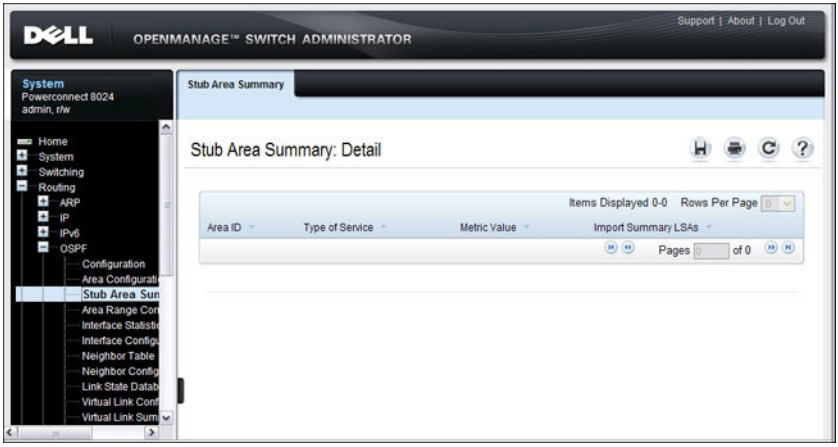
Use the **NSSA Delete** button to remove the NSSA area.

OSPF Stub Area Summary

The **Stub Area Summary** page displays OSPF stub area detail.

To display the page, click **Routing** → **OSPF** → **Stub Area Summary** in the navigation panel.

Figure 35-5. OSPF Stub Area Summary

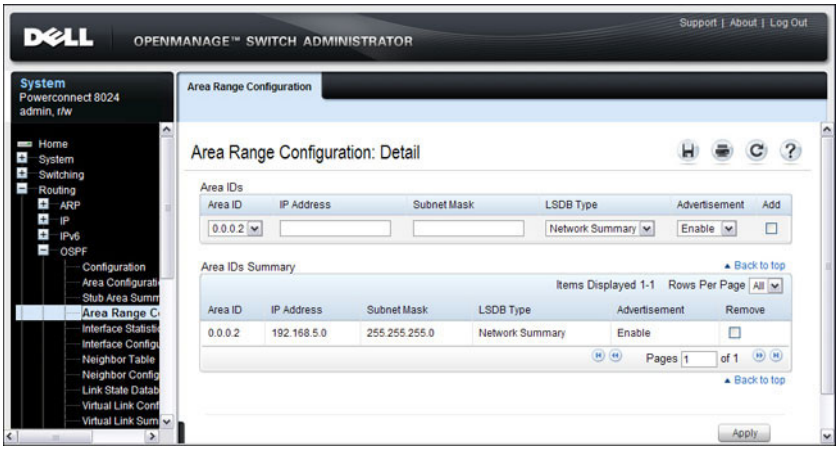


OSPF Area Range Configuration

Use the **Area Range Configuration** page to configure and display an area range for a specified NSSA.

To display the page, click **Routing** → **OSPF** → **Area Range Configuration** in the navigation panel.

Figure 35-6. OSPF Area Range Configuration

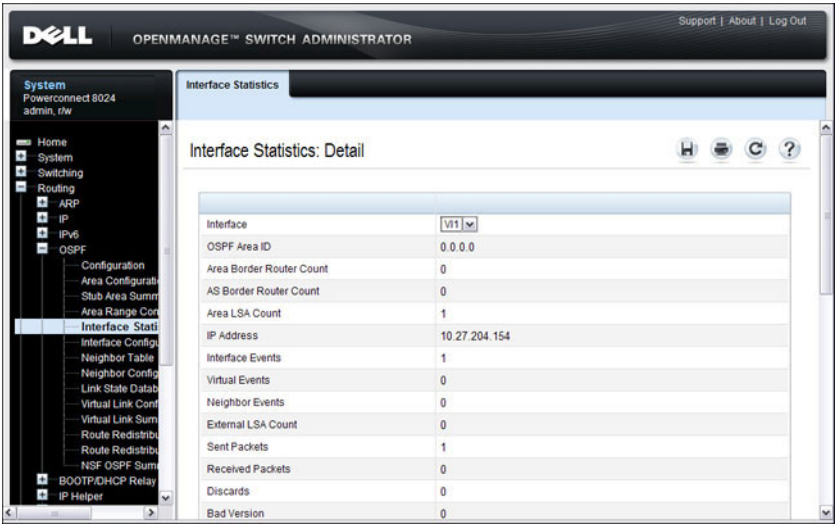


OSPF Interface Statistics

Use the **Interface Statistics** page to display statistics for the selected interface. The information is displayed only if OSPF is enabled.

To display the page, click **Routing → OSPF → Interface Statistics** in the navigation panel.

Figure 35-7. OSPF Interface Statistics

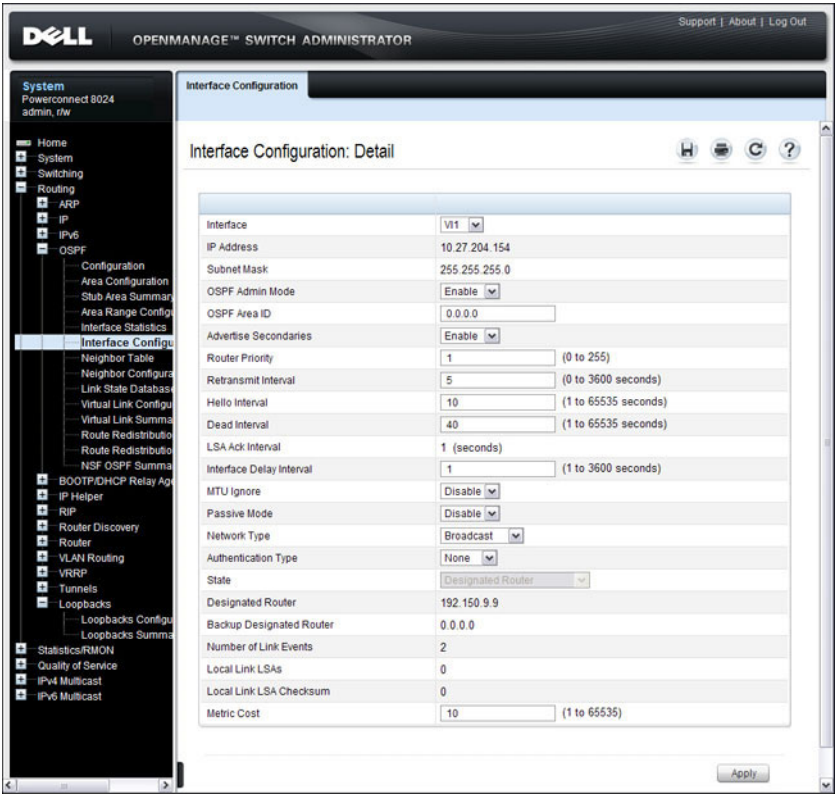


OSPF Interface Configuration

Use the **Interface Configuration** page to configure an OSPF interface.

To display the page, click **Routing → OSPF → Interface Configuration** in the navigation panel.

Figure 35-8. OSPF Interface Configuration

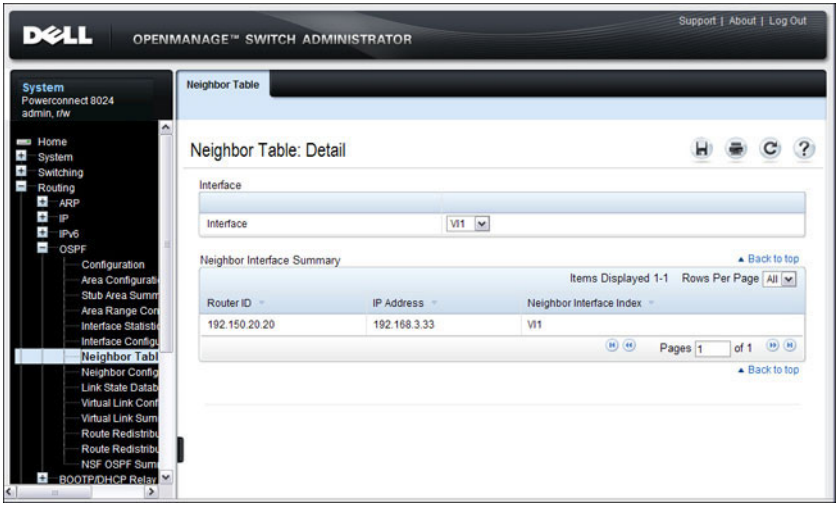


OSPF Neighbor Table

Use the **Neighbor Table** page to display the OSPF neighbor table list. When a particular neighbor ID is specified, detailed information about a neighbor is given. The information below is only displayed if OSPF is enabled.

To display the page, click **Routing** → **OSPF** → **Neighbor Table** in the navigation panel.

Figure 35-9. OSPF Neighbor Table

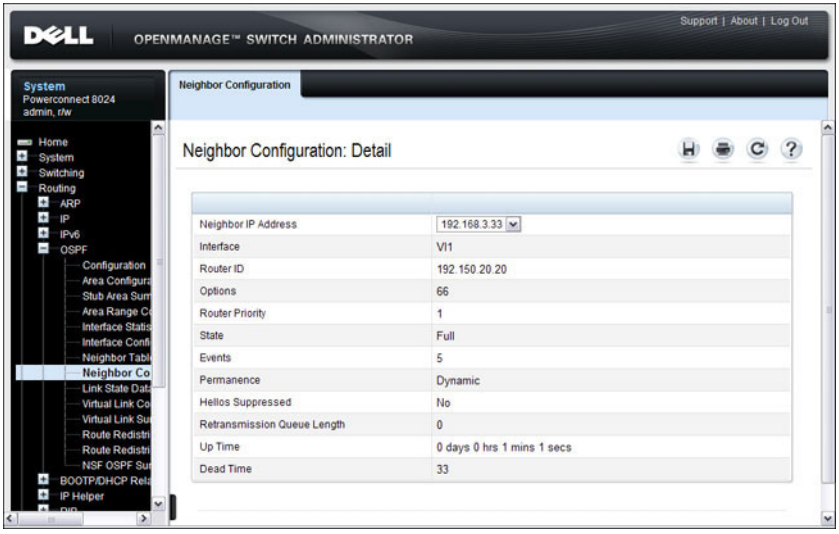


OSPF Neighbor Configuration

Use the **Neighbor Configuration** page to display the OSPF neighbor configuration for a selected neighbor ID. When a particular neighbor ID is specified, detailed information about a neighbor is given. The information below is only displayed if OSPF is enabled and the interface has a neighbor. The IP address is the IP address of the neighbor.

To display the page, click **Routing** → **OSPF** → **Neighbor Configuration** in the navigation panel.

Figure 35-10. OSPF Neighbor Configuration

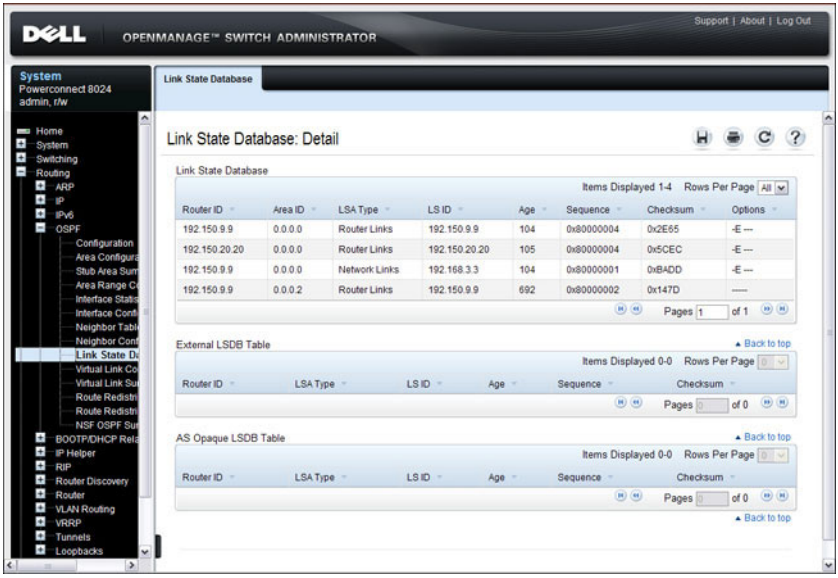


OSPF Link State Database

Use the **Link State Database** page to display OSPF link state, external LSDB table, and AS opaque LSDB table information.

To display the page, click **Routing → OSPF → Link State Database** in the navigation panel.

Figure 35-11. OSPF Link State Database

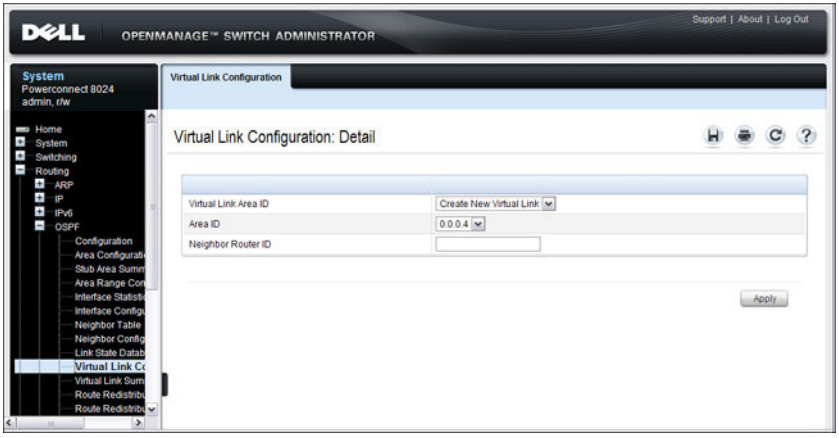


OSPF Virtual Link Configuration

Use the **Virtual Link Configuration** page to create or configure virtual interface information for a specific area and neighbor. A valid OSPF area must be configured before this page can be displayed.

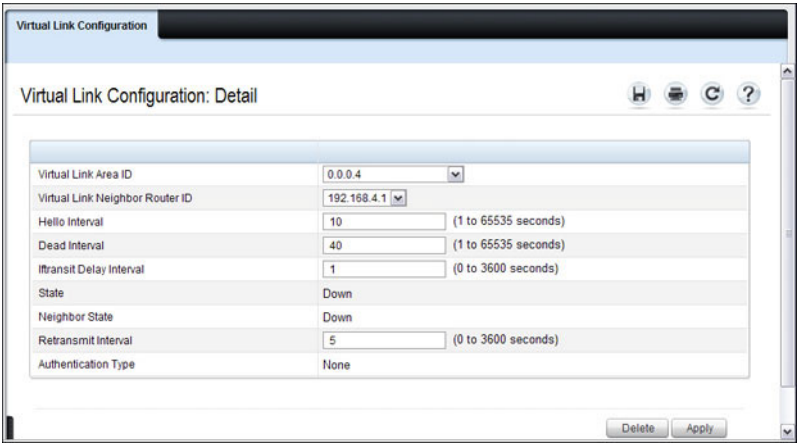
To display the page, click **Routing → OSPF → Virtual Link Configuration** in the navigation panel.

Figure 35-12. OSPF Virtual Link Creation



After you create a virtual link, additional fields display, as the Figure 35-13 shows.

Figure 35-13. OSPF Virtual Link Configuration

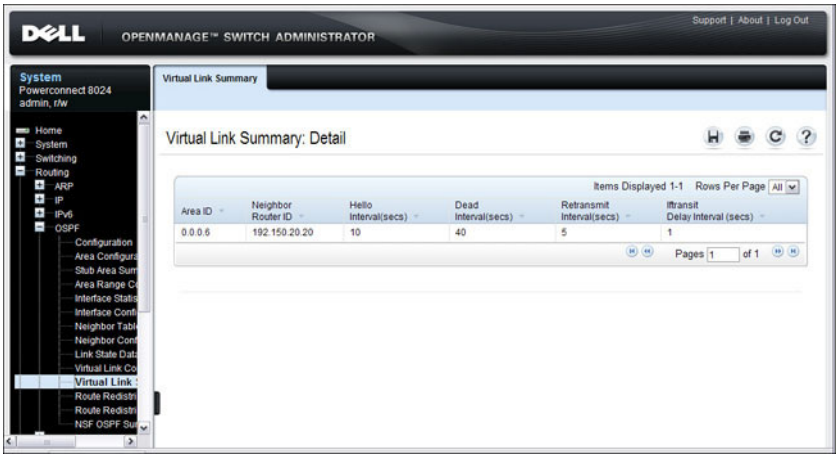


OSPF Virtual Link Summary

Use the **Virtual Link Summary** page to display all of the configured virtual links.

To display the page, click **Routing → OSPF → Virtual Link Summary** in the navigation panel.

Figure 35-14. OSPF Virtual Link Summary

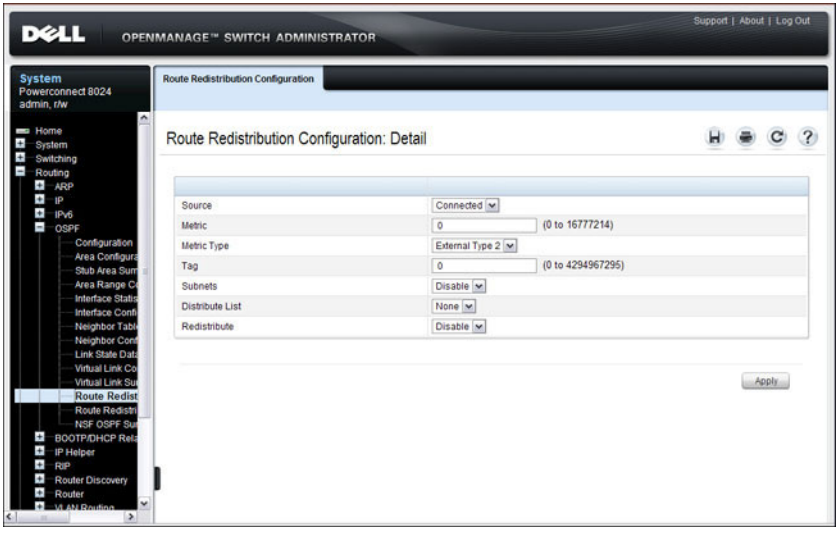


OSPF Route Redistribution Configuration

Use the **Route Redistribution Configuration** page to configure redistribution in OSPF for routes learned through various protocols. You can choose to redistribute routes learned from all available protocols or from selected ones.

To display the page, click **Routing → OSPF → Route Redistribution Configuration** in the navigation panel.

Figure 35-15. OSPF Route Redistribution Configuration

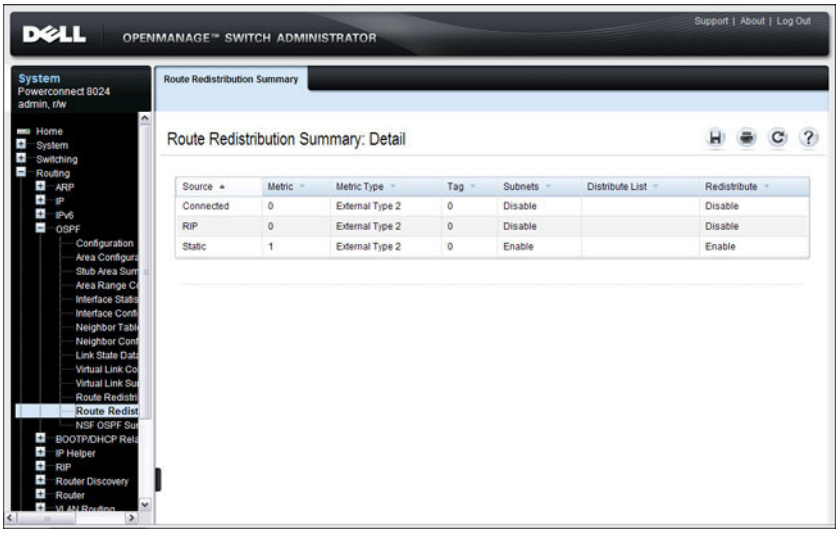


OSPF Route Redistribution Summary

Use the **Route Redistribution Summary** page to display OSPF Route Redistribution configurations.

To display the page, click **Routing → OSPF → Route Redistribution Summary** in the navigation panel.

Figure 35-16. OSPF Route Redistribution Summary



NSF OSPF Configuration

Use the **NSF OSPF Configuration** page to configure the non-stop forwarding (NSF) support mode and to view NSF summary information for the OSPF feature. NSF is a feature used in switch stacks to maintain switching and routing functions in the event of a stack unit failure. For information about NSF, see "What is Nonstop Forwarding?" on page 147 in the Managing a Switch Stack chapter.

To display the page, click **Routing → OSPF → NSF OSPF Configuration** in the navigation panel.

Figure 35-17. NSF OSPF Configuration



Configuring OSPFv3 Features (Web)

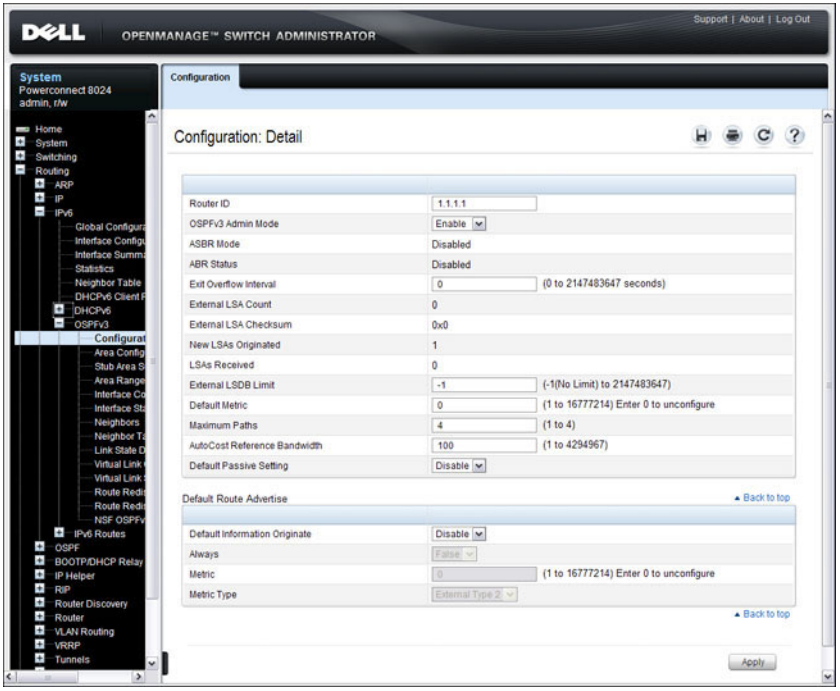
This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring OSPFv3 features on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

OSPFv3 Configuration

Use the **Configuration** page to activate and configure OSPFv3 for a switch.

To display the page, click **IPv6 → OSPFv3 → Configuration** in the navigation panel.

Figure 35-18. OSPFv3 Configuration

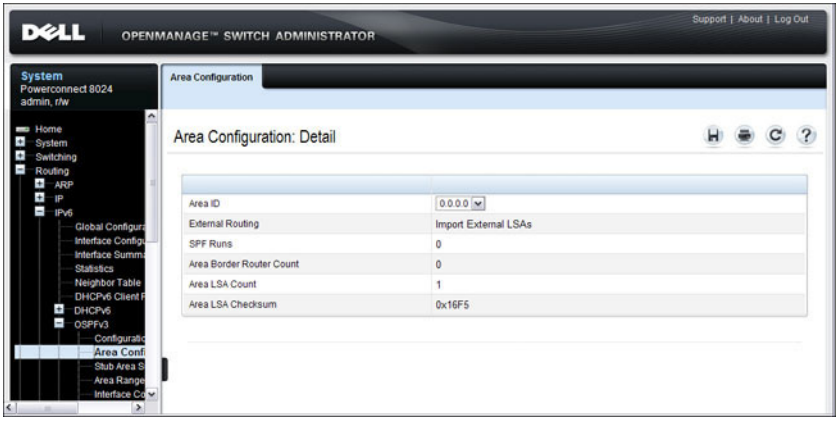


OSPFv3 Area Configuration

Use the **Area Configuration** page to create and configure an OSPFv3 area.

To display the page, click **IPv6 → OSPFv3 → Area Configuration** in the navigation panel.

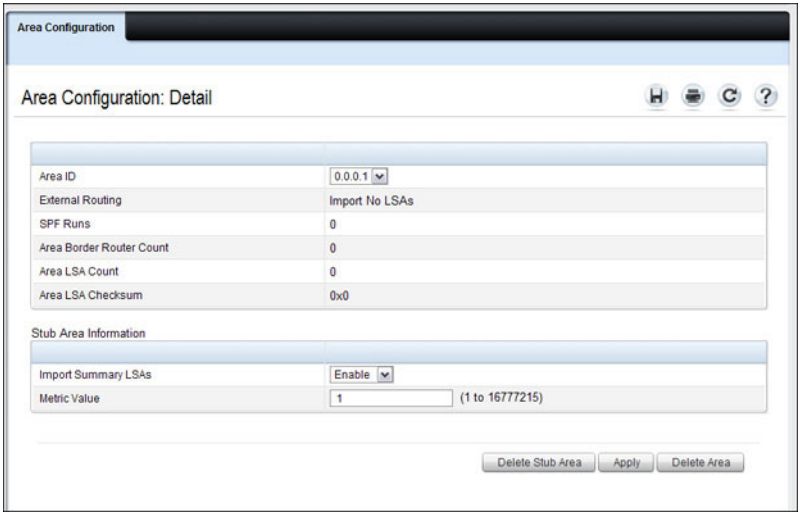
Figure 35-19. OSPFv3 Area Configuration



Configuring an OSPFv3 Stub Area

To configure the area as an OSPFv3 stub area, click **Create Stub Area**. The pages refreshes, and displays additional fields that are specific to the stub area.

Figure 35-20. OSPFv3 Stub Area Configuration

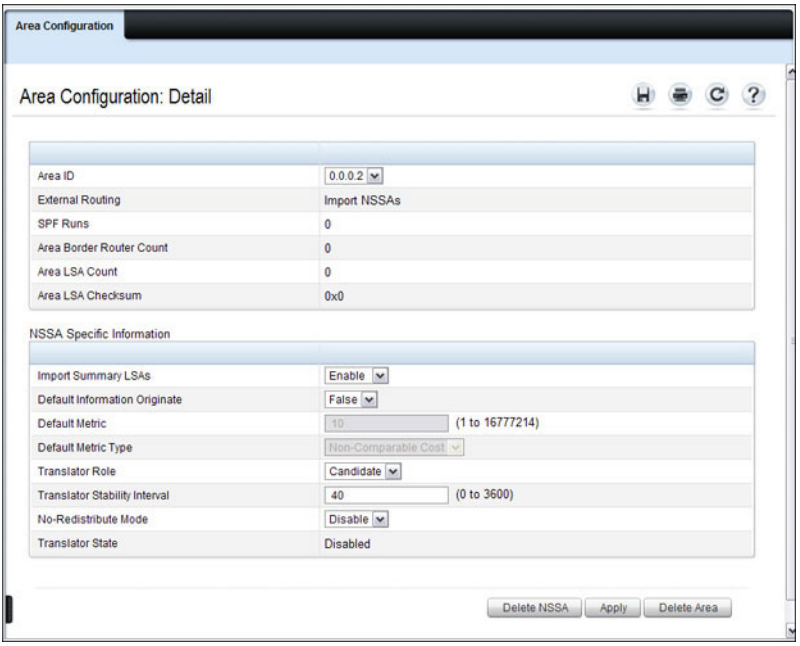


Use the **Delete Stub Area** button to remove the stub area.

Configuring an OSPFv3 Not-So-Stubby Area

To configure the area as an OSPFv3 not-so-stubby area (NSSA), click **Create NSSA**. The pages refreshes, and displays additional fields that are specific to the NSSA.

Figure 35-21. OSPFv3 NSSA Configuration



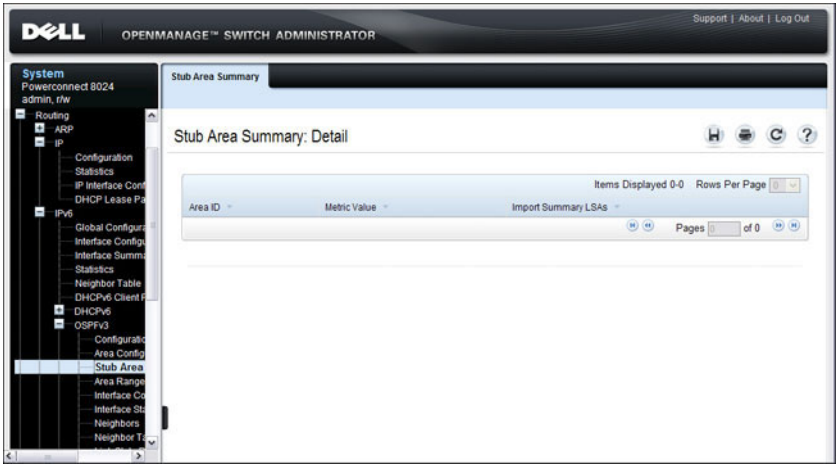
Use the **Delete NSSA** button to remove the NSSA area.

OSPFv3 Stub Area Summary

Use the **Stub Area Summary** page to display OSPFv3 stub area detail.

To display the page, click **IPv6 → OSPFv3 → Stub Area Summary** in the navigation panel.

Figure 35-22. OSPFv3 Stub Area Summary

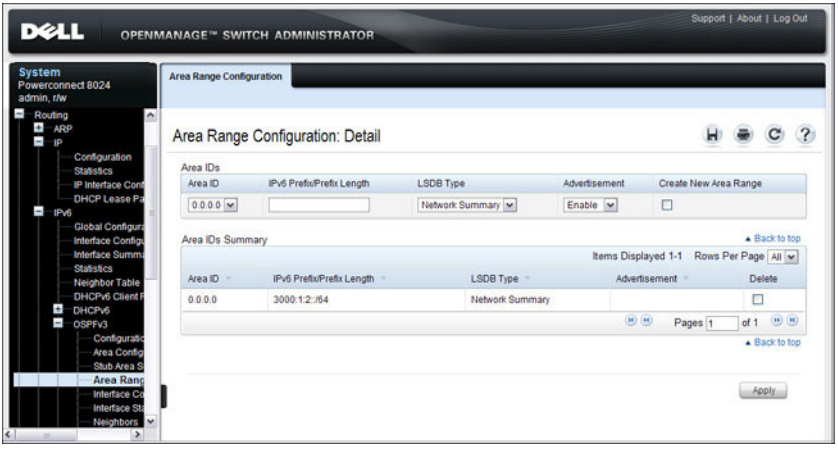


OSPFv3 Area Range Configuration

Use the **Area Range Configuration** page to configure OSPFv3 area ranges.

To display the page, click **IPv6 → OSPFv3 → Area Range Configuration** in the navigation panel.

Figure 35-23. OSPFv3 Area Range Configuration

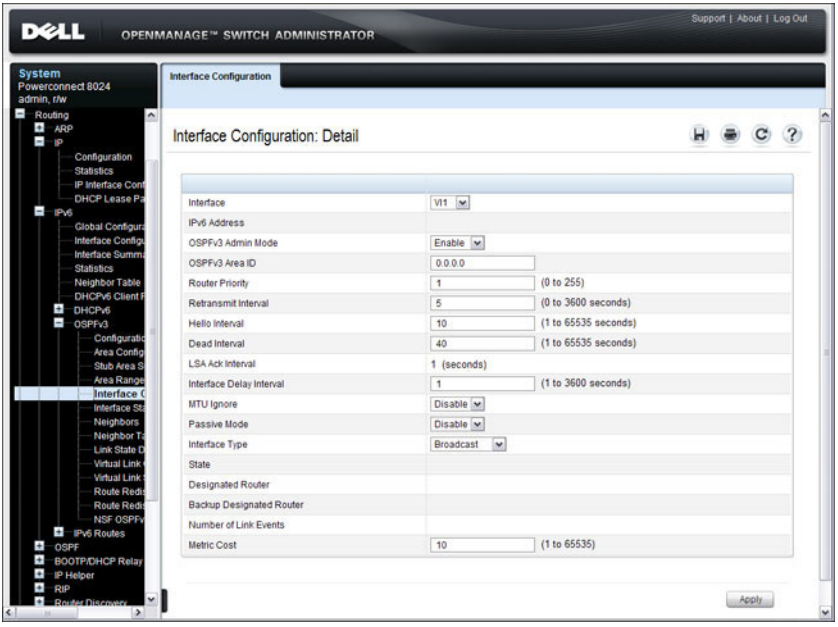


OSPFv3 Interface Configuration

Use the **Interface Configuration** page to create and configure OSPFv3 interfaces. This page has been updated to include the Passive Mode field.

To display the page, click **IPv6 → OSPFv3 → Interface Configuration** in the navigation panel.

Figure 35-24. OSPFv3 Interface Configuration



OSPFv3 Interface Statistics

Use the **Interface Statistics** page to display OSPFv3 interface statistics. Information is only displayed if OSPF is enabled. Several fields have been added to this page.

To display the page, click **IPv6 → OSPFv3 → Interface Statistics** in the navigation panel.

Figure 35-25. OSPFv3 Interface Statistics

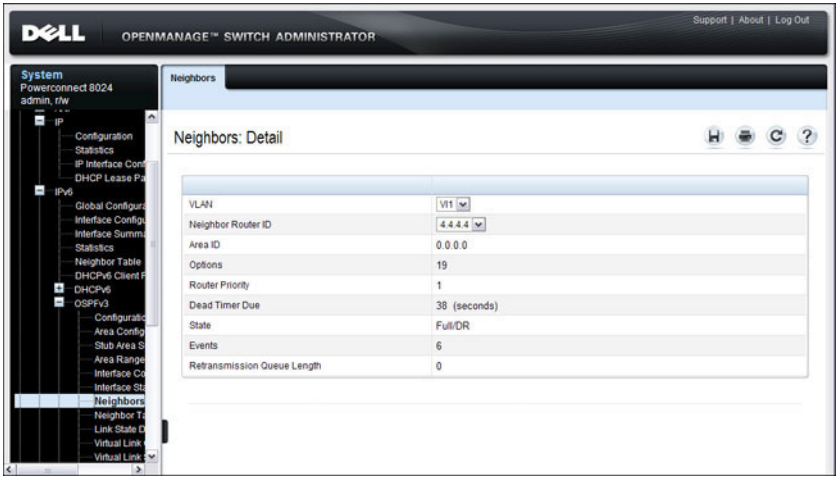


OSPFv3 Neighbors

Use the **Neighbors** page to display the OSPF neighbor configuration for a selected neighbor ID. When a particular neighbor ID is specified, detailed information about that neighbor is given. Neighbor information only displays if OSPF is enabled and the interface has a neighbor. The IP address is the IP address of the neighbor.

To display the page, click **IPv6 → OSPFv3 → Neighbors** in the navigation panel.

Figure 35-26. OSPFv3 Neighbors

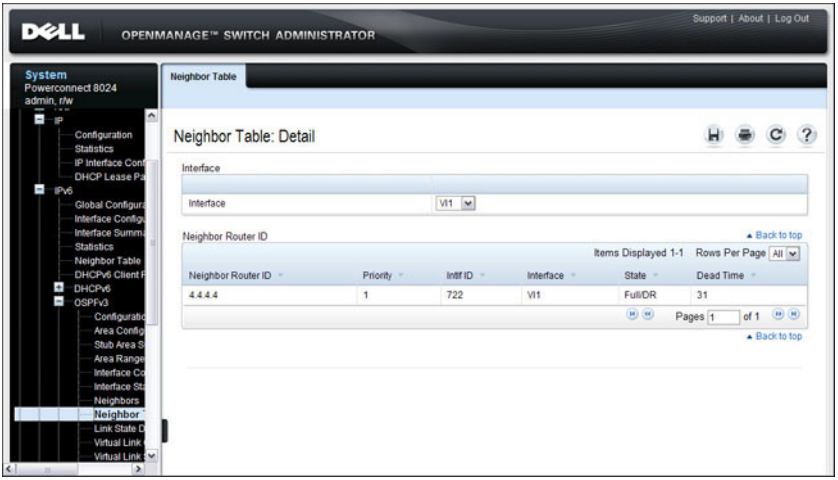


OSPFv3 Neighbor Table

Use the **Neighbor Table** page to display the OSPF neighbor table list. When a particular neighbor ID is specified, detailed information about a neighbor is given. The neighbor table is only displayed if OSPF is enabled.

To display the page, click **IPv6 → OSPFv3 → Neighbor Table** in the navigation panel.

Figure 35-27. OSPFv3 Neighbor Table

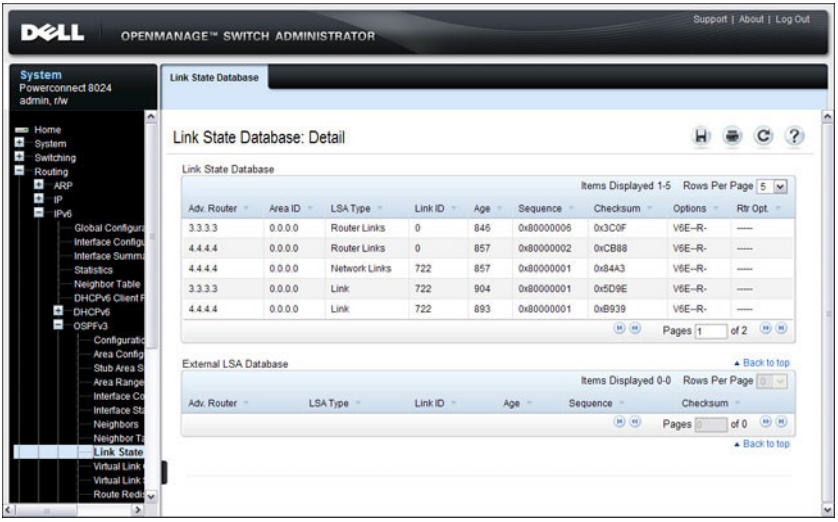


OSPFv3 Link State Database

Use the **Link State Database** page to display the link state and external LSA databases. The OSPFv3 **Link State Database** page has been updated to display external LSDB table information in addition to OSPFv3 link state information.

To display the page, click **IPv6 → OSPFv3 → Link State Database** in the navigation panel.

Figure 35-28. OSPFv3 Link State Database

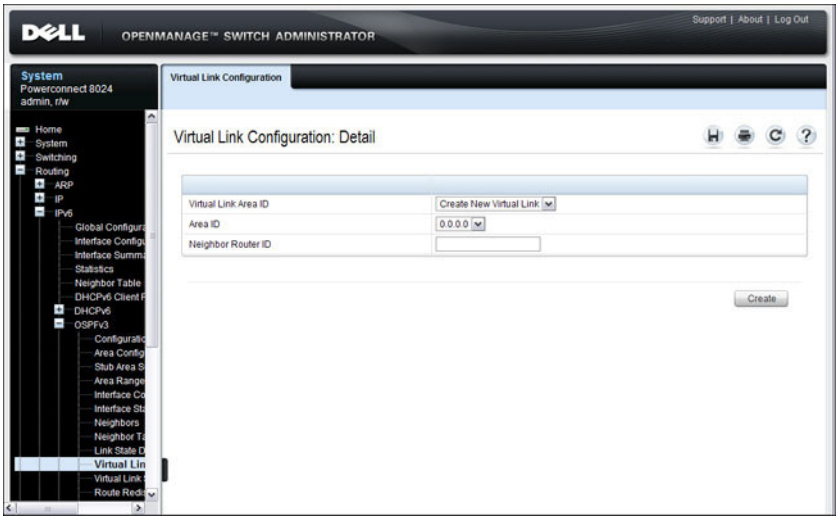


OSPFv3 Virtual Link Configuration

Use the **Virtual Link Configuration** page to define a new or configure an existing virtual link. To display this page, a valid OSPFv3 area must be defined through the OSPFv3 Area Configuration page.

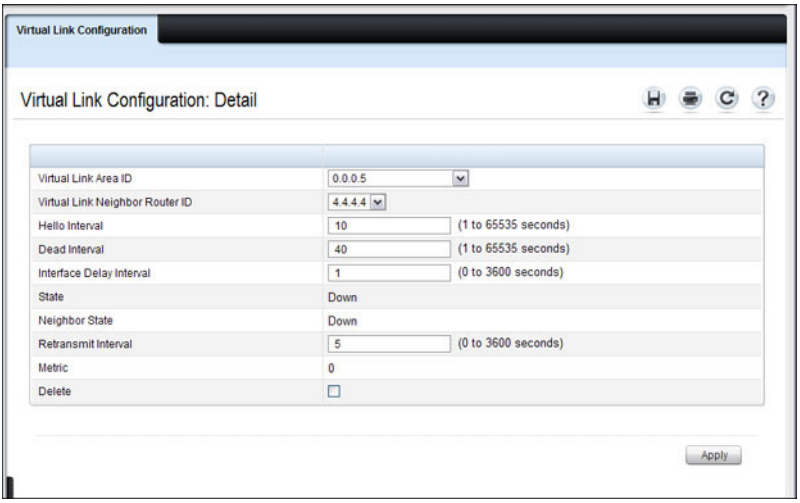
To display the page, click **IPv6 → OSPFv3 → Virtual Link Configuration** in the navigation panel.

Figure 35-29. OSPFv3 Virtual Link Configuration



After you create a virtual link, additional fields display, as the Figure 35-30 shows.

Figure 35-30. OSPFv3 Virtual Link Configuration

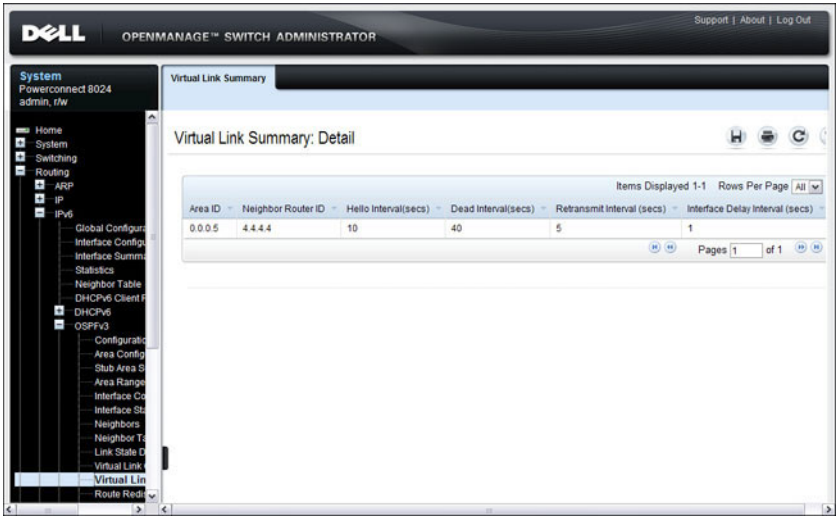


OSPFv3 Virtual Link Summary

Use the **Virtual Link Summary** page to display virtual link data by Area ID and Neighbor Router ID.

To display the page, click **IPv6 → OSPFv3 → Virtual Link Summary** in the navigation panel.

Figure 35-31. OSPFv3 Virtual Link Summary

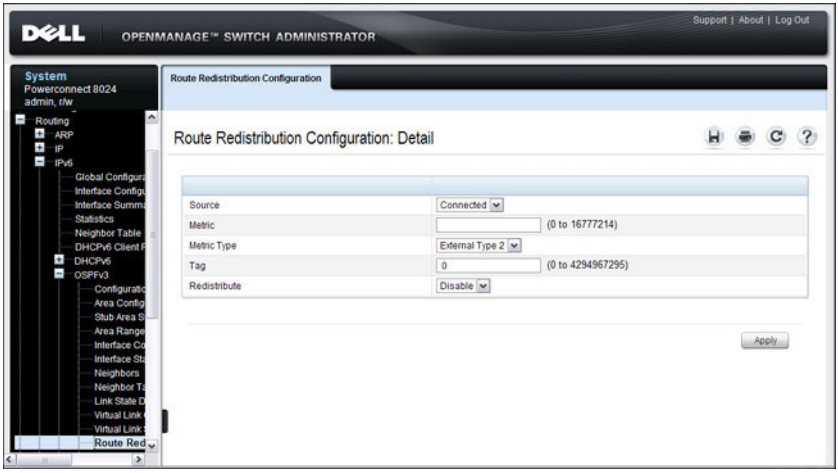


OSPFv3 Route Redistribution Configuration

Use the **Route Redistribution Configuration** page to configure route redistribution.

To display the page, click **IPv6 → OSPFv3 → Route Redistribution Configuration** in the navigation panel.

Figure 35-32. OSPFv3 Route Redistribution Configuration

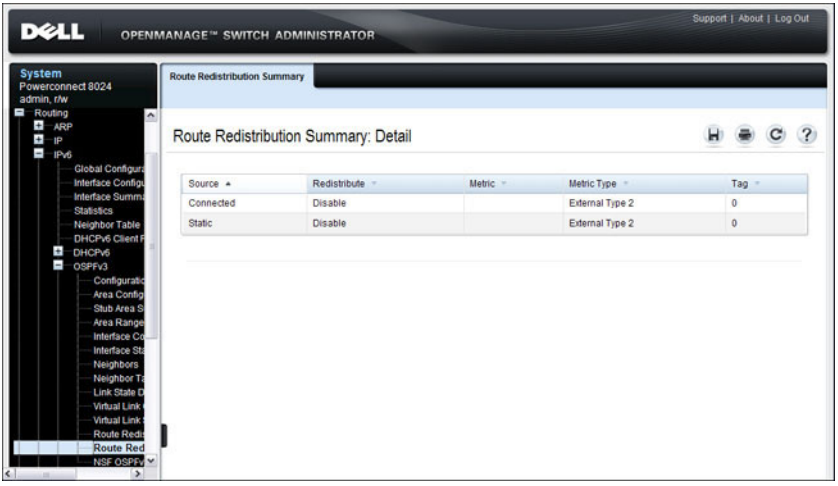


OSPFv3 Route Redistribution Summary

Use the **Route Redistribution Summary** page to display route redistribution settings by source.

To display the page, click **IPv6 → OSPFv3 → Route Redistribution Summary** in the navigation panel.

Figure 35-33. OSPFv3 Route Redistribution Summary

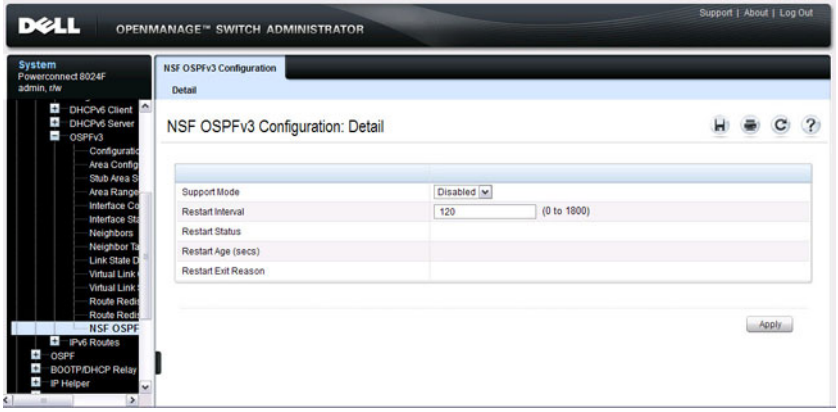


NSF OSPFv3 Configuration

Use the **NSF OSPFv3 Configuration** page to configure the non-stop forwarding (NSF) support mode and to view NSF summary information for the OSPFv3 feature. NSF is a feature used in switch stacks to maintain switching and routing functions in the event of a stack unit failure. For information about NSF, see "What is Nonstop Forwarding?" on page 147 in the Managing a Switch Stack chapter.

To display the page, click **Routing → OSPFv3 → NSF OSPFv3 Configuration** in the navigation panel.

Figure 35-34. NSF OSPFv3 Configuration



Configuring OSPF Features (CLI)

This section provides information about the commands you use to configure and view OSPF settings on the switch. This section does not describe all available **show** commands. For more information about all available OSPF commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Configuring Global OSPF Settings

Beginning in Privileged EXEC mode, use the following commands to configure various global OSPF settings for the switch.

Command	Purpose
configure	Enter global configuration mode.
router ospf	Enter OSPF configuration mode.
router-id *ip-address*	Set the 4-digit dotted-decimal number that uniquely identifies the router.
auto-cost reference-bandwidth *ref_bw*	Set the reference bandwidth used in the formula to compute link cost for an interface: link cost = ref_bw÷interface bandwidth The *ref_bw* variable is the reference bandwidth in Mbps (Range: 1–4294967).
capability opaque	Allow OSPF to store and flood opaque LSAs. An opaque LSA is used for flooding user defined information within an OSPF router domain.
compatible rfc1583	(Optional) Enable compatibility with RFC 1583. If all OSPF routers in the routing domain are capable of operating according to RFC 2328, OSPF 1583 compatibility mode should be disabled.

Command	Purpose
default-information originate [**always**] [**metric** *metric-value*] [**metric-type** *type-value*]	Control the advertisement of default routes. • **always** — Normally, OSPF originates a default route only if a default route is redistributed into OSPF (and default-information originate is configured). When the always option is configured, OSPF originates a default route, even if no default route is redistributed. • *metric-value* — The metric (or preference) value of the default route. (Range: 1–16777214) • *type-value* — The value is either 1 or 2: External type-1 route or External type-2 route.
default-metric *metric-value*	Set a default for the metric of distributed routes (Range: 1–16777214).
distance ospf {**external** | **inter-area** | **intra-area** } *distance*	Set the preference values of OSPF route types in the router. The range for the *distance* variable is 1–255. Lower route preference values are preferred when determining the best route.
enable	Enable OSPF.
exit-overflow-interval *seconds*	Specify the exit overflow interval for OSPF as defined in RFC 1765. The interval is the number of seconds after entering overflow state that a router will wait before attempting to leave the overflow state. (Range: 0–2147483647)
external-lsdb-limit *limit*	Configure the external LSDB limit for OSPF as defined in RFC 1765. If the value is –1, then there is no limit. The *limit* variable is the maximum number of non-default AS external-LSAs allowed in the router's link-state database. (Range: 1 to 2147483647)
maximum-paths *integer*	Set the number of paths that OSPF can report for a given destination (Range: 1–4).

Command	Purpose
passive-interface default	Configure OSPF interfaces as passive by default. This command overrides any interface-level passive mode settings.OSPF does not form adjacencies on passive interfaces but does advertise attached networks as stub networks.
timers spf *delay-time hold-time*	Specify the SPF delay and hold time. • *delay-time* — SPF delay time. (Range: 0–65535 seconds) • *hold-time* — SPF hold time. (Range: 0–65535 seconds)
exit	Exit to Global Configuration mode.
exit	Exit to Privileged EXEC mode.
show ip ospf	View OSPF global configuration and status.
show ip ospf statistics	View OSPF routing table calculation statistics.
clear ip ospf [{**configuration** | **redistribution** | **counters** | **neighbor** [**interface vlan** *vlan-id* [*neighbor-id*]]}]	Reset specific OSPF states. If no parameters are specified, OSPF is disabled and then re-enabled.

Configuring OSPF Interface Settings

Beginning in Privileged EXEC mode, use the following commands to configure per-interface OSPF settings.

Command	Purpose
configure	Enter global configuration mode.
interface vlan *vlan-id*	Enter Interface Configuration mode for the specified VLAN.
ip ospf area *area-id* [**secondaries none**]	Enables OSPFv2 on the interface and sets the area ID of an interface. This command supersedes the effects of network area command. The *area-id* variable is the ID of the area (Range: IP address or decimal from 0 –4294967295) Use the **secondaries none** keyword to prevent the interface from advertising its secondary addresses into the OSPFv2 domain.
ip ospf priority *number-value*	Set the OSPF priority for the interface. The *number-value* variable specifies the priority of an interface (Range: 0 to 255). The default priority is 1, which is the highest router priority. A value of 0 indicates that the router is not eligible to become the designated router on this network.
ip ospf retransmit-interval *seconds*	Set the OSPF retransmit interval for the interface. The *seconds* variable is the number of seconds between link-state advertisements for adjacencies belonging to this router interface. This value is also used when retransmitting database descriptions and link-state request packets. Valid values range from 0 to 3600 seconds (1 hour).
ip ospf hello-interval *seconds*	Set the OSPF hello interval for the interface. This parameter must be the same for all routers attached to a network. The *seconds* variable indicates the number of seconds to wait before sending Hello packets from the interface. (Range: 1–65535).

Command	Purpose
ip ospf dead-interval *seconds*	Set the OSPF dead interval for the interface. The *seconds* variable indicates the number of seconds a router waits to see a neighbor router's Hello packets before declaring that the router is down (Range: 1–65535). This parameter must be the same for all routers attached to a network. This value should be some multiple of the Hello Interval.
ip ospf transmit-delay *seconds*	Set the OSPF Transit Delay for the interface. The *seconds* variable sets the estimated number of seconds it takes to transmit a link state update packet over this interface. (Range: 1–3600 seconds)
ip ospf mtu-ignore	Disable OSPF MTU mismatch detection on the received database description.
ip ospf network {broadcast | point-to-point}	Set the OSPF network type on the interface to broadcast or point-to-point. OSPF selects a designated router and originates network LSAs only for broadcast networks. No more than two OSPF routers may be present on a point-to-point link.
ip ospf authentication {none | {simple *key***} | {encrypt** *key key-id***}}**	Set the OSPF Authentication Type and Key for the specified interface. • **encrypt** — MD5 encrypted authentication key. • *key* — Authentication key for the specified interface. (Range: 8 bytes or less if the authentication type is **simple** and 16 bytes or less if the type is **encrypt**.) • *key-id* — Authentication key identifier for the authentication type **encrypt**. (Range: 0–25)
ip ospf cost *interface-cost*	Set the metric cost of the interface. The *interface-cost* variable specifies the cost (link-state metric) of the OSPF interface. (Range: 1–65535)
bandwidth *bw*	Set the interface bandwidth used in the formula to compute link cost for an interface: link cost = ref_bw÷interface bandwidth The *bw* variable is the interface bandwidth (Range: 1–10000000 Kbps).

Command	Purpose
exit	Exit to Global Configuration Mode
router ospf	Enter OSPF configuration mode.
passive-interface vlan *vlan-id*	Make an interface passive to prevent OSPF from forming an adjacency on an interface. OSPF advertises networks attached to passive interfaces as stub networks.
network *ip-address wildcard-mask* **area** *area-id*	Enable OSPFv2 on interfaces whose primary IP address matches this command, and make the interface a member of the specified area. • *ip-address* — Base IPv4 address of the network area. • *wildcard-mask* — The network mask indicating the subnet. • *area-id* — The ID of the area (Range: IP address or decimal from 0–4294967295).
exit	Exit to Global Config mode.
exit	Exit to Privileged EXEC mode.
show ip ospf interface [**vlan** *vlan-id*]	View summary information for all OSPF interfaces configured on the switch or for the specified routing interface.
show ip ospf interface stats vlan *vlan-id*	View per-interface OSPF statistics.

Configuring Stub Areas and NSSAs

Beginning in Privileged EXEC mode, use the following commands to configure OSPF stub areas and NSSAs.

Command	Purpose
configure	Enter global configuration mode.
router ospf	Enter OSPF configuration mode.
area *area-id* **stub**	Create a stub area for the specified area ID.
area *area-id* **stub no-summary**	Prevent Summary LSAs from being advertised into the stub area.

Command	Purpose
area *area-id* **default-cost** *integer*	Configure the metric value (default cost) for the type 3 summary LSA sent into the stub area. Range: 1–16777215)
area *area-id* **nssa**	Create an NSSA for the specified area ID.
area *area-id* **nssa no-summary**	Configure the NSSA so that summary LSAs are not advertised into the NSSA.
area *area-id* **nssa translator-role** {**always** | **candidate**}	Configure the translator role of the NSSA. • **always** — The router assumes the role of the translator when it becomes a border router. • **candidate** — The router can participate in the translator election process when it attains border router status.
area *area-id* **nssa translator-stab-intv** *integer*	Configure the translator stability interval of the NSSA. The *integer* variable is the period of time that an elected translator continues to perform its duties after it determines that its translator status has been deposed by another router. (Range: 0–3600)
area *area-id* **nssa default-information-originate** [**metric** *metric-value*] [**metric-type** *metric-type-value*]	Configure the metric value and type for the default route advertised into the NSSA. The metric type can be comparable (nssa-external 1) or non-comparable (nssa-external 2).
area *area-id* **nssa no-redistribution**	Prevent learned external routes from being redistributed to the NSSA.
exit	Exit to Global Config mode.
exit	Exit to Privileged EXEC mode.
show ip ospf area *area-id*	View the configuration and status of an OSPF area.

Configuring Virtual Links

Beginning in Privileged EXEC mode, use the following commands to configure OSPF Virtual Links.

Command	Purpose
configure	Enter global configuration mode.
router ospf	Enter OSPF configuration mode.
area *area-id* **virtual-link** *neighbor-id*	Create the OSPF virtual interface for the specified area-id and neighbor router. The *neighbor-id* variable is the IP address of the neighboring router.
area *area-id* **virtual-link** *router-id* [**authentication** [**message-digest** | **null**]] [[**authentication-key** *key*] | [**message-digest-key** *key-id* **md5** *key*]]	Create the OSPF virtual interface for the specified area-id and neighbor router. Use the optional parameters to configure authentication for the virtual link. If the area has not been previously created, it is created by this command. If the area already exists, the virtual-link information is added or modified. • **authentication**—Specifies authentication type. • **message-digest**—Specifies that message-digest authentication is used. • **null**—No authentication is used. Overrides password or message-digest authentication if configured for the area. • **md5**—Use MD5 Encryption for an OSPF Virtual Link • *key*—Authentication key for the specified interface. (Range: 8 bytes or less if the authentication type is simple and 16 bytes or less if the type is encrypt.) • *key-id*—Authentication key identifier for the authentication type encrypt. (Range: 0-255)
area *area-id* **virtual-link** *neighbor-id* **retransmit-interval** *seconds*	Set the OSPF retransmit interval for the virtual link interface. The *seconds* variable is the number of seconds to wait between retransmitting LSAs if no acknowledgement is received. (Range: 0–3600)

Command	Purpose
area *area-id* **virtual-link** *neighbor-id* **hello-interval** *seconds*	Set the OSPF hello interval for the virtual link. The *seconds* variable indicates the number of seconds to wait before sending Hello packets from the virtual interface. (Range: 1–65535).
area *area-id* **virtual-link** *neighbor-id* **dead-interval** *seconds*	Set the OSPF dead interval for the virtual link. The *seconds* variable indicates the number of seconds to wait before the virtual interface is assumed to be dead. (Range: 1–65535)
area *area-id* **virtual-link** *neighbor-id* **transmit-delay** *seconds*	Set the OSPF Transit Delay for the interface. The *seconds* variable is the number of seconds to increment the age of the LSA before sending, based on the estimated time it takes to transmit from the interface. (Range: 0–3600)
exit	Exit to Global Config mode.
exit	Exit to Privileged EXEC mode.
show ip ospf virtual-link brief	View summary information about all virtual links configured on the switch.

Configuring OSPF Area Range Settings

Beginning in Privileged EXEC mode, use the following commands to configure an OSPF area range.

Command	Purpose
configure	Enter global configuration mode.
router ospf	Enter OSPF configuration mode.
area *area-id* **range** *ip-address mask* {**summarylink** | **nssaexternallink**} [**advertise** |**not-advertise**]	Configure a summary prefix for routes learned in a given area. • *area-id* — Identifies the OSPF NSSA to configure. (Range: IP address or decimal from 0–4294967295) • *ip-address* — IP address. • *subnet-mask* — Subnet mask associated with IP address. • **summarylink** — Specifies a summary link LSDB type. • **nssaexternallink** — Specifies an NSSA external link LSDB type. • **advertise** — Advertisement of the area range. • **not-advertise** — Suppresses advertisement of the area range.
exit	Exit to Global Config mode.
exit	Exit to Privileged EXEC mode.
show ip ospf range *area-id*	View information about the area ranges for the specified *area-id*.

Configuring OSPF Route Redistribution Settings

Beginning in Privileged EXEC mode, use the following commands to configure OSPF route redistribution settings.

Command	Purpose
configure	Enter global configuration mode.
router ospf	Enter OSPF configuration mode.

Command	Purpose
distribute-list *accesslistname* **out** {**rip** | **static** | **connected**}	Specify the access list to filter routes received from the source protocol. The ACL must already exist on the switch. For information about the commands you use to configure ACLs, see "Configuring ACLs (CLI)" on page 521. • *accesslistname* — The name used to identify an existing ACL. • **rip** — Apply the specified access list when RIP is the source protocol. • **static** — Apply the specified access list when packets come through the static route. • **connected** — Apply the specified access list when packets come from a directly connected route.
redistribute {**rip** | **static** | **connected**} [**metric** *integer*] [**metric-type** {1 | 2}] [**tag** *integer*] [**subnets**]	Configure OSPF to allow redistribution of routes from the specified source protocol/routers. • **rip** — Specifies RIP as the source protocol. • **static** — Specifies that the source is a static route. • **connected** — Specifies that the source is a directly connected route. • *metric* — Specifies the metric to use when redistributing the route. (Range: 0–16777214) • **metric-type 1** — Type 1 external route. • **metric-type 2** — Type 2 external route. • *tag* — Value attached to each external route. (Range: 0–4294967295) • **subnets**—Unless this keyword is configured, OSPF distributes only class A, class B, and class C prefixes.
exit	Exit to Global Config mode.
exit	Exit to Privileged EXEC mode.
show ip ospf	View OSPF configuration and status information, including route distribution information.

Configuring OSPFv3 Features (CLI)

This section provides information about the commands you use to configure OSPFv3 settings on the switch. For more information about the commands and about additional **show** commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Configuring Global OSPFv3 Settings

Beginning in Privileged EXEC mode, use the following commands to configure various global OSPFv3 settings for the switch.

Command	Purpose
configure	Enter global configuration mode.
ipv6 router ospf	Enter OSPFv3 configuration mode.
router-id *ip-address*	Set the 4-digit dotted-decimal number that uniquely identifies the router.
auto-cost reference-bandwidth *ref_bw*	Set the reference bandwidth used in the formula to compute link cost for an interface: link cost = ref_bw÷interface bandwidth The *ref_bw* variable is the reference bandwidth in Mbps (Range: 1–4294967).
default-information originate [**always**] [**metric** *metric-value*] [**metric-type** *type-value*]	Control the advertisement of default routes. • **always** — Normally, OSPFv3 originates a default route only if a default route is redistributed into OSPFv3 (and default-information originate is configured). When the always option is configured, OSPFv3 originates a default route, even if no default route is redistributed. • *metric-value* — The metric (or preference) value of the default route. (Range: 1–16777214) • *type-value* — The value is either 1 or 2: External type-1 route or External type-2 route.
default-metric *metric-value*	Set a default for the metric of distributed routes. (Range: 1–16777214).

Command	Purpose
distance ospf {**external** | **inter-area** | **intra-area** } *distance*	Set the preference values of OSPFv3 route types in the router. The range for the *distance* variable is 1–255. Lower route preference values are preferred when determining the best route.
enable	Enable OSPFv3.
exit-overflow-interval *seconds*	Specify the exit overflow interval for OSPFv3 as defined in RFC 1765. The interval is the number of seconds after entering overflow state that a router will wait before attempting to leave the overflow state. (Range: 0–2147483647)
external-lsdb-limit *limit*	Configure the external LSDB limit for OSPFv3 as defined in RFC 1765. If the value is -1, then there is no limit. The *limit* variable is the maximum number of non-default AS external-LSAs allowed in the router's link-state database. (Range: –1 to 2147483647)
maximum-paths *maxpaths*	Set the number of paths that OSPFv3 can report for a given destination. (Range: 1–4.)
passive-interface default	Configure OSPFv3 interfaces as passive by default. This command overrides any interface-level passive mode settings. OSPFv3 does not form adjacencies on passive interfaces but does advertise attached networks as stub networks.
exit	Exit to Global Configuration mode.
exit	Exit to Privileged EXEC mode.
show ipv6 ospf	View OSPFv3 global configuration and status.
clear ipv6 ospf [{**configuration** | **redistribution** | **counters** | **neighbor** [**interface vlan** *vlan-id* [*neighbor-id*]]}]	Reset specific OSPFv3 states. If no parameters are specified, OSPFv3 is disabled and then re-enabled.

Configuring OSPFv3 Interface Settings

Beginning in Privileged EXEC mode, use the following commands to configure per-interface OSPFv3 settings.

Command	Purpose
configure	Enter global configuration mode.
interface vlan *vlan-id*	Enter Interface Configuration mode for the specified VLAN.
ipv6 ospf areaid *area-id*	Enables OSPFv3 on the interface and sets the area ID of an interface. This command supersedes the effects of network area command. The *area-id* variable is the ID of the area (Range: IP address or decimal from 0 –4294967295)
ipv6 ospf priority *number-value*	Set the OSPFv3 priority for the interface. The *number-value* variable specifies the priority of an interface (Range: 0 to 255). The default priority is 1, which is the highest router priority. A value of 0 indicates that the router is not eligible to become the designated router on this network.
ipv6 ospf retransmit-interval *seconds*	Set the OSPFv3 retransmit interval for the interface. The *seconds* variable is the number of seconds between link-state advertisements for adjacencies belonging to this router interface. This value is also used when retransmitting database descriptions and link-state request packets. Valid values range from 0 to 3600 seconds (1 hour).
ipv6 ospf hello-interval *seconds*	Set the OSPFv3 hello interval for the interface. This parameter must be the same for all routers attached to a network. The *seconds* variable indicates the number of seconds to wait before sending Hello packets from the interface. (Range: 1–65535).

Command	Purpose
ipv6 ospf dead-interval *seconds*	Set the OSPFv3 dead interval for the interface. The *seconds* variable indicates the number of seconds a router waits to see a neighbor router's Hello packets before declaring that the router is down (Range: 1–65535). This parameter must be the same for all routers attached to a network. This value should be some multiple of the Hello Interval.
ipv6 ospf transmit-delay *seconds*	Set the OSPFv3 Transit Delay for the interface. The *seconds* variable sets the estimated number of seconds it takes to transmit a link state update packet over this interface. (Range: 1–3600 seconds)
ip ospf mtu-ignore	Disable OSPFv3 MTU mismatch detection on received database description packets.
ipv6 ospf network {broadcast | point-to-point }	Set the OSPFv3 network type on the interface to broadcast or point-to-point. OSPFv3 selects a designated router and originates network LSAs only for broadcast networks. No more than two OSPFv3 routers may be present on a point-to-point link.
ipv6 ospf cost *interface-cost*	Set the metric cost of the interface. The *interface-cost* variable specifies the cost (link-state metric) of the OSPFv3 interface. (Range: 1–65535)
bandwidth *bw*	Set the interface bandwidth used in the formula to compute link cost for an interface: link cost = ref_bw÷interface bandwidth The *bw* variable is the interface bandwidth (Range: 1–10000000 Kbps).
exit	Exit to Global Configuration Mode
ipv6 router ospf	Enter OSPFv3 configuration mode.
passive-interface {vlan *vlan-id* | tunnel *tunnel-id*}	Make an interface passive to prevent OSPFv3 from forming an adjacency on an interface. OSPFv3 advertises networks attached to passive interfaces as stub networks.
exit	Exit to Global Config mode.
exit	Exit to Privileged EXEC mode.

Command	Purpose
show ipv6 ospf interface [*interface-type interface-number*]	View summary information for all OSPFv3 interfaces configured on the switch or for the specified routing interface.
show ipv6 ospf interface stats *interface-type interface-number*	View per-interface OSPFv3 statistics.

Configuring Stub Areas and NSSAs

Beginning in Privileged EXEC mode, use the following commands to configure OSPFv3 stub areas and NSSAs.

Command	Purpose
configure	Enter global configuration mode.
ipv6 router ospf	Enter OSPFv3 configuration mode.
area *area-id* **stub**	Create a stub area for the specified area ID.
area *area-id* **stub no-summary**	Prevent Summary LSAs from being advertised into the stub area.
area *area-id* **default-cost** *cost*	Configure the metric value (default cost) for the type 3 summary LSA sent into the stub area. Range: 1–16777215)

Command	Purpose

area *area-id* **nssa** [**no-redistribution**] [**default-information-originate** [**metric** *metric-value*] [**metric-type** *metric-type-value*]] [**no-summary**] [**translator-role** *role*] [**translator-stab-intv** *interval*]	Create and configure an NSSA for the specified area ID.
	• *metric-value*—Specifies the metric of the default route advertised to the NSSA. (Range: 1–16777214)
	• *metric-type-value*—The metric type can be one of the following :
	  • A metric type of nssa-external 1 (comparable)
	  • A metric type of nssa-external 2 (non-comparable)
	• **no-summary—**Summary LSAs are not advertised into the NSSA
	• **role**—The translator role where role is one of the following :
	  • **always**—The router assumes the role of the translator when it becomes a border router.
	  • **candidate**—The router to participate in the translator election process when it attains border router status.
	• **interval**—The period of time that an elected translator continues to perform its duties after it determines that its translator status has been deposed by another router. (Range: 0–3600)

area *area-id* **nssa no-redistribution**	Prevent learned external routes from being redistributed to the NSSA.
exit	Exit to Global Config mode.
exit	Exit to Privileged EXEC mode.
show ipv6 ospf area *area-id*	Show configuration and status of an OSPF area.

Configuring Virtual Links

Beginning in Privileged EXEC mode, use the following commands to configure OSPFv3 Virtual Links.

Command	Purpose
configure	Enter global configuration mode.
ipv6 router ospf	Enter OSPFv3 configuration mode.
area *area-id* **virtual-link** *neighbor-id*	Create the OSPFv3 virtual interface for the specified *area-id* and neighbor router. The *neighbor-id* variable is the IP address of the neighboring router.
area *area-id* **virtual-link** *neighbor-id* **retransmit-interval** *seconds*	Set the OSPFv3 retransmit interval for the virtual link interface. The *seconds* variable is the number of seconds to wait between retransmitting LSAs if no acknowledgement is received. (Range: 0–3600)
area *area-id* **virtual-link** *neighbor-id* **hello-interval** *seconds*	Set the OSPFv3 hello interval for the virtual link. The *seconds* variable indicates the number of seconds to wait before sending Hello packets from the virtual interface. (Range: 1–65535).
area *area-id* **virtual-link** *neighbor-id* **dead-interval** *seconds*	Set the OSPFv3 dead interval for the virtual link. The *seconds* variable indicates the number of seconds to wait before the virtual interface is assumed to be dead. (Range: 1–65535)
area *area-id* **virtual-link** *neighbor-id* **transmit-delay** *seconds*	Set the OSPFv3 Transit Delay for the interface. The *seconds* variable is the number of seconds to increment the age of the LSA before sending, based on the estimated time it takes to transmit from the interface. (Range: 0–3600)
exit	Exit to Global Config mode.
exit	Exit to Privileged EXEC mode.
show ipv6 ospf virtual-link brief	View summary information about all virtual links configured on the switch.

Configuring an OSPFv3 Area Range

Beginning in Privileged EXEC mode, use the following commands to configure an OSPFv3 area range.

Command	Purpose
configure	Enter global configuration mode.
ipv6 router ospf	Enter OSPFv3 configuration mode.
area *area-id* **range** *ipv6-prefix/prefix-length* **{summarylink | nssaexternallink} [advertise |not-advertise]**	Configure a summary prefix for routes learned in a given area. • *area-id* — Identifies the OSPFv3 NSSA to configure. (Range: IP address or decimal from 0–4294967295) • *ipv6-prefix/prefix-length* — IPv6 address and prefix length. • **summarylink** — Specifies a summary link LSDB type. • **nssaexternallink** — Specifies an NSSA external link LSDB type. • **advertise** — Advertisement of the area range. • **not-advertise** — Suppresses advertisement of the area range.
exit	Exit to Global Config mode.
exit	Exit to Privileged EXEC mode.
show ipv6 ospf range *area-id*	View information about the area ranges for the specified area-id.

Configuring OSPFv3 Route Redistribution Settings

Beginning in Privileged EXEC mode, use the following commands to configure OSPFv3 route redistribution settings.

Command	Purpose
configure	Enter global configuration mode.
ipv6 router ospf	Enter OSPFv3 configuration mode.
redistribute {**static** | **connected**} [**metric** *metric*] [**metric-type** {1 | 2}] [**tag** *tag*]	Configure OSPFv3 to allow redistribution of routes from the specified source protocol/routers. • **static** — Specifies that the source is a static route. • **connected** — Specifies that the source is a directly connected route. • *metric* — Specifies the metric to use when redistributing the route. (Range: 0–16777214) • **metric-type 1** — Type 1 external route. • **metric-type 2** — Type 2 external route. • *tag* — Value attached to each external route, which might be used to communicate information between ASBRs. (Range: 0–4294967295)
exit	Exit to Global Config mode.
exit	Exit to Privileged EXEC mode.
show ipv6 ospf	View OSPFv3 configuration and status information, including information about redistributed routes.

OSPF Configuration Examples

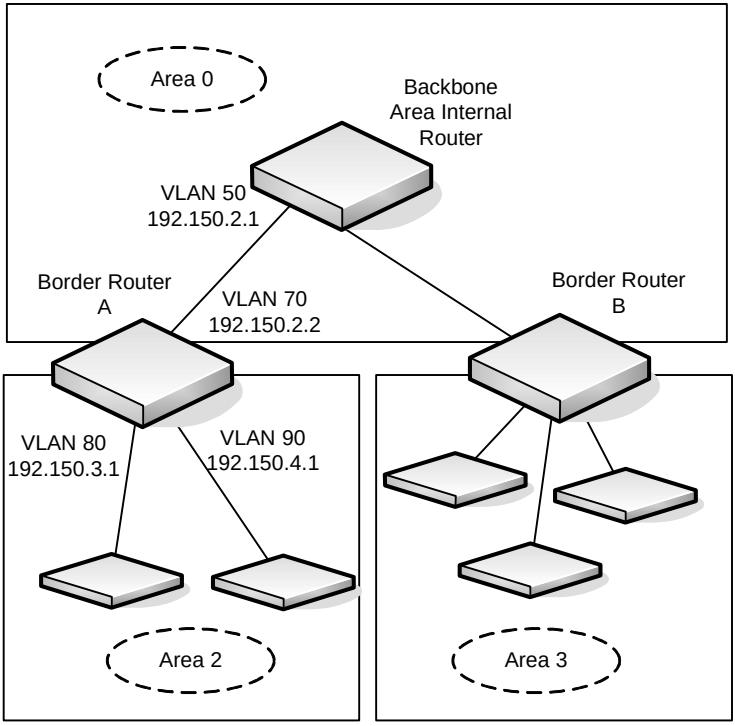
This section contains the following examples:

- Configuring an OSPF Border Router and Setting Interface Costs
- Configuring Stub and NSSA Areas for OSPF and OSPFv3
- Configuring a Virtual Link for OSPF and OSPFv3

Configuring an OSPF Border Router and Setting Interface Costs

This example shows how to configure the PowerConnect switch as an OSPF border router. The commands in this example configure the areas and interfaces on Border Router A shown in Figure 35-35.

Figure 35-35. OSPF Area Border Router



To Configure Border Router A:

1. Enable routing on the switch.

```
console#configure
console(config)#ip routing
```

2. Create VLANS 70, 80, and 90.

```
console(config)#vlan 70,80,90
```

3. Assign IP addresses for VLANs 70, 80 and 90.

```
console(config)#interface vlan 70
console(config-if-vlan70)#ip address 192.150.2.2
255.255.255.0
console(config-if-vlan70)#exit

console(config)#interface vlan 80
console(config-if-vlan80)#ip address 192.150.3.1
255.255.255.0
console(config-if-vlan80)#exit

console(config)#interface vlan 90
console(config-if-vlan90)#ip address 192.150.4.1
255.255.255.0
console(config-if-vlan90)#exit
```

4. Enable OSPF on the switch and specify a router ID.

```
console(config)#router ospf
console(config-router)#router-id 192.150.9.9
console(config-router)#exit
```

5 Configure the OSPF area ID, priority, and cost for each interface.



NOTE: OSPF is globally enabled by default. To make it operational on the router, you configure OSPF for particular interfaces and identify which area the interface is associated with.

```
console(config)#interface vlan 70
console(config-if-vlan70)#ip ospf area 0.0.0.0
console(config-if-vlan70)#ip ospf priority 128
console(config-if-vlan70)#ip ospf cost 32
console(config-if-vlan70)#exit

console(config)#interface vlan 80
console(config-if-vlan80)#ip ospf area 0.0.0.2
console(config-if-vlan80)#ip ospf priority 255
console(config-if-vlan80)#ip ospf cost 64
console(config-if-vlan80)#exit

console(config)#interface vlan 90
console(config-if-vlan90)#ip ospf area 0.0.0.2
console(config-if-vlan90)#ip ospf priority 255
console(config-if-vlan90)#ip ospf cost 64
console(config-if-vlan90)#exit
```

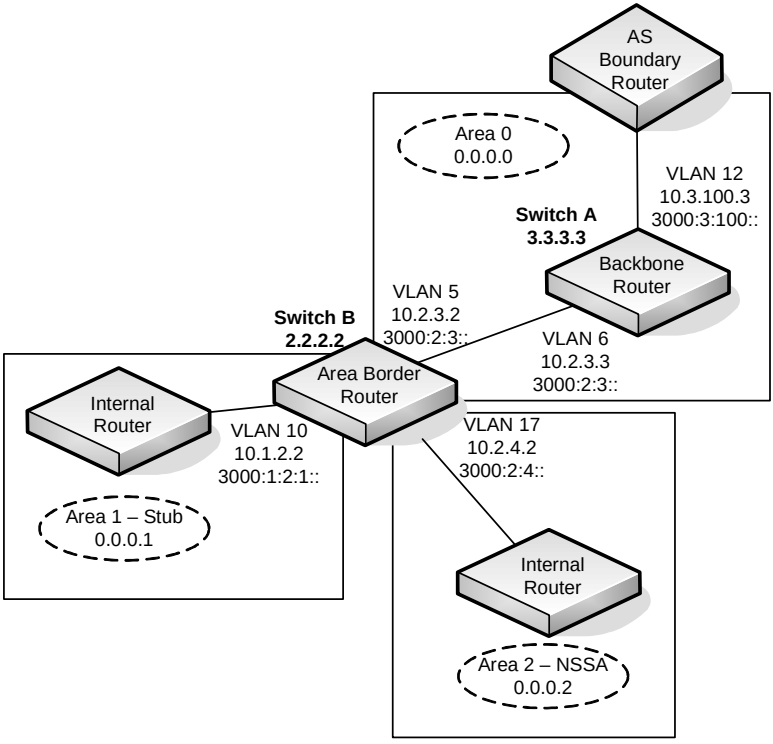
Configuring Stub and NSSA Areas for OSPF and OSPFv3

In this example, Area 0 connects directly to two other areas: Area 1 is defined as a stub area and Area 2 is defined as an NSSA area.

NOTE: OSPFv2 and OSPFv3 can operate concurrently on a network and on the same interfaces (although they do not interact). This example configures both protocols simultaneously.

Figure 35-36 illustrates this example OSPF configuration.

Figure 35-36. OSPF Configuration—Stub Area and NSSA Area



Switch A is a backbone router. It links to an ASBR (not defined here) that routes traffic outside the AS.

To configure Switch A:

1. Globally enable IPv6 and IPv4 routing:

```
console#configure
console(config)#ipv6 unicast-routing
console(config)#ip routing
```

2. Create VLANs 6 and 12.

```
console(config)#vlan 6,12
```

3. Configure IP and IPv6 addresses on VLAN routing interface 6.

```
console(config-if)#interface vlan 6
console(config-if-vlan6)#ip address 10.2.3.3
255.255.255.0
console(config-if-vlan6)#ipv6 address
3000:2:3::/64 eui64
```

4. Associate the interface with area 0.0.0.0 and enable OSPFv3.

```
console(config-if-vlan6)#ip ospf area 0.0.0.0
console(config-if-vlan6)#ipv6 ospf
console(config-if-vlan6)#exit
```

5. Configure IP and IPv6 addresses on VLAN routing interface 12.

```
console(config)#interface vlan 12
console(config-if-vlan12)#ip address  10.3.100.3
255.255.255.0
console(config-if-vlan12)#ipv6 address
3000:3:100::/64 eui64
```

6. Associate the interface with area 0.0.0.0 and enable OSPFv3.

```
console(config-if-vlan12)#ip ospf area 0.0.0.0
console(config-if-vlan12)#ipv6 ospf
console(config-if-vlan12)#exit
```

7. Define the OSPF and OSPFv3 router IDs for the switch:

```
console(config)#ipv6 router ospf
console(config-rtr)#router-id 3.3.3.3
console(config-rtr)#exit
```

```
console(config)#router ospf
console(config-router)#router-id 3.3.3.3
console(config-router)#exit
```

Switch B is a ABR that connects Area 0 to Areas 1 and 2.

To configure Switch B:

1. Configure IPv6 and IPv4 routing. The static routes are included for illustration only: Redistributed static routes, like routes distributed from other protocols, are not injected into stub areas such as Area 1:

   ```
   console#configure
   console(config)#ipv6 unicast-routing

   console(config)#ipv6 route 3000:44:44::/64
   3000:2:3::210:18ff:fe82:c14
   console(config)#ip route 10.23.67.0 255.255.255.0
   10.2.3.3
   ```

2. Create VLANs 5, 10, and 17.

   ```
   console(config)#vlan 5,10,17
   ```

3. On VLANs 5, 10, and 17, configure IPv4 and IPv6 addresses and enable OSPFv3. For IPv6, associate VLAN 5 with Area 0, VLAN 10 with Area 1, and VLAN 17 with Area 2.

   ```
   console(config)#interface vlan 5
   console(config-if-vlan5)#ip address  10.2.3.2
   255.255.255.0
   console(config-if-vlan5)#ipv6 address
   3000:2:3::/64 eui64
   console(config-if-vlan5)#ipv6 ospf
   console(config-if-vlan5)#ipv6 ospf areaid 0
   console(config-if-vlan5)#exit
   console(config)#interface vlan 10
   console(config-if-vlan10)#ip address  10.1.2.2
   255.255.255.0
   console(config-if-vlan10)#ipv6 address
   3000:1:2::/64 eui64
   console(config-if-vlan10)#ipv6 ospf
   console(config-if-vlan10)#ipv6 ospf areaid 1
   console(config-if-vlan10)#exit
   ```

```
console(config)#interface vlan 17
console(config-if-vlan17)#ip address  10.2.4.2
255.255.255.0
console(config-if-vlan17)#ipv6 address
3000:2:4::/64 eui64
console(config-if-vlan17)#ipv6 ospf
console(config-if-vlan17)#ipv6 ospf areaid 2
console(config-if-vlan17)#exit
```

4 For IPv4: Configure the router ID, define an OSPF router, and define Area 1 as a stub., and define Area 2 as an NSSA.

```
console(config)#router ospf
console(config-router)#router-id 2.2.2.2
console(config-router)#area 0.0.0.1 stub
console(config-router)#area 0.0.0.2 nssa
```

5 For IPv4: Enable OSPF for IPv4 on VLANs 10, 5, and 17 by globally defining the range of IP addresses associated with each interface, and then associating those ranges with Areas 1, 0, and 2, respectively.

```
console(config-router)#network 10.1.2.0 0.0.0.255
area 0.0.0.1
console(config-router)#network 10.2.3.0 0.0.0.255
area 0.0.0.0
console(config-router)#network 10.2.4.0 0.0.0.255
area 0.0.0.2
```

6 For IPv4: Configure a metric cost to associate with static routes when they are redistributed via OSPF:

```
console(config-router)#redistribute static metric
1 subnets
console(config-router)#exit
```

7 For IPv6: Define an OSPF router. Define Area 1 as a stub and area 2 as a Not-So-Stubby-Area (NSSA). Configure a metric cost to associate with static routes when they are redistributed via OSPF:

```
console(config)#ipv6 router ospf
console(config-rtr)#router-id 2.2.2.2
console(config-rtr)#area 0.0.0.1 stub
console(config-rtr)#area 0.0.0.2 nssa
```

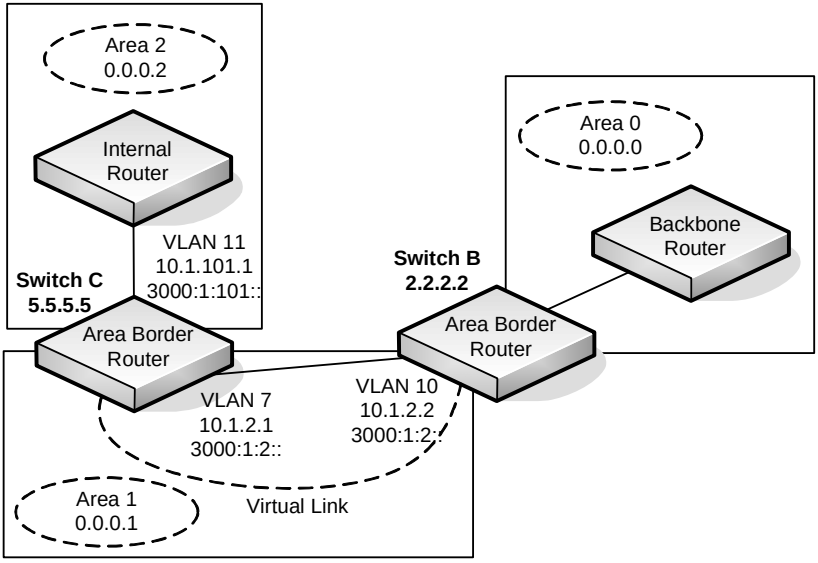
```
console(config-rtr)#redistribute static metric 105
metric-type 1
console(config-rtr)#exit
```

Configuring a Virtual Link for OSPF and OSPFv3

In this example, Area 0 connects directly to Area 1. A virtual link is defined that traverses Area 1 and connects to Area 2. This example assumes other OSPF settings, such as area and interface configuration, have already been configured.

Figure 35-37 illustrates the relevant components in this example OSPF configuration.

Figure 35-37. OSPF Configuration—Virtual Link



Switch B is an ABR that directly connects Area 0 to Area 1. Note that in the previous example, Switch B connected to a stub area and an NSSA. Virtual links cannot be created across stub areas or NSSAs.

The following commands define a virtual link that traverses Area 1 to Switch C (5.5.5.5).

To configure Switch B:

1. Configure the virtual link to Switch C for IPv4.

```
console#configure
console(config)#router ospf
console(config-router)#area 0.0.0.1 virtual-link
5.5.5.5
console(config-router)#exit
```

2. Configure the virtual link to Switch C for IPv6.

```
console#configure
console(config)#ipv6 router ospf
console(config-rtr)#area 0.0.0.1 virtual-link
5.5.5.5
console(config-rtr)#exit
```

Switch C is a ABR that enables a virtual link from the remote Area 2 in the AS to Area 0. The following commands define a virtual link that traverses Area 1 to Switch B (2.2.2.2).

To configure Switch C:

1. For IPv4, assign the router ID, create the virtual link to Switch B, and associate the VLAN routing interfaces with the appropriate areas.

```
console(config)#router ospf
console(config-router)#area 0.0.0.1 virtual-link
2.2.2.2
console(config-router)#exit
```

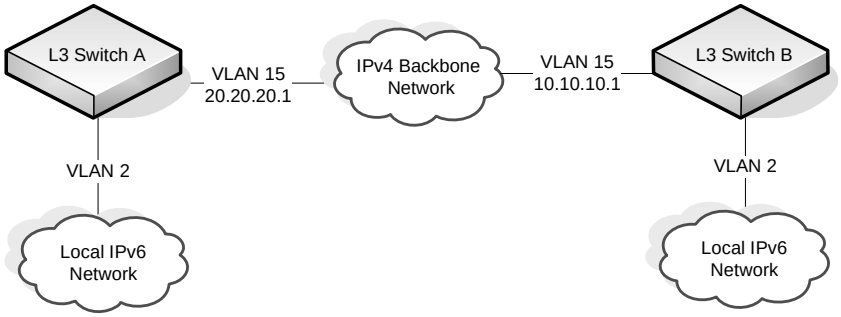
2. For IPv6, assign the router ID and create the virtual link to Switch B.

```
console(config)#ipv6 router ospf
console(config-rtr)#area 0.0.0.1 virtual-link
2.2.2.2
console(config-rtr)#exit
```

Interconnecting an IPv4 Backbone and Local IPv6 Network

In Figure 35-38, two PowerConnect L3 switches are connected as shown in the diagram. The VLAN 15 routing interface on both switches connects to an IPv4 backbone network where OSPF is used as the dynamic routing protocol to exchange IPv4 routes. OSPF allows device 1 and device 2 to learn routes to each other (from the 20.20.20.x network to the 10.10.10.x network and vice versa). The VLAN 2 routing interface on both devices connects to the local IPv6 network. OSPFv3 is used to exchange IPv6 routes between the two devices. The tunnel interface allows data to be transported between the two remote IPv6 networks over the IPv4 network.

Figure 35-38. IPv4 and IPv6 Interconnection Example



To configure Switch A:

1. Create the VLANs.

   ```
   console(config)#vlan 2,15
   ```

2. Enable IPv4 and IPv6 routing on the switch.

   ```
   console(config)#ip routing
   console(config)#ipv6 unicast-routing
   ```

3. Set the OSPF router ID.

   ```
   console(config)#router ospf
   console(config-router)#router-id 1.1.1.1
   console(config-router)#exit
   ```

4 Set the OSPFv3 router ID.

```
console(config)#ipv6 router ospf
console(config-rtr)#router-id 1.1.1.1
console(config-rtr)#exit
```

5 Configure the IPv4 address and OSPF area for VLAN 15.

```
console(config)#interface vlan 15
console(config-if-vlan15)#ip address 20.20.20.1
255.255.255.0
console(config-if-vlan15)#ip ospf area 0.0.0.0
console(config-if-vlan15)#exit
```

6 Configure the IPv6 address and OSPFv3 information for VLAN 2.

```
console(config)#interface vlan 2
console(config-if-vlan2)#ipv6 address 2020:1::1/64
console(config-if-vlan2)#ipv6 ospf
console(config-if-vlan2)#ipv6 ospf network point-
to-point
console(config-if-vlan2)#exit
```

7 Configure the tunnel.

```
console(config)#interface tunnel 0
console(config-if-tunnel0)#ipv6 address 2001::1/64
console(config-if-tunnel0)#tunnel mode ipv6ip
console(config-if-tunnel0)#tunnel source
20.20.20.1
console(config-if-tunnel0)#tunnel destination
10.10.10.1
console(config-if-tunnel0)#ipv6 ospf
console(config-if-tunnel0)#ipv6 ospf network
point-to-point
console(config-if-tunnel0)#exit
```

8 Configure the loopback interface. The switch uses the loopback IP address as the OSPF and OSPFv3 router ID.

```
console(config)#interface loopback 0
console(config-if-loopback0)#ip address 1.1.1.1
255.255.255.0
console(config-if-loopback0)#exit
console(config)#exit
```

To configure Switch B:

1. Create the VLANs.

```
console(config)#vlan 2,15
```

2. Enable IPv4 and IPv6 routing on the switch.

```
console(config)#ip routing
console(config)#ipv6 unicast-routing
```

3. Set the OSPF router ID.

```
console(config)#router ospf
console(config-router)#router-id 2.2.2.2
console(config-router)#exit
```

4. Set the OSPFv3 router ID.

```
console(config)#ipv6 router ospf
console(config-rtr)#router-id 2.2.2.2
console(config-rtr)#exit
```

5. Configure the IPv4 address and OSPF area for VLAN 15.

```
console(config)#interface vlan 15
console(config-if-vlan15)#ip address 10.10.10.1
255.255.255.0
console(config-if-vlan15)#ip ospf area 0.0.0.0
console(config-if-vlan15)#exit
```

6. Configure the IPv6 address and OSPFv3 information for VLAN 2.

```
console(config)#interface vlan 2
console(config-if-vlan2)#ipv6 address 2020:2::2/64
console(config-if-vlan2)#ipv6 ospf
console(config-if-vlan2)#ipv6 ospf network point-
to-point
console(config-if-vlan2)#exit
```

7. Configure the tunnel.

```
console(config)#interface tunnel 0
console(config-if-tunnel0)#ipv6 address 2001::2/64
console(config-if-tunnel0)#tunnel mode ipv6ip
console(config-if-tunnel0)#tunnel source
10.10.10.1
```

```
console(config-if-tunnel0)#tunnel destination
20.20.20.1
console(config-if-tunnel0)#ipv6 ospf
console(config-if-tunnel0)#ipv6 ospf network
point-to-point
console(config-if-tunnel0)#exit
```

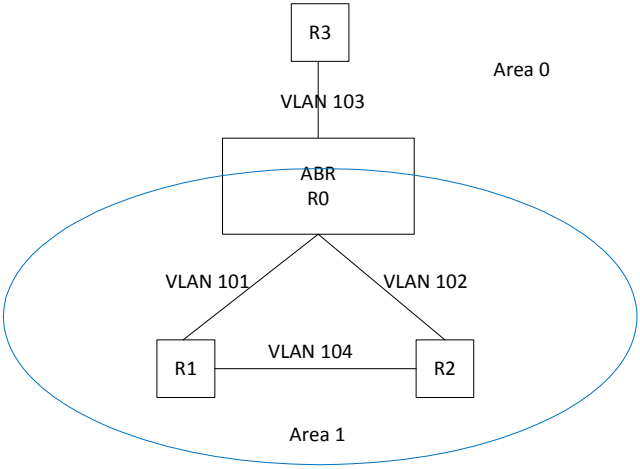
8 Configure the loopback interface. The switch uses the loopback IP address as the OSPF and OSPFv3 router ID.

```
console(config)#interface loopback 0
console(config-if-loopback0)#ip address 2.2.2.2
255.255.255.0
console(config-if-loopback0)#exit
console(config)#exit
```

Configuring the Static Area Range Cost

Figure 35-39 shows a topology for the configuration that follows.

Figure 35-39. Static Area Range Cost Example Topology



1 Configure R0.

```
terminal length 0
config
```

```
hostname ABR-R0
line console
exec-timeout 0
exit
vlan 101-103
exit
ip routing
router ospf
router-id 10.10.10.10
network 172.20.0.0 0.0.255.255 area 0
network 172.21.0.0 0.0.255.255 area 1
area 1 range 172.21.0.0 255.255.0.0 summarylink
timers spf 3 5
exit
interface vlan 101
ip address 172.21.1.10 255.255.255.0
ip ospf hello-interval 1
ip ospf dead-interval 4
ip ospf network point-to-point
exit
interface te1/0/21
switchport mode trunk
description "R1"
exit
interface vlan 102
ip address 172.21.2.10 255.255.255.0
ip ospf hello-interval 1
ip ospf dead-interval 4
ip ospf network point-to-point
exit
interface te1/0/22
description "R2"
switchport mode trunk
exit
interface vlan 103
ip address 172.20.1.10 255.255.255.0
ip ospf hello-interval 1
ip ospf dead-interval 4
ip ospf network point-to-point
exit
interface te1/0/23
switchport mode trunk
description "R3"
exit
exit
```

2 Configure R1.

```
terminal length 0
config
hostname R1
line console
exec-timeout 0
exit
vlan 101,104
exit
ip routing
router ospf
router-id 1.1.1.1
network 172.21.0.0 0.0.255.255 area 1
timers spf 3 5
exit
interface vlan 101
ip address 172.21.1.1 255.255.255.0
routing
ip ospf hello-interval 1
ip ospf dead-interval 4
ip ospf network point-to-point
exit
interface te1/0/21
switchport mode trunk
exit
interface vlan 104
ip address 172.21.3.1 255.255.255.0
routing
ip ospf hello-interval 1
ip ospf dead-interval 4
ip ospf network point-to-point
exit
interface te1/0/22
switchport mode trunk
exit
interface loopback 0
ip address 172.21.254.1 255.255.255.255
exit
exit
```

3 Configure R2.

```
terminal length 0
config
line console
serial timeout 0
exit
ip routing
```

```
router ospf
router-id 2.2.2.2
network 172.21.0.0 0.0.255.255 area 1
timers spf 3 5
exit
vlan 102,104
exit
interface vlan 102
ip address 172.21.2.2 255.255.255.0
routing
ip ospf hello-interval 1
ip ospf dead-interval 4
ip ospf network point-to-point
exit
interface te1/0/21
switchport mode trunk
exit
interface vlan 104
ip address 172.21.3.2 255.255.255.0
routing
ip ospf hello-interval 1
ip ospf dead-interval 4
ip ospf network point-to-point
exit
interface te1/0/22
switchport mode trunk
exit
interface loopback 0
ip address 172.21.254.2 255.255.255.255
exit
exit
```

4 R3 config:

```
terminal length 0
config
line console
serial timeout 0
exit
ip routing
router ospf
router-id 3.3.3.3
network 172.21.0.0 0.0.255.255 area 0
timers spf 3 5
exit
vlan 103
exit
```

```
interface vlan 103
ip address 172.21.1.1 255.255.255.0
routing
ip ospf hello-interval 1
ip ospf dead-interval 4
ip ospf network point-to-point
exit
interface te1/0/21
switchport mode trunk
exit
interface loopback 0
ip address 172.21.254.2 255.255.255.255
exit
exit
```

Discussion

With no area range cost specified, the range uses auto cost:

```
(ABR-R0) #show ip ospf range 1

          Prefix      Subnet Mask   Type     Action       Cost   Active
      172.21.0.0      255.255.0.0      S  Advertise       Auto        Y

(ABR-R0) #show ip ospf database summary

                Network Summary States (Area 0.0.0.0)
LS Age: 644
LS options: (E-Bit)
LS Type: Network Summary LSA
LS Id: 172.21.0.0 (network prefix)
Advertising Router: 10.10.10.10
LS Seq Number: 0x80000002
Checksum: 0x8ee1
Length: 28
 Network Mask: 255.255.0.0
 Metric: 2
```

Min—The cost can be set to 0, the minimum value. OSPF re-advertises the summary LSA with a metric of 0:

```
(ABR-R0) (config-router)#area 1 range 172.21.0.0 255.255.0.0 summarylink
advertise cost ?

<0-16777215>             Set area range cost

(ABR-R0) (config-router)#area 1 range 172.21.0.0 255.255.0.0 summarylink
advertise cost 0
```

```
(ABR-R0) #show ip ospf range 1

          Prefix      Subnet Mask   Type     Action       Cost   Active
      172.21.0.0      255.255.0.0      S  Advertise          0        Y

(ABR-R0) #show ip ospf 0 database summary

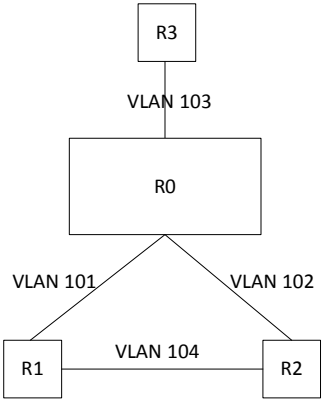
                Network Summary States (Area 0.0.0.0)
LS Age: 49
LS options: (E-Bit)
LS Type: Network Summary LSA
LS Id: 172.21.0.0 (network prefix)
Advertising Router: 10.10.10.10
LS Seq Number: 0x80000003
Checksum: 0x78f8
Length: 28
Network Mask: 255.255.0.0
Metric: 0
```

The cost can be set to the maximum value, 16,777,215, which is LSInfinity. Since OSPF cannot send a type 3 summary LSA with this metric (according to RFC 2328), the summary LSA is flushed. The individual routes are not re-advertised.

Configuring Flood Blocking

Figure 35-40 shows an example topology for flood blocking. The configuration follows.

Figure 35-40. Flood Blocking Topology



1 Configure R0:

```
terminal length 0
config
hostname R0
line console
exec-timeout 0
exit
vlan 101-103
exit
ip routing
router ospf
router-id 10.10.10.10
network 172.20.0.0 0.0.255.255 area 0
network 172.21.0.0 0.0.255.255 area 0
timers spf 3 5
exit
interface vlan 101
ip address 172.21.1.10 255.255.255.0
ip ospf hello-interval 1
ip ospf dead-interval 4
ip ospf network point-to-point
exit
interface te1/0/21
switchport mode trunk
description "R1"
exit
interface vlan 102
```

```
ip address 172.21.2.10 255.255.255.0
ip ospf hello-interval 1
ip ospf dead-interval 4
ip ospf network point-to-point
exit
interface te1/0/22
description "R2"
switchport mode trunk
exit
interface vlan 103
ip address 172.20.1.10 255.255.255.0
ip ospf hello-interval 1
ip ospf dead-interval 4
ip ospf network point-to-point
exit
interface te1/0/23
switchport mode trunk
description "R3"
exit
exit
```

2 Configure R1:

```
terminal length 0
config
hostname R1
line console
exec-timeout 0
exit
vlan 101,104
exit
ip routing
router ospf
router-id 1.1.1.1
network 172.21.0.0 0.0.255.255 area 0
timers spf 3 5
exit
interface vlan 101
ip address 172.21.1.1 255.255.255.0
routing
ip ospf hello-interval 1
ip ospf dead-interval 4
ip ospf network point-to-point
exit
interface te1/0/21
switchport mode trunk
```

```
exit
interface vlan 104
ip address 172.21.3.1 255.255.255.0
routing
ip ospf hello-interval 1
ip ospf dead-interval 4
ip ospf network point-to-point
exit
interface te1/0/22
switchport mode trunk
exit
interface loopback 0
ip address 172.21.254.1 255.255.255.255
exit
exit
```

3 Configure R2:

```
terminal length 0
config
line console
serial timeout 0
exit
ip routing
router ospf
router-id 2.2.2.2
network 172.21.0.0 0.0.255.255 area 0
timers spf 3 5
exit
vlan 102,104
exit
interface vlan 102
ip address 172.21.2.2 255.255.255.0
routing
ip ospf hello-interval 1
ip ospf dead-interval 4
ip ospf network point-to-point
exit
interface te1/0/21
switchport mode trunk
exit
interface vlan 104
ip address 172.21.3.2 255.255.255.0
routing
ip ospf hello-interval 1
ip ospf dead-interval 4
```

```
ip ospf network point-to-point
exit
interface te1/0/22
switchport mode trunk
exit
interface loopback 0
ip address 172.21.254.2 255.255.255.255
exit
exit
```

4 Configure R3:

```
terminal length 0
config
line console
serial timeout 0
exit
ip routing
router ospf
router-id 3.3.3.3
network 172.21.0.0 0.0.255.255 area 0
timers spf 3 5
exit
vlan 103
exit
interface vlan 103
ip address 172.21.1.1 255.255.255.0
routing
ip ospf hello-interval 1
ip ospf dead-interval 4
ip ospf network point-to-point
exit
interface te1/0/21
switchport mode trunk
exit
interface loopback 0
ip address 172.21.254.2 255.255.255.255
exit
exit
```

Discussion

With flood blocking disabled on all interfaces, sending a T3 summary LSA from R3 to R0 will cause R0 to forward the LSA on its interface to R1. Enabling flood blocking on R0's interface to R1 will inhibit this behavior.

```
(R0)(config-if-vlan101)ip ospf database-filter all out
```

A trace on the R3-R0 link shows that the LSA is actually flooded from R1 to R0, since R1 received the LSA via R2. Even though R1 does not receive this LSA directly from R0, it still correctly computes the route through the R0:

```
(R1) #show ip route

Route Codes: R - RIP Derived, O - OSPF Derived, C - Connected, S - Static
       B - BGP Derived, IA - OSPF Inter Area
       E1 - OSPF External Type 1, E2 - OSPF External Type 2
       N1 - OSPF NSSA External Type 1, N2 - OSPF NSSA External Type 2

O IA   100.0.0.0/24 [110/2] via 172.21.1.10,   00h:01m:35s,  0/25
```

OSPF also blocks external LSAs on the blocked interface. Stopping and restarting R3's OSPF protocol causes R3 to re-originate its router LSA. R0 does not send R3's router LSA on the blocked interface.

With flood blocking enabled on the R0 interface, if the link from R0 to R1 bounces, R0 Database Description packets do not include any LSAs. However, database synchronization still occurs (through R2) and R1 computes the correct routes after the link is restored.

36

Configuring RIP

This chapter describes how to configure Routing Information Protocol (RIP) on the switch. RIP is a dynamic routing protocol for IPv4 networks.

The topics covered in this chapter include:

- RIP Overview
- Default RIP Values
- Configuring RIP Features (Web)
- Configuring RIP Features (CLI)
- RIP Configuration Example

RIP Overview

RIP is an Interior Gateway Protocol (IGP) that performs dynamic routing within a network. PowerConnect 8000-series and 8100-series switches support two dynamic routing protocols: OSPF and Routing Information Protocol (RIP).

Unlike OSPF, RIP is a distance-vector protocol and uses UDP broadcasts to maintain topology information and hop counts to determine the best route to transmit IP traffic. RIP is best suited for small, homogenous networks.

How Does RIP Determine Route Information?

The routing information is propagated in RIP update packets that are sent out both periodically and in the event of a network topology change. On receipt of a RIP update, depending on whether the specified route exists or does not exist in the route table, the router may modify, delete or add the route to its route table.

RIP uses hop count, which is the number of routers an IP packet must pass through, to calculate the best route for a packet. A route with a low hop count is preferred over a route with a higher hop count. A directly-connected route

has a hop-count of 0. With RIP, the maximum number of hops from source to destination is 15. Packets with a hop count greater than 15 are dropped because the destination network is considered unreachable.

What Is Split Horizon?

RIP uses a technique called split horizon to avoid problems caused by including routes in updates sent to the router from which the route was originally learned. With simple split horizon, a route is not included in updates sent on the interface on which it was learned. In split horizon with poison reverse, a route is included in updates sent on the interface where it was learned, but the metric is set to infinity.

What RIP Versions Are Supported?

There are two versions of RIP:

- RIP-1 defined in RFC 1058
  - Routes are specified by IP destination network and hop count
  - The routing table is broadcast to all stations on the attached network
- RIP-2 defined in RFC 1723
  - Route specification is extended to include subnet mask and gateway
  - The routing table is sent to a multicast address, reducing network traffic
  - An authentication method is used for security

The PowerConnect 8000-series and 8100-series switches support both versions of RIP. You may configure a given port:

- To receive packets in either or both formats
- To transmit packets formatted for RIP-1 or RIP-2 or to send RIP-2 packets to the RIP-1 broadcast address
- To prevent any RIP packets from being received
- To prevent any RIP packets from being transmitted

Default RIP Values

RIP is globally enabled by default. To make it operational on the router, you configure and enable RIP for particular VLAN routing interfaces.

Table 36-1 shows the global default values for RIP.

Table 36-1. RIP Global Defaults

Parameter	Default Value
Admin Mode	Enabled
Split Horizon Mode	Simple
Auto Summary Mode	Disabled
Host Routes Accept Mode	Enabled
Default Information Originate	Disabled
Default Metric	None configured
Route Redistribution	Disabled for all sources.

Table 36-2 shows the per-interface default values for RIP.

Table 36-2. RIP Per-Interface Defaults

Parameter	Default Value
Admin Mode	Disabled
Send Version	RIPv2
Receive Version	Both
Authentication Type	None

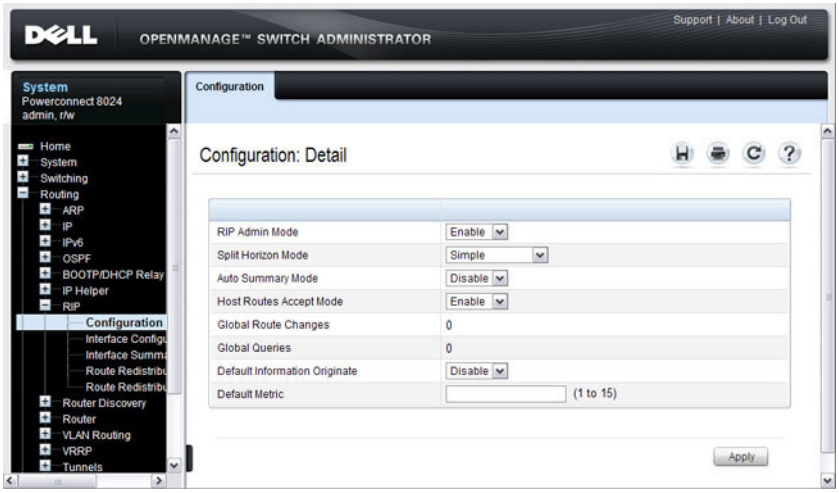
Configuring RIP Features (Web)

This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring RIP features on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

RIP Configuration

Use the **Configuration** page to enable and configure or disable RIP in Global mode. To display the page, click **Routing** → **RIP** → **Configuration** in the navigation panel.

Figure 36-1. RIP Configuration

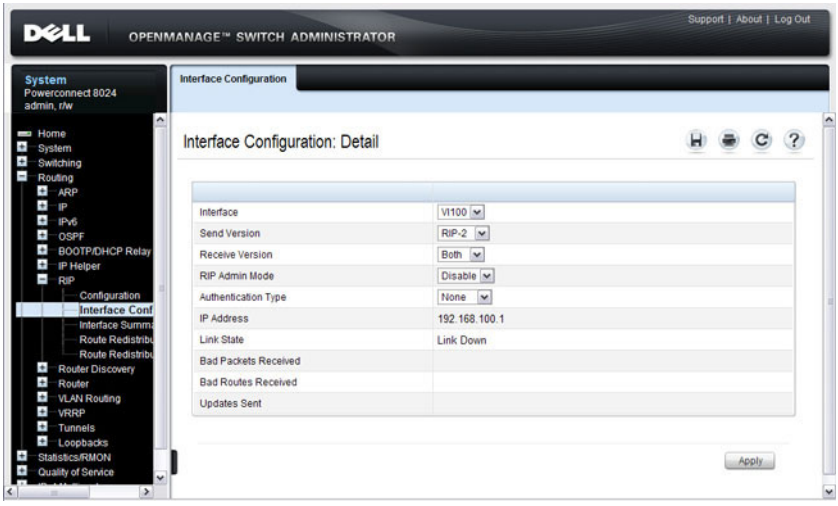


RIP Interface Configuration

Use the **Interface Configuration** page to enable and configure or to disable RIP on a specific interface.

To display the page, click **Routing** → **RIP** → **Interface Configuration** in the navigation panel.

Figure 36-2. RIP Interface Configuration

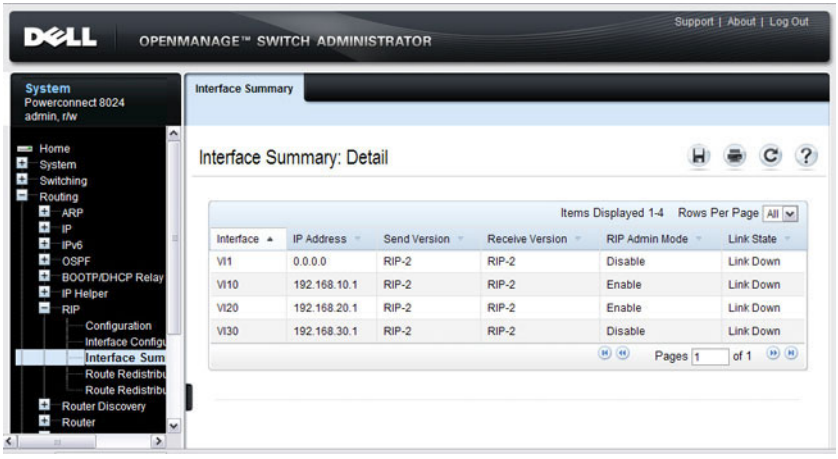


RIP Interface Summary

Use the **Interface Summary** page to display RIP configuration status on an interface.

To display the page, click **Routing** → **RIP** → **Interface Summary** in the navigation panel.

Figure 36-3. RIP Interface Summary

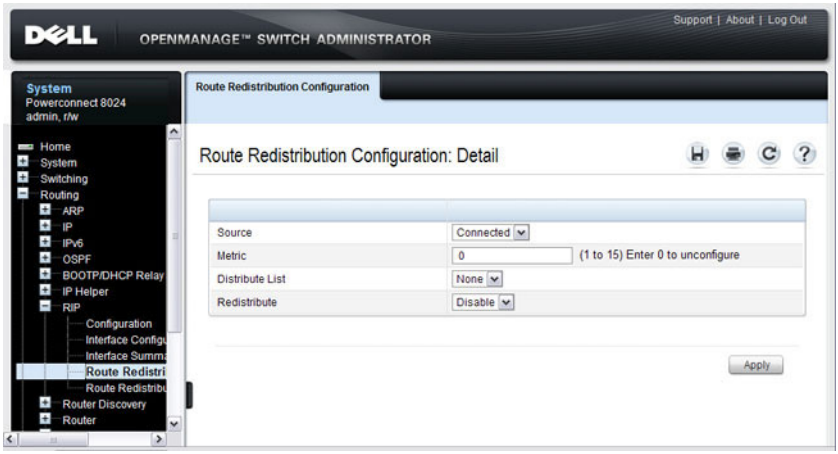


RIP Route Redistribution Configuration

Use the **Route Redistribution Configuration** page to configure the RIP Route Redistribution parameters. The allowable values for each fields are displayed next to the field. If any invalid values are entered, an alert message is displayed with the list of all the valid values.

To display the page, click **Routing → RIP → Route Redistribution Configuration** in the navigation panel.

Figure 36-4. RIP Route Redistribution Configuration



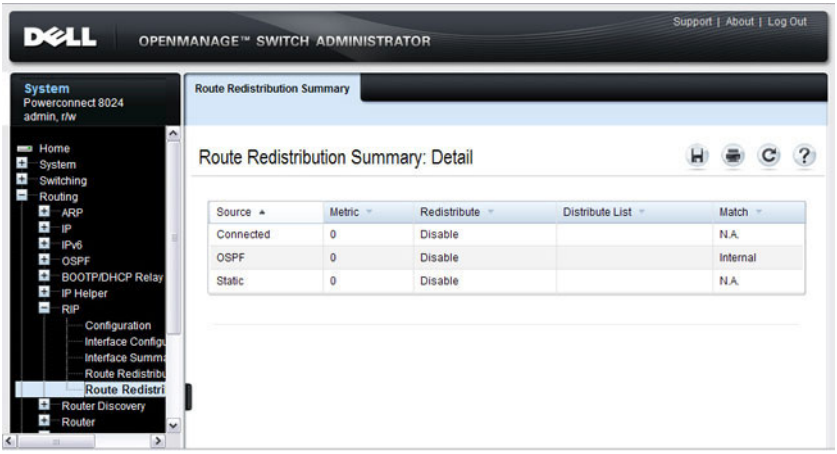
NOTE: Static reject routes are not redistributed by RIP. For a static reject route, the next hop interface value is Null0. Packets to the network address specified in static reject routes are intentionally dropped.

RIP Route Redistribution Summary

Use the **Route Redistribution Summary** page to display Route Redistribution configurations.

To display the page, click **Routing → RIP → Route Redistribution Summary** in the navigation panel.

Figure 36-5. RIP Route Redistribution Summary



Configuring RIP Features (CLI)

This section provides information about the commands you use to configure RIP settings on the switch. For more information about the commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Configuring Global RIP Settings

Beginning in Privileged EXEC mode, use the following commands to configure various global RIP settings for the switch.

NOTE: RIP is enabled by default. The Global RIP Settings are optional.

Command	Purpose
configure	Enter global configuration mode.
router rip	Enter OSPF configuration mode.
split-horizon {none | simple | poison}	Set the RIP split horizon mode. • **none** — RIP does not use split horizon to avoid routing loops. • **simple** — RIP uses split horizon to avoid routing loops. • **poison** — RIP uses split horizon with poison reverse (increases routing packet update size).
auto-summary	Enable the RIP auto-summarization mode.
no hostroutesaccept	Prevent the switch from accepting host routes.
default-information originate	Control the advertisement of default routes.
default-metric *metric-value*	Set a default for the metric of distributed routes. The *metric-value* variable is the metric (or preference) value of the default route. (Range: 1–15)
enable	Reset the default administrative mode of RIP in the router (active)
CTRL + Z	Exit to Privileged EXEC mode.
show ip rip	View various RIP settings for the switch.

Configuring RIP Interface Settings

Beginning in Privileged EXEC mode, use the following commands to configure per-interface RIP settings.

Command	Purpose
configure	Enter global configuration mode.
interface vlan *vlan-id*	Enter Interface Configuration mode for the specified VLAN.
ip rip	Enable RIP on the interface.
ip rip send version {rip1 rip1c | rip2 |none}	Configure the interface to allow RIP control packets of the specified version(s) to be sent.
ip rip receive version {rip1 | rip2 | both | none}	Configure the interface to allow RIP control packets of the specified version(s) to be received.
ip rip authentication {none | {simple *key***} | {encrypt** *key key-id***}**	set the RIP Version 2 Authentication Type and Key for the interface. • *key* — Authentication key for the specified interface. (Range: 16 bytes or less) • **encrypt** — Specifies the Ethernet unit/port of the interface to view information. • *key-id* — Authentication key identifier for authentication type encrypt. (Range: 0-255)
exit	Exit to Global Configuration Mode
exit	Exit to Privileged Exec mode.
show ip rip interface vlan *vlan-id*	View RIP configuration information for the specified routing interface.
show ip rip interface brief	View summary information about the RIP configuration on all interfaces.

Configuring Route Redistribution Settings

Beginning in Privileged EXEC mode, use the following commands to configure an OSPF area range and to configure route redistribution settings.

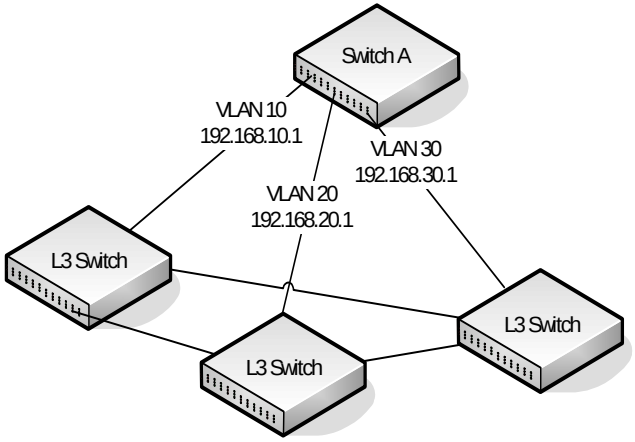
Command	Purpose
configure	Enter global configuration mode.
router rip	Enter RIP configuration mode.
distribute-list *accesslistname* **out** {**ospf** | **static** | **connected**}	Specify the access list to filter routes received from the source protocol. The ACL must already exist on the switch. For information about the commands you use to configure ACLs, see "Configuring ACLs (CLI)" on page 521. • *accesslistname* — The name used to identify an existing ACL. • **ospf** — Apply the specified access list when OSPF is the source protocol. • **static** — Apply the specified access list when packets come through the static route. • **connected** — Apply the specified access list when packets come from a directly connected route.
redistribute {**static** | **connected**} [**metric** *integer*]	Configure RIP to allow redistribution of routes from the specified source protocol/routers. • **static** — Specifies that the source is a static route. • **connected** — Specifies that the source is a directly connected route. • *metric* — Specifies the metric to use when redistributing the route. Range: 1-15.

Command	Purpose
redistribute ospf [**metric** *metric*] [**match** [**internal**] [**external 1**] [**external 2**] [**nssa-external 1**] [**nssa-external 2**]]	Configure RIP to allow redistribution of routes from the OSPF. • **ospf**— Specifies OSPF as the source protocol. • *metric* — Specifies the metric to use when redistributing the route. Range: 1-15. • **internal** — Adds internal matches to any match types presently being redistributed. • **external 1** — Adds routes imported into OSPF as Type-1 external routes into any match types presently being redistributed. • **external 2** — Adds routes imported into OSPF as Type-2 external routes into any match types presently being redistributed. • **nssa-external 1** — Adds routes imported into OSPF as NSSA Type-1 external routes into any match types presently being redistributed. • **nssa-external 2** — Adds routes imported into OSPF as NSSA Type-2 external routes into any match types presently being redistributed.
distance rip *integer*	Set the route preference value of RIP in the router. Lower route preference values are preferred when determining the best route.
exit	Exit to Global Config mode.
exit	Exit to Privileged Exec mode.
show ip rip	View information about the RIP route distribution configuration.

RIP Configuration Example

This example includes four PowerConnect switches that use RIP to determine network topology and route information. The commands in this example configure Switch A shown in Figure 36-6.

Figure 36-6. RIP Network Diagram



To configure the switch:

1. Enable routing on the switch

   ```
   console#config
   console(config)#ip routing
   ```

2. Create VLANs 10, 20, and 30.

   ```
   console(config)#vlan 10,20,30
   ```

3. Assign an IP address and enable RIP on each interface. Additionally, the commands specify that each interface can receive both RIP-1 and RIP-2 frames but send only RIP-2 formatted frames.

   ```
   console(config)#interface vlan 10
   console(config-if-vlan10)#ip address 192.168.10.1
   255.255.255.0
   console(config-if-vlan10)#ip rip
   ```

```
console(config-if-vlan10)#ip rip receive version
both
console(config-if-vlan10)#ip rip send version rip2
console(config-if-vlan10)#exit

console(config)#interface vlan 20
console(config-if-vlan20)#ip address 192.168.20.1
255.255.255.0
console(config-if-vlan20)#ip rip
console(config-if-vlan20)#ip rip receive version
both
console(config-if-vlan20)#ip rip send version rip2
console(config-if-vlan20)#exit

console(config)#interface vlan 30
console(config-if-vlan30)#ip address 192.168.30.1
255.255.255.0
console(config-if-vlan30)#ip rip
console(config-if-vlan30)#ip rip receive version
both
console(config-if-vlan30)#ip rip send version rip2
console(config-if-vlan30)#exit
```

4 Enable auto summarization of subprefixes when crossing classful boundaries.

```
console(config)#router rip
console(config-router)#auto-summary
console(config-router)#exit
console(config)#exit
```

5 Verify the configuration

```
console#show ip rip

RIP Admin Mode..................... Enable
Split Horizon Mode................. Simple
Auto Summary Mode.................. Enable
Host Routes Accept Mode............ Enable
Global route changes............... 0
Global queries..................... 0
```

```
Default Metric..................... Not configured
Default Route Advertise........... 0
```

```
console#show ip rip interface brief

Interface  IP Address     Send       Receive  RIP        Link
                         Version    Version  Mode       State
---------- ------------  ---------  -------  ---------  ------
Vl1        0.0.0.0       RIP-2      RIP-2    Disable    Down
Vl10       192.168.10.1  RIP-2      Both     Enable     Down
Vl20       192.168.10.1  RIP-2      Both     Enable     Down
Vl30       192.168.10.1  RIP-2      Both     Disable    Down
```

37

Configuring VRRP

This chapter describes how to configure Virtual Routing Redundancy Protocol (VRRP) on the switch. VRRP can help create redundancy on networks in which end-stations are statically configured with the default gateway IP address.

The topics covered in this chapter include:

- VRRP Overview
- Default VRRP Values
- Configuring VRRP Features (Web)
- Configuring VRRP Features (CLI)
- VRRP Configuration Example

VRRP Overview

The Virtual Router Redundancy (VRRP) protocol is designed to handle default router (L3 switch) failures by providing a scheme to dynamically elect a backup router. VRRP can help minimize black hole periods due to the failure of the default gateway router during which all traffic directed towards it is lost until the failure is detected.

How Does VRRP Work?

VRRP eliminates the single point of failure associated with static default routes by enabling a backup router to take over from a master router without affecting the end stations using the route. The end stations will use a virtual IP address that will be recognized by the backup router if the master router fails. Participating routers use an election protocol to determine which router is the master router at any given time. A given port may appear as more than one virtual router to the network, also, more than one port on a switch may be configured as a virtual router. Either a physical port or a routed VLAN may participate.

With VRRP, a virtual router is associated with one or more IP addresses that serve as default gateways. In the event that the VRRP router controlling these IP addresses (formally known as the master) fails, the group of IP addresses and the default forwarding role is taken over by a Backup VRRP router.

NOTE: It is not possible to ping the VRRP IP address from the VRRP master. Use the `show vrrp` command to display the status of the VRRP router

What Is the VRRP Router Priority?

The VRRP router priority is a value from 1–255 that determines which router is the master. The greater the number, the higher the priority. If the virtual IP address is the IP address of a VLAN routing interface on one of the routers in the VRRP group, the router with IP address that is the same as the virtual IP address is the interface owner and automatically has a priority of 255. By default, this router is the VRRP master in the group.

If no router in the group owns the VRRP virtual IP address, the router with the highest configured priority is the VRRP master. If multiple routers have the same priority, the router with the highest IP address becomes the VRRP master.

If the VRRP master fails, other members of the VRRP group will elect a master based on the configured router priority values. For example, router A is the interface owner and master, and it has a priority of 255. Router B is configured with a priority of 200, and Router C is configured with a priority of 190. If Router A fails, Router B assumes the role of VRRP master because it has a higher priority.

What Is VRRP Preemption?

If preempt mode is enabled and a router with a higher priority joins the VRRP group, it takes over the VRRP master role if the current VRRP master is not the owner of the virtual IP address. The preemption delay controls how long to wait to determine whether a higher priority Backup router preempts a lower priority master. In certain cases, for example, during periods of network congestion, a backup router might fail to receive advertisements from the master. This could cause members in the VRRP group to change their states frequently, i.e. flap. The problem can be resolved by setting the VRRP preemption delay timer to a non-zero value.

What Is VRRP Accept Mode?

The accept mode allows the switch to respond to pings (ICMP Echo Requests) sent to the VRRP virtual IP address. The VRRP specification (RFC 3768) indicates that a router may accept IP packets sent to the virtual router IP address only if the router is the address owner. In practice, this restriction makes it more difficult to troubleshoot network connectivity problems. When a host cannot communicate, it is common to ping the host's default gateway to determine whether the problem is in the first hop of the path to the destination. When the default gateway is a virtual router that does not respond to pings, this troubleshooting technique is unavailable. In the PowerConnect switch VRRP feature, you can enable Accept Mode to allow the system to respond to pings that are sent to the virtual IP address.

This capability adds support for responding to pings, but does not allow the VRRP master to accept other types of packets. The VRRP master responds to both fragmented and un-fragmented ICMP Echo Request packets. The VRRP master responds to Echo Requests sent to the virtual router's primary address or any of its secondary addresses.

Members of the virtual router who are in backup state discard ping packets destined to VRRP addresses, just as they discard any Ethernet frame sent to a VRRP MAC address.

When the VRRP master responds with an Echo Reply, the source IPv4 address is the VRRP address and source MAC address is the virtual router's MAC address.

What Are VRRP Route and Interface Tracking?

The VRRP Route/Interface Tracking feature extends VRRP capability to allow tracking of specific routes and interface IP states within the router that can alter the priority level of a virtual router for a VRRP group.

VRRP interface tracking monitors a specific interface IP state within the router. Depending on the state of the tracked interface, the feature can alter the VRRP priority level of a virtual router for a VRRP group.

NOTE: An exception to the priority level change is that if the VRRP group is the IP address owner, its priority is fixed at 255 and cannot be reduced through the tracking process.

With standard VRRP, the backup router takes over only if the router goes down. With VRRP interface tracking, if a tracked interface goes down on the VRRP master, the priority decrement value is subtracted from the router priority. If the master router priority becomes less than the priority on the backup router, the backup router takes over. If the tracked interface becomes up, the value of the priority decrement is added to the current router priority. If the resulting priority is more than the backup router priority, the original VRRP master resumes control.

VRRP route tracking monitors the reachability of an IP route. A tracked route is considered up when a routing table entry exists for the route and the route is accessible. When the tracked route is removed from the routing table, the priority of the VRRP router will be reduced by the priority decrement value. When the tracked route is added to the routing table, the priority will be incremented by the same.

Default VRRP Values

Table 37-1 shows the global default values for VRRP.

Table 37-1. VRRP Defaults

Parameter	Default Value
Admin Mode	Disabled
Virtual Router ID (VRID)	None
Preempt Mode	Enabled
Preempt Delay	0 Seconds
Learn Advertisement Timer Interval	Enabled
Accept Mode	Disabled
Configured Priority	100
Advertisement Interval	1
Authentication	None
Route Tracking	No routes tracked
Interface Tracking	No interfaces tracked

Configuring VRRP Features (Web)

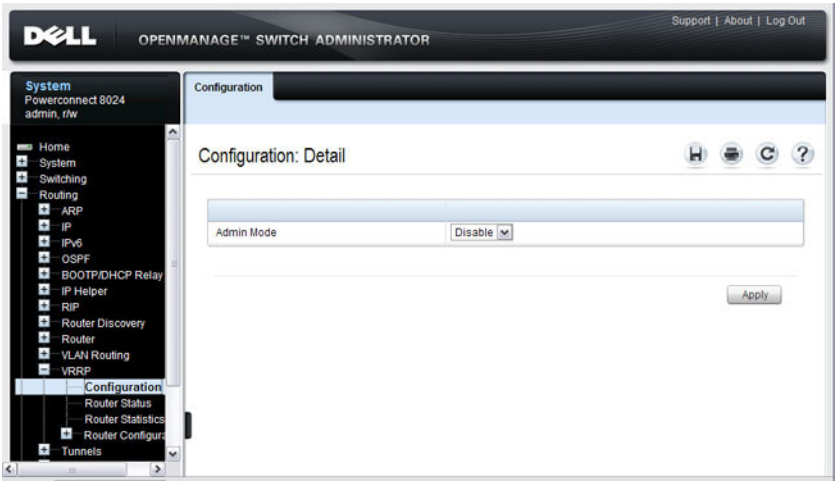
This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring VRRP features on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

VRRP Configuration

Use the **Configuration** page to enable or disable the administrative status of a virtual router.

To display the page, click **Routing** → **VRRP** → **Configuration** in the navigation panel.

Figure 37-1. VRRP Configuration

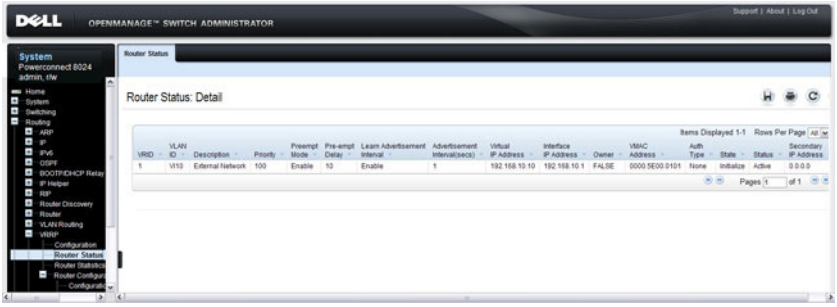


VRRP Virtual Router Status

Use the **Router Status** page to display virtual router status.

To display the page, click **Routing → VRRP → Router Status** in the navigation panel.

Figure 37-2. Virtual Router Status

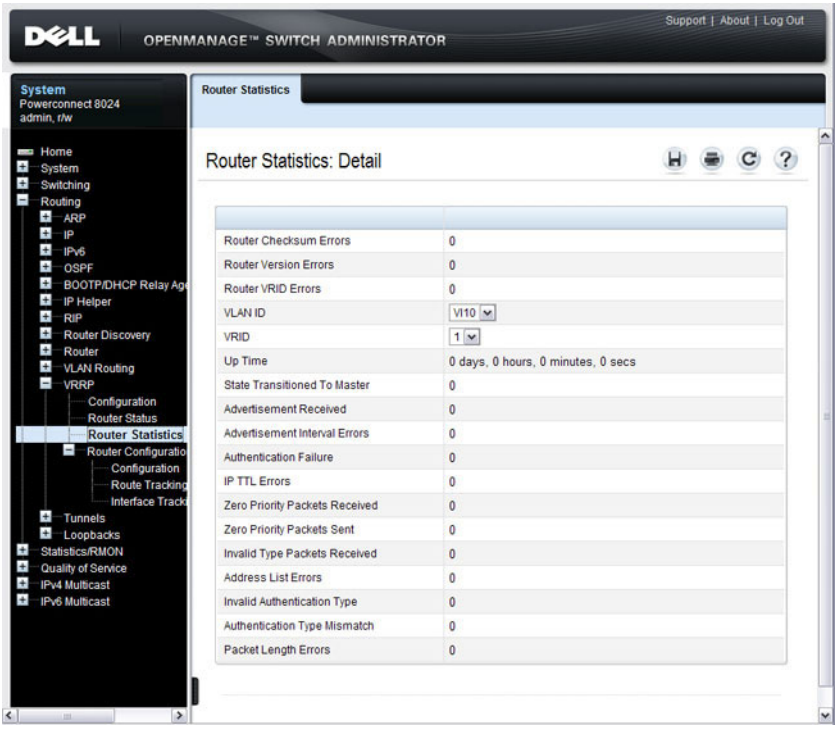


VRRP Virtual Router Statistics

Use the **Router Statistics** page to display statistics for a specified virtual router.

To display the page, click **Routing** → **VRRP** → **Router Statistics** in the navigation panel.

Figure 37-3. Virtual Router Statistics

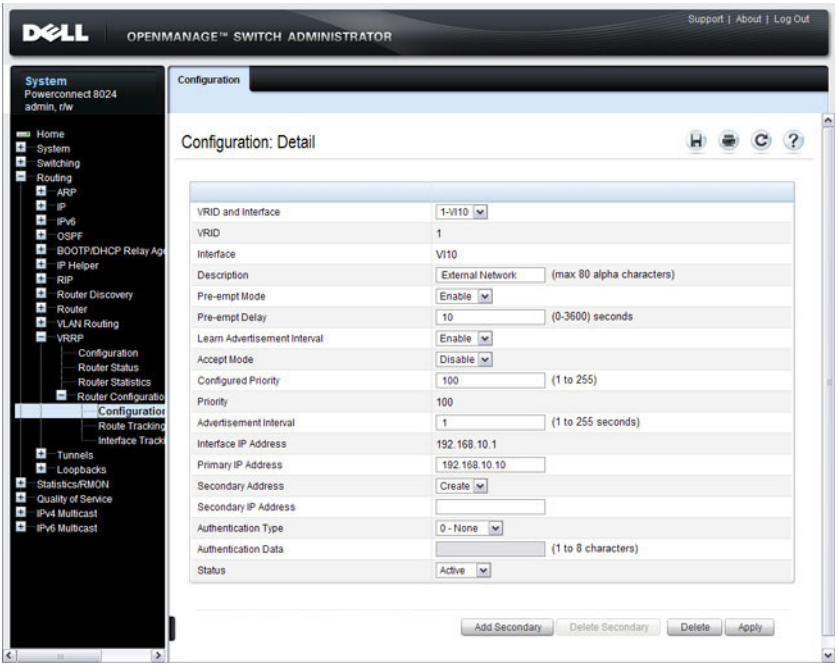


VRRP Router Configuration

Use the **Configuration** page to configure a virtual router.

To display the page, click **Routing** → **VRRP** → **Router Configuration** → **Configuration** in the navigation panel.

Figure 37-4. VRRP Router Configuration

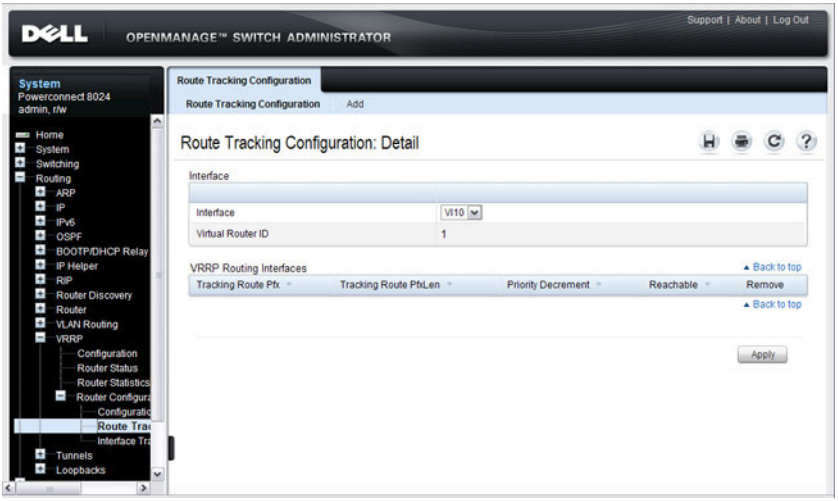


VRRP Route Tracking Configuration

Use the **Route Tracking Configuration** page to view routes that are tracked by VRRP and to add new tracked routes.

To display the page, click **Routing** → **VRRP** → **Router Configuration** → **Route Tracking Configuration** in the navigation panel.

Figure 37-5. VRRP Route Tracking Configuration



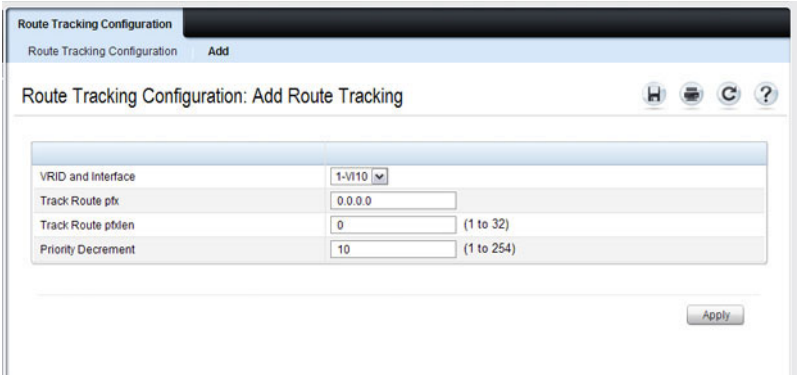
Configuring VRRP Route Tracking

To configure VRRP route tracking:

1 From the **Route Tracking Configuration** page, click **Add**.

The **Add Route Tracking** page displays.

Figure 37-6. Add Route Tracking



2 Select the virtual router ID and VLAN routing interface that will track the route.

3 Specify the destination network address (track route prefix) for the route to track. Use dotted decimal format, for example 192.168.10.0.

4 Specify the prefix length for the tracked route.

5 Specify a value for the **Priority Decrement** to define the amount that the router priority will be decreased when a tracked route becomes unreachable.

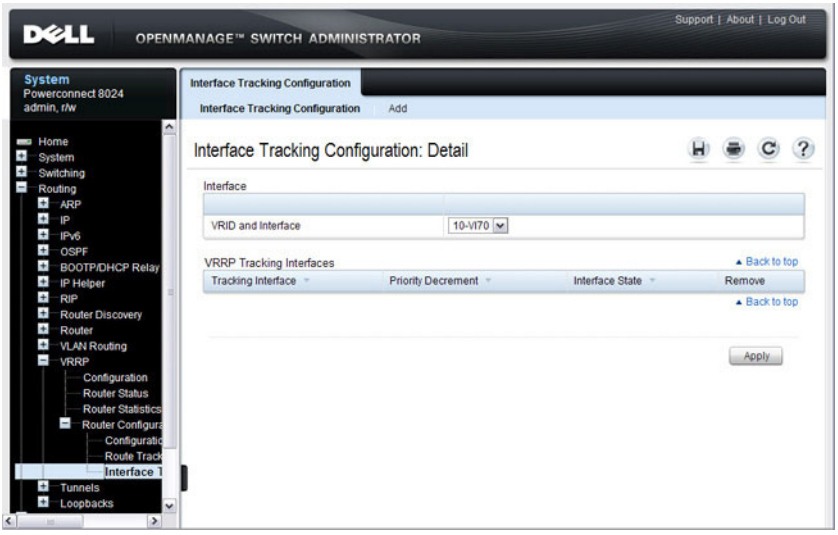
6. Click **Apply** to update the switch.

VRRP Interface Tracking Configuration

Use the **Interface Tracking Configuration** page to view interfaces that are tracked by VRRP and to add new tracked interfaces.

To display the page, click **Routing** → **VRRP** → **Router Configuration** → **Interface Tracking Configuration** in the navigation panel.

Figure 37-7. VRRP Interface Tracking Configuration



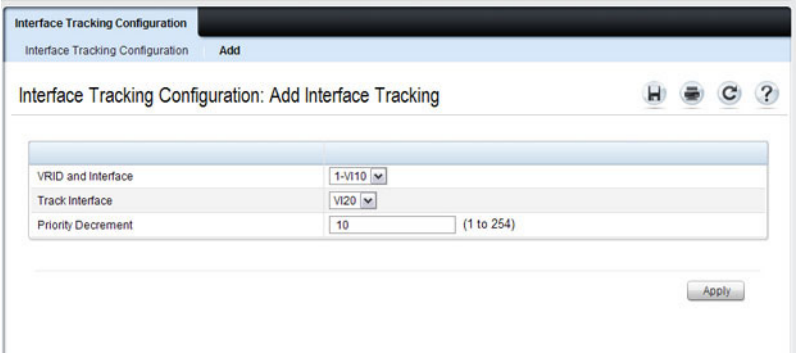
Configuring VRRP Interface Tracking

To configure VRRP interface tracking:

1 From the **Interface Tracking Configuration** page, click **Add**.

The **Add Interface Tracking** page displays.

Figure 37-8. VRRP Interface Tracking Configuration



2 Select the virtual router ID and VLAN routing interface that will track the interface.

3 Specify the interface to track.

4 Specify a value for the **Priority Decrement** to define the amount that the router priority will be decreased when a tracked interface goes down.

5. Click **Apply** to update the switch.

Configuring VRRP Features (CLI)

This section provides information about the commands you use to configure VRRP settings on the switch. For more information about the commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Configuring VRRP Settings

Beginning in Privileged EXEC mode, use the following commands to configure switch and interface VRRP settings. This set of commands also describes how to configure VRRP interface and route tracking.

Command	Purpose
configure	Enter global configuration mode.
ip vrrp	Enable the administrative mode of VRRP for the router (L3 switch).
interface vlan *vlan-id*	Enter Interface Configuration mode for the specified VLAN.
vrrp *vr-id*	Allow the interface to create in the VRRP group specified by the *vr-id* parameter, which is a number from 1–255.
vrrp *vr-id* **description**	(Optional) Create a text description that identifies the VRRP group.
vrrp *vr-id* **preempt** [**delay** *seconds*]	Enable the preemption mode value for the virtual router configured on a specified interface. You can optionally configure a preempt delay, which is the number of seconds the VRRP router waits before the VRRP router sends an advertisement to claim master ownership.
vrrp *vr-id* **accept-mode**	Allow the VRRP master to accept ping packets sent to one of the virtual router's IP addresses.
vrrp *vr-id* **priority** *priority*	Set the priority value for the virtual router configured on the interface.
vrrp *vr-id* **ip** *ip-address* [**secondary**]	Set the virtual router IP address value for an interface.

Command	Purpose
vrrp *vr-id* **timers {learn | advertise** *seconds*}	Configure the VRRP timer settings. Use the keyword **learn** to enable VRRP to learn the advertisement timer interval of the master router. Use the keyword **advertise** to set the frequency, in seconds, that an interface on the specified virtual router sends a virtual router advertisement.
vrrp *vr-id* **authentication {none | simple** *key*}	Set the authorization details value for the virtual router configured on a specified interface. • *vr-id* — The virtual router identifier. (Range: 1-255) • **none** — Indicates authentication type is none. • **simple** — Authentication type is a simple text password. • *key* — The key for simple authentication. (Range: String values)
vrrp *vr-id* **mode**	Enable the virtual router configured on an interface, which starts the virtual router.
vrrp *vr-id* **track interface vlan** *vlan-id* [**decrement** *priority*]	Specify an interface the virtual router (*vr-id*) on the interface will track. If the interface goes down, the virtual router priority is decreased by the amount specified by the *priority* value.
vrrp *vr-id* **track ip route** *ip-address/prefix-length* [**decrement** *priority*]	Specify a route that the virtual router (*vr-id*) on the interface will track. If the route to the destination network specified by the *ip-address/prefix-length* variable is removed from the routing table, the virtual router priority is decreased by the amount specified by the *priority* value.
CTRL + Z	Exit to Privileged EXEC mode.
show vrrp [*vr-id*]	View settings for all VRRP groups or for the specified VRRP group for the switch.
show vrrp brief	View a summary of interfaces configured to participate in VRRP groups.
show vrrp interface {brief | vlan *vlan-id* [**stats**]}	View information about VRRP settings configured on all interfaces or on the specified interface. If you specify an interface, use the keyword **stats** to view VRRP statistics for the interface.

VRRP Configuration Example

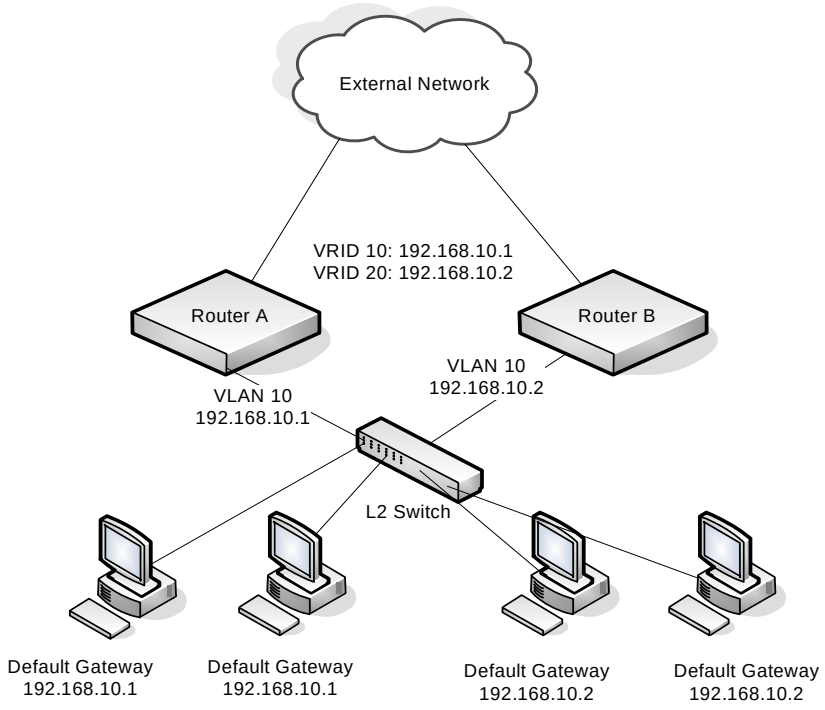
This section contains the following VRRP examples:

- VRRP with Load Sharing
- VRRP with Route and Interface Tracking

VRRP with Load Sharing

In Figure 37-9, two L3 PowerConnect switches are performing the routing for network clients. Router A is the default gateway for some clients, and Router B is the default gateway for other clients.

Figure 37-9. VRRP with Load Sharing Network Diagram



This example configures two VRRP groups on each router. Router A is the VRRP master for the VRRP group with VRID 10 and the backup for VRID 20. Router B is the VRRP master for VRID 20 and the backup for VRID 10. If Router A fails, Router B will become the master of VRID 10 and will use the virtual IP address 192.168.10.1. Traffic from the clients configured to use Router A as the default gateway will be handled by Router B.

To configure Router A:

1. Enable routing for the switch.

   ```
   console#config
   console(config)#ip routing
   ```

2. Create and configure the VLAN routing interface to use as the default gateway for network clients. This example assumes all other routing interfaces, such as the interface to the external network, have been configured.

   ```
   console(config)#interface vlan 10
   console(config-if-vlan10)#ip address 192.168.10.1
   255.255.255.0
   console(config-if-vlan10)#exit
   ```

3. Enable VRRP for the switch.

   ```
   console(config)#ip vrrp
   ```

4. Assign a virtual router ID to the VLAN routing interface for the first VRRP group.

   ```
   console(config)#interface vlan 10
   console(config-if-vlan10)#vrrp 10
   ```

5. Specify the IP address that the virtual router function will use. The router is the virtual IP address owner (the routing interface has the same IP address as the virtual IP address for the VRRP group), so the priority value is 255.

   ```
   console(config-if-vlan10)#vrrp 10 ip 192.168.10.1
   ```

6. Configure an optional description to help identify the VRRP group.

   ```
   console(config-if-vlan10)#vrrp 10 description
   master
   ```

7 Assign a virtual router ID to the VLAN routing interface for the second VRRP group.

```
console(config-if-vlan10)#vrrp 20
```

8 Specify the IP address that the virtual router function will use.

```
console(config-if-vlan10)#vrrp 20 ip 192.168.10.2
```

9 Configure an optional description to help identify the VRRP group.

```
console(config-if-vlan10)#vrrp 20 description
backup
```

10 Enable the VRRP groups on the interface.

```
console(config-if-vlan10)#ip vrrp 10 mode
console(config-if-vlan10)#ip vrrp 20 mode
console(config-if-vlan10)#exit
console(config)#exit
```

The only difference between the Router A and Router B configurations is the IP address assigned to VLAN 10. On Router B, the IP address of VLAN 10 is 192.168.10.2. Because this is also the virtual IP address of VRID 20, Router B is the interface owner and VRRP master of VRRP group 20.

To configure Router B:

1 Enable routing for the switch.

```
console#config
console(config)#ip routing
```

2 Create and configure the VLAN routing interface to use as the default gateway for network clients. This example assumes all other routing interfaces, such as the interface to the external network, have been configured.

```
console(config)#interface vlan 10
console(config-if-vlan10)#ip address 192.168.10.2
255.255.255.0
console(config-if-vlan10)#exit
```

3 Enable VRRP for the switch.

```
console(config)#ip vrrp
```

4 Assign a virtual router ID to the VLAN routing interface for the first VRRP group.

```
console(config)#interface vlan 10
console(config-if-vlan10)#vrrp 10
```

5 Specify the IP address that the virtual router function will use.

```
console(config-if-vlan10)#vrrp 10 ip 192.168.10.1
```

6 Configure an optional description to help identify the VRRP group.

```
console(config-if-vlan10)#vrrp 10 description
master
```

7 Assign a virtual router ID to the VLAN routing interface for the second VRRP group.

```
console(config-if-vlan10)#vrrp 20
```

8 Specify the IP address that the virtual router function will use.

The router is the virtual IP address owner of this address, so the priority value is 255 by default.

```
console(config-if-vlan10)#vrrp 20 ip 192.168.10.2
```

9 Configure an optional description to help identify the VRRP group.

```
console(config-if-vlan10)#vrrp 20 description
backup
```

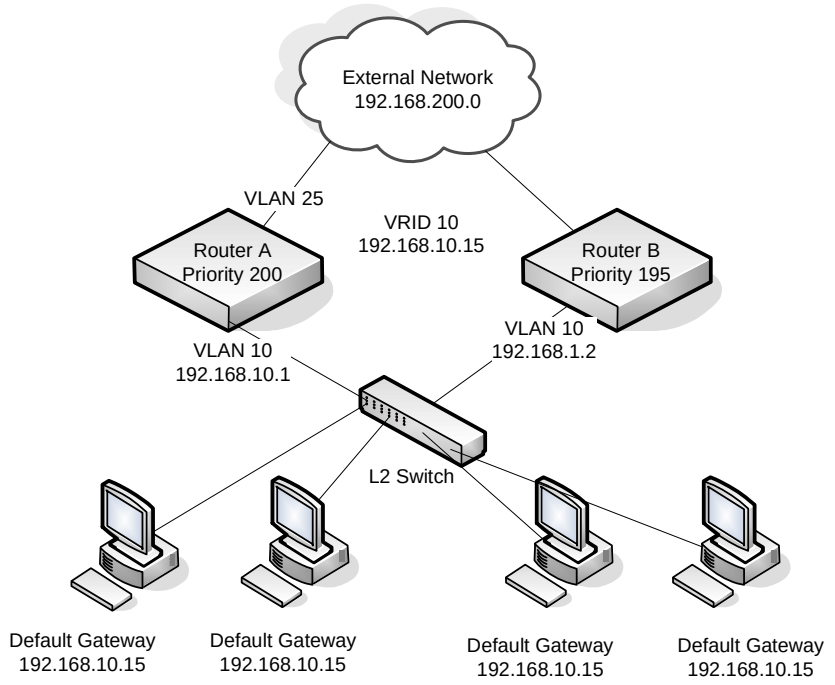
10 Enable the VRRP groups on the interface.

```
console(config-if-vlan10)#ip vrrp 10 mode
console(config-if-vlan10)#ip vrrp 20 mode
console(config-if-vlan10)#exit
console(config)#exit
```

VRRP with Route and Interface Tracking

In Figure 37-10, the VRRP priorities are configured so that Router A is the VRRP master, and Router B is the VRRP backup. Router A forwards IP traffic from clients to the external network through the VLAN 25 routing interface. The clients are configured to use the virtual IP address 192.168.10.15 as the default gateway.

Figure 37-10. VRRP with Tracking Network Diagram



Without VRRP interface or route tracking, if something happened to VLAN 25 or the route to the external network, as long as Router A remains up, it will continue to be the VRRP master even though traffic from the clients does not have a path to the external network. However, if the interface and/or route tracking features are configured, Router A can decrease its priority value when the problems occur so that Router B becomes the master.

To configure Router A:

1. Enable routing for the switch.

```
console#config
console(config)#ip routing
```

2. Create and configure the VLAN routing interface to use as the default gateway for network clients. This example assumes all other routing interfaces, such as the interface to the external network, have been configured.

```
console(config)#interface vlan 10
console(config-if-vlan10)#ip address 192.168.10.1
255.255.255.0
console(config-if-vlan10)#exit
```

3. Enable VRRP for the switch.

```
console(config)#ip vrrp
```

4. Assign a virtual router ID to the VLAN routing interface for the VRRP group.

```
console(config)#interface vlan 10
console(config-if-vlan10)#vrrp 10
```

5. Specify the IP address that the virtual router function will use.

```
console(config-if-vlan10)#vrrp 10 ip 192.168.10.15
```

6. Configure the router priority.

```
console(config-if-vlan10)#vrrp 10 priority 200
```

7. Enable preempt mode so that the router can regain its position as VRRP master if its priority is greater than the priority of the backup router.

```
console(config-if-vlan10)#vrrp 10 preempt
```

8. Enable the VRRP groups on the interface.

```
console(config-if-vlan10)#ip vrrp 10 mode
console(config-if-vlan10)#exit
```

9. Track the routing interface VLAN 25 on VRID 10 so that if it goes down, the priority of VRID 10 on Router A is decreased by 10, which is the default decrement priority value.

```
console(config-if-vlan10)#vrrp 10 track interface
vlan 25
```

10 Track the route to the 192.168.200.0 network. If it becomes unavailable, the priority of VRID 10 on Router A is decreased by 10, which is the default decrement priority value.

```
console(config-if-vlan10)#vrrp 10 track ip route
192.168.200.0/24
console(config-if-vlan10)#exit
```

Router B is the backup router for VRID 10. The configured priority is 195. If the VLAN 25 routing interface or route to the external network on Router A go down, the priority of Router A will become 190 (or 180, if both the interface and router are down). Because the configured priority of Router B is greater than the actual priority of Router A, Router B will become the master for VRID 10. When VLAN 25 and the route to the external network are back up, the priority of Router A returns to 200, and it resumes its role as VRRP master.

To configure Router B:

1 Enable routing for the switch.

```
console#config
console(config)#ip routing
```

2 Create and configure the VLAN routing interface to use as the default gateway for network clients. This example assumes all other routing interfaces, such as the interface to the external network, have been configured.

```
console(config)#interface vlan 10
console(config-if-vlan10)#ip address 192.168.10.2
255.255.255.0
console(config-if-vlan10)#exit
```

3 Enable VRRP for the switch.

```
console(config)#ip vrrp
```

4 Assign a virtual router ID to the VLAN routing interface for the VRRP group.

```
console(config)#interface vlan 10
console(config-if-vlan10)#vrrp 10
```

5 Specify the IP address that the virtual router function will use.

```
console(config-if-vlan10)#vrrp 10 ip 192.168.10.15
```

6 Configure the router priority.

```
console(config-if-vlan10)#vrrp 10 priority 195
```

7 Enable preempt mode so that the router can regain its position as VRRP master if its priority is greater than the priority of the backup router.

```
console(config-if-vlan10)#vrrp 10 preempt
```

8 Enable the VRRP groups on the interface.

```
console(config-if-vlan10)#ip vrrp 10 mode
console(config-if-vlan10)#exit
console(config)#exit
```

38

Configuring IPv6 Routing

This chapter describes how to configure general IPv6 routing information on the switch, including global routing settings and IPv6 static routes. The topics covered in this chapter include:

- IPv6 Routing Overview
- Default IPv6 Routing Values
- Configuring IPv6 Routing Features (Web)
- Configuring IPv6 Routing Features (CLI)

The PowerConnect 8000-series and 8100-series switches support additional features to help manage IPv6 networks, including OSPFv3, DHCPv6, and IPv6 multicast. For information about OSPFv3, see "Configuring OSPF and OSPFv3" on page 931. For information about DHCPv6, see "Configuring DHCPv6 Server and Relay Settings" on page 1083. For information about IPv6 multicast, see "Managing IPv4 and IPv6 Multicast" on page 1157.

For configuration examples that include IPv6 interface configuration, see "OSPF Configuration Examples" on page 996

IPv6 Routing Overview

IPv6 is the next generation of the Internet Protocol. With 128-bit addresses, versus 32-bit addresses for IPv4, IPv6 solves the address depletion issues seen with IPv4 and removes the requirement for Network Address Translation (NAT), which is used in IPv4 networks to reduce the number of globally unique IP addresses required for a given network.

On the PowerConnect 8000-series and 8100-series switches, IPv6 coexists with IPv4. As with IPv4, IPv6 routing can be enabled on loopback and VLAN interfaces. Each L3 routing interface can be used for IPv4, IPv6, or both. IP protocols running over L3 (for example, UDP and TCP) are common to both IPv4 and IPv6.

How Does IPv6 Compare with IPv4?

There are many conceptual similarities between IPv4 and IPv6 network operation. Addresses still have a network prefix portion (network) and a device interface specific portion (host). While the length of the network portion is still variable, most users have standardized on using a network prefix length of 64 bits. This leaves 64 bits for the interface specific portion, called an Interface ID in IPv6. Depending upon the underlying link addressing, the Interface ID can be automatically computed from the link (e.g., MAC address). Such an automatically computed Interface ID is called an EUI-64 identifier, which is the interface MAC address with ff:fe inserted in the middle.

IPv6 packets on the network are of an entirely different format than traditional IPv4 packets and are also encapsulated in a different EtherType (86DD rather than 0800 which is used with IPv4). The details for encapsulating IPv6 in Ethernet frames are described in RFC2462.

Unlike IPv4, IPv6 does not have broadcasts. There are two types of IPv6 addresses — unicast and multicast. Unicast addresses allow direct one-to-one communication between two hosts, whereas multicast addresses allow one-to-many communication. Multicast addresses are used as destinations only. Unicast addresses will have 00 through fe in the most significant octets and multicast addresses will have ff in the most significant octets.

How Are IPv6 Interfaces Configured?

In PowerConnect 8000-series and 8100-series switches software, IPv6 coexists with IPv4. As with IPv4, IPv6 routing can be enabled on VLAN interfaces. Each L3 routing interface can be used for IPv4, IPv6, or both simultaneously.

Neighbor Discovery (ND) protocol is the IPv6 replacement for Address Resolution Protocol (ARP) in IPv4. The IPv6 Neighbor Discovery protocol is described in detail in RFC4861. Router advertisement is part of the Neighbor Discovery process and is required for IPv6. As part of router advertisement, PowerConnect 8000-series and 8100-series switches software supports stateless auto configuration of end nodes. The switch supports both EUI-64 interface identifiers and manually configured interface IDs.

While optional in IPv4, router advertisement is mandatory in IPv6. Router advertisements specify the network prefix(es) on a link which can be used by receiving hosts, in conjunction with an EUI-64 identifier, to autoconfigure a

host's address. Routers have their network prefixes configured and may use EUI-64 or manually configured interface IDs. In addition to zero or more global addresses, each IPv6 interface also has an autoconfigured "link-local" address which is:

- fe80::/10, with the EUI-64 address in the least significant bits.
- Reachable only on the local VLAN — link-local addresses are never routed.
- Not globally unique

Next hop addresses computed by routing protocols are usually link-local addresses.

During the period of transitioning the Internet to IPv6, a global IPv6 Internet backbone may not be available. One transition mechanism is to tunnel IPv6 packets inside IPv4 to reach remote IPv6 islands. When a packet is sent over such a link, it is encapsulated in IPv4 in order to traverse an IPv4 network and has the IPv4 headers removed at the other end of the tunnel.

Default IPv6 Routing Values

IPv6 is disabled by default on the switch and on all interfaces.

Table 38-1 shows the default values for the IP routing features this chapter describes.

Table 38-1. IPv6 Routing Defaults

Parameter	Default Value
IPv6 Unicast Routing Mode	Disabled
IPv6 Hop Limit	Unconfigured
ICMPv6 Rate Limit Error Interval	1000 milliseconds
ICMPv6 Rate Limit Burst Size	100
Interface IPv6 Mode	Disabled
IPv6 Router Route Preferences	Local—0
	Static—1
	OSPFv3 Intra—110
	OSPFv3 Inter—110
	OSPFv3 External—110

Table 38-2 shows the default IPv6 interface values after a VLAN routing interface has been created.

Table 38-2. IPv6 Interface Defaults

Parameter	Default Value
IPv6 Mode	Disabled
DHCPv6 Client Mode	Disabled
Stateless Address AutoConfig Mode	Disabled
Routing Mode	Enabled
Interface Maximum Transmit Unit	1500
Router Duplicate Address Detection Transmits	1
Router Advertisement NS Interval	Not configured
Router Lifetime Interval	1800 seconds
Router Advertisement Reachable Time	0 seconds
Router Advertisement Interval	600 seconds
Router Advertisement Managed Config Flag	Disabled
Router Advertisement Other Config Flag	Disabled
Router Advertisement Suppress Flag	Disabled
IPv6 Destination Unreachables	Enabled

Configuring IPv6 Routing Features (Web)

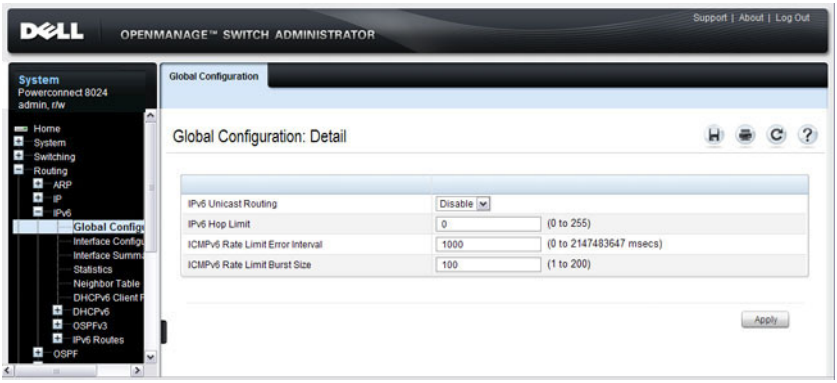
This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring IPv6 unicast routing features on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

Global Configuration

Use the **Global Configuration** page to enable IPv6 forwarding on the router, enable the forwarding of IPv6 unicast datagrams, and configure global IPv6 settings.

To display the page, click **Routing → IPv6 → Global Configuration** in the navigation panel.

Figure 38-1. IPv6 Global Configuration

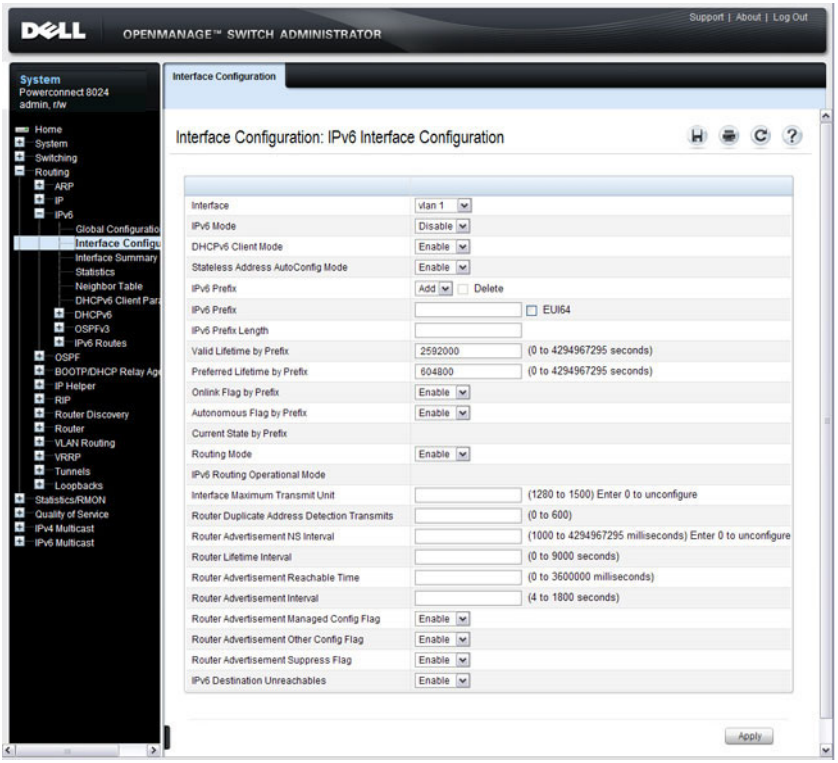


Interface Configuration

Use the **Interface Configuration** page to configure IPv6 interface parameters. This page has been updated to include the IPv6 Destination Unreachables field.

To display the page, click **Routing → IPv6 → Interface Configuration** in the navigation panel.

Figure 38-2. IPv6 Interface Configuration

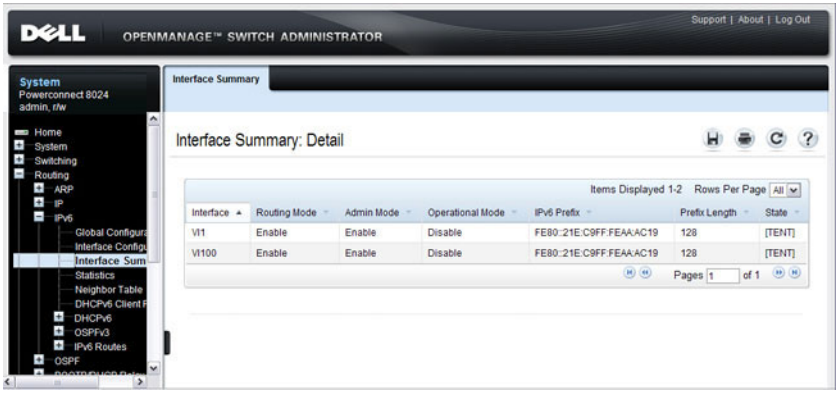


Interface Summary

Use the **Interface Summary** page to display settings for all IPv6 interfaces.

To display the page, click **Routing** → **IPv6** → **Interface Summary** in the navigation panel.

Figure 38-3. IPv6 Interface Summary

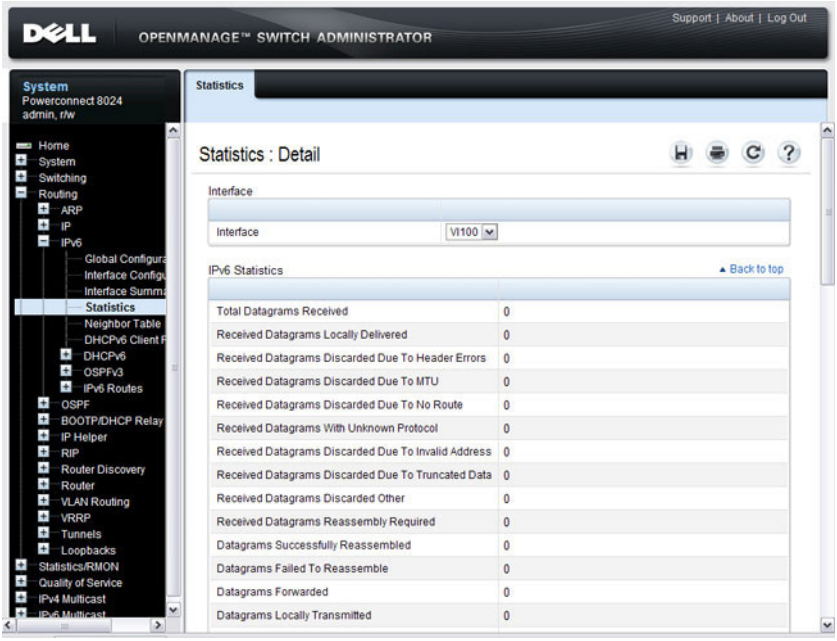


IPv6 Statistics

Use the **IPv6 Statistics** page to display IPv6 traffic statistics for one or all interfaces.

To display the page, click **Routing → IPv6 → IPv6 Statistics** in the navigation panel.

Figure 38-4. IPv6 Statistics

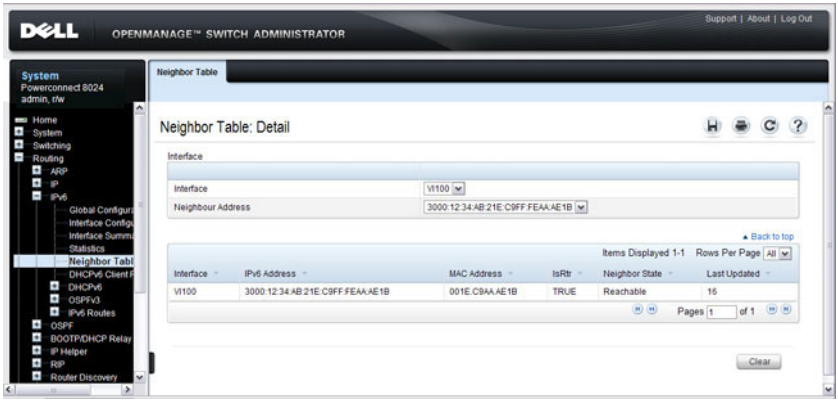


IPv6 Neighbor Table

Use the **IPv6 Neighbor Table** page to display IPv6 neighbor details for a specified interface.

To display the page, click **IPv6** → **IPv6 Neighbor Table** in the navigation panel.

Figure 38-5. IPv6 Neighbor Table

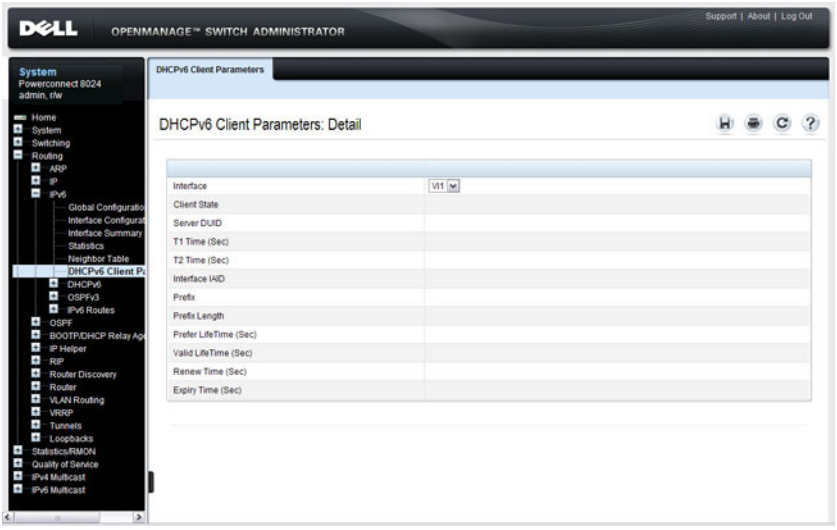


DHCPv6 Client Parameters

Use the **DHCPv6 Client Parameters** page to view information about the network information automatically assigned to an interface by the DHCPv6 server. This page displays information only if the DHCPv6 client has been enabled on an IPv6 routing interface.

To display the page, click **Routing → IPv6 → DHCPv6 Client Parameters** in the navigation panel.

Figure 38-6. DHCPv6 Client Parameters

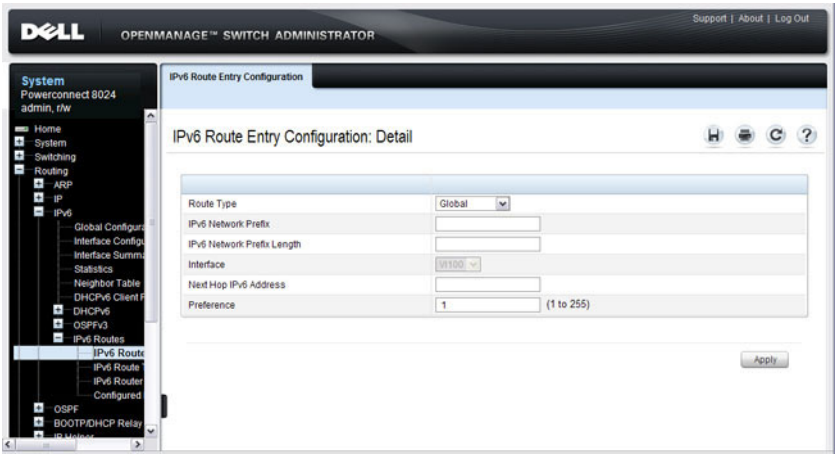


IPv6 Route Entry Configuration

Use the **IPv6 Route Entry Configuration** page to configure information for IPv6 routes.

To display the page, click **Routing → IPv6 → IPv6 Routes → IPv6 Route Entry Configuration** in the navigation panel.

Figure 38-7. IPv6 Route Entry Configuration

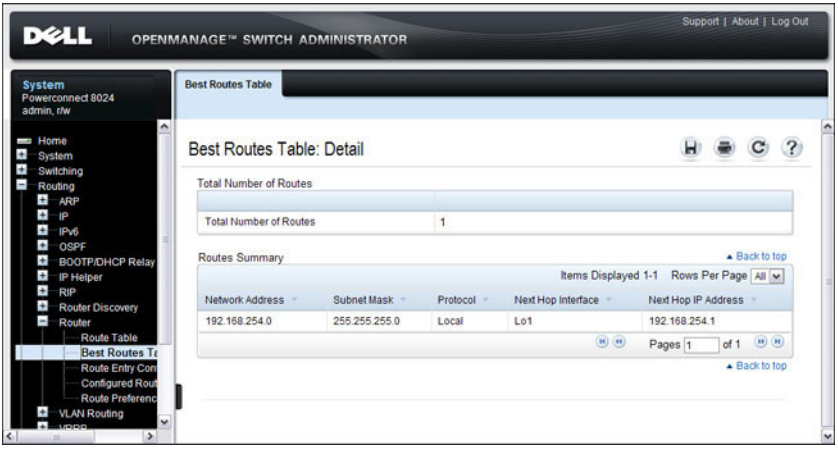


IPv6 Route Table

Use the **IPv6 Route Table** page to display all active IPv6 routes and their settings.

To display the page, click **Routing → IPv6 → IPv6 Routes → IPv6 Route Table** in the navigation panel.

Figure 38-8. IPv6 Route Table

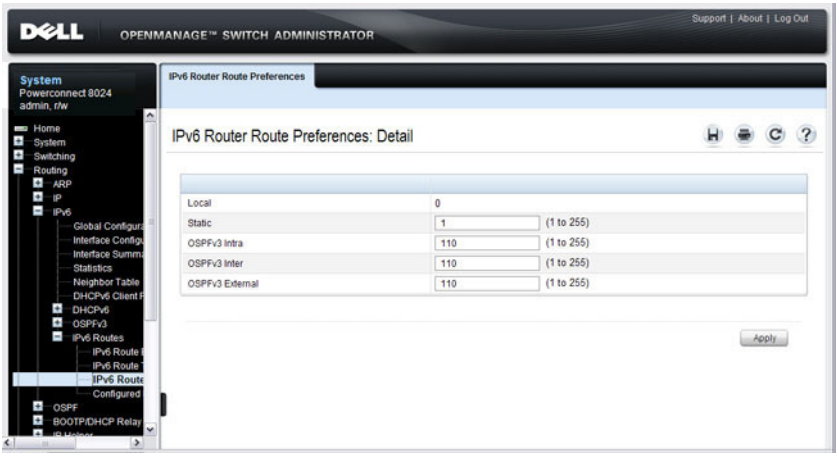


IPv6 Route Preferences

Use the **IPv6 Route Preferences** page to configure the default preference for each protocol. These values are arbitrary values in the range of 1 to 255 and are independent of route metrics. Most routing protocols use a route metric to determine the shortest path known to the protocol, independent of any other protocol. The best route to a destination is chosen by selecting the route with the lowest preference value. When there are multiple routes to a destination, the preference values are used to determine the preferred route. If there is still a tie, the route with the best route metric is chosen. To avoid problems with mismatched metrics, you must configure different preference values for each of the protocols.

To display the page, click **Routing** → **IPv6** → **IPv6 Routes** → **IPv6 Route Preferences** in the navigation panel.

Figure 38-9. IPv6 Route Preferences



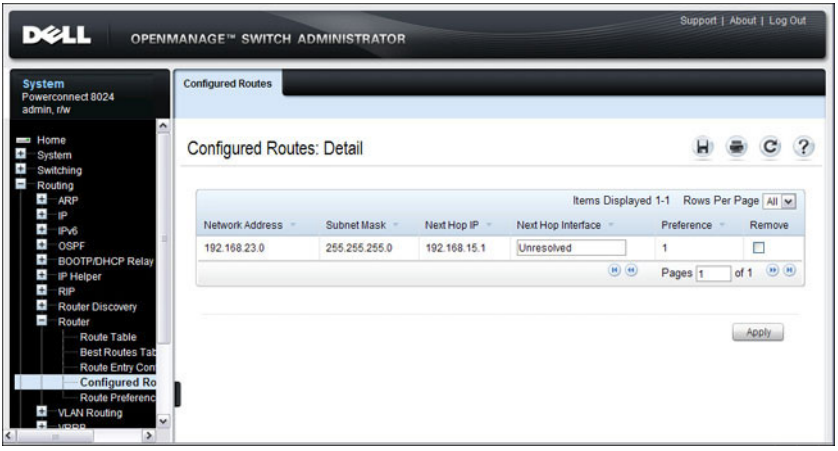
Configured IPv6 Routes

Use the **Configured IPv6 Routes** page to display selected IPv6 routes.

NOTE: For a static reject route, the next hop interface value is Null0. Packets to the network address specified in static reject routes are intentionally dropped.

To display the page, click **Routing** → **IPv6** → **IPv6 Routes** → **Configured IPv6 Routes** in the navigation panel.

Figure 38-10. Configured IPv6 Routes



To remove a configured route, select the check box in the **Delete** column of the route to remove, and click **Apply**.

Configuring IPv6 Routing Features (CLI)

This section provides information about the commands you use to configure IPv6 routing on the switch. For more information about the commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Configuring Global IP Routing Settings

Beginning in Privileged EXEC mode, use the following commands to configure various global IP routing settings for the switch.

Command	Purpose
configure	Enter global configuration mode.
sdm prefer dual-ipv4-and-ipv6 default	Select a Switch Database Management (SDM) template to enable support for both IPv4 and IPv6. Changing the SDM template requires a system reload.
ipv6 unicast-routing	Globally enable IPv6 routing on the switch.
ipv6 hop-limit *limit*	Set the TTL value for the router. The valid range is 0 to 255.
ipv6 icmp error-interval *burst-interval* [*burst-size*]	Limit the rate at which IPv4 ICMP error messages are sent. • *burst-interval* — How often the token bucket is initialized (Range: 0–2147483647 milliseconds). • *burst-size* — The maximum number of messages that can be sent during a burst interval (Range: 1–200).
exit	Exit to Privileged EXEC mode.

Configuring IPv6 Interface Settings

Beginning in Privileged EXEC mode, use the following commands to configure IPv6 settings for VLAN, tunnel, or loopback interfaces.

Command	Purpose
configure	Enter Global Configuration mode.
interface {**vlan** | **tunnel** | **loopback**} *interface-id*	Enter Interface Configuration mode for the specified VLAN, tunnel, or loopback interface.
ipv6 enable	Enable IPv6 on the interface. Configuring an IPv6 address will automatically enable IPv6 on the interface.
ipv6 address {**autoconfig** | **dhcp** | *prefix/prefix-length* [**eui64**]}	Configure the IPv6 address and network prefix length. Setting an IPv6 address enables IPv6 on the interface. You can also use the **ipv6 enable command** to enable IPv6 on the interface without setting an address. Link-local, multicast, IPv4-compatible, and IPv4-mapped addresses are not allowed to be configured. Include the **EUI-64** keyword to have the system add the 64-bit interface ID to the address. You must use a network prefix length of 64 in this case. For VLAN interfaces, use the **dhcp** keyword to enable the DHCPv6 client and obtain an IP address form a network DHCPv6 server.
ipv6 mtu *size*	(VLAN interfaces only) Set the IPv6 Maximum Transmission Unit (MTU) on a routing interface. The IPv6 MTU is the size of the largest IPv6 packet that can be transmitted on the interface without fragmentation. The range is 1280–1500 bytes.
ipv6 traffic-filter *ACL name*	Add an access-list filter to this interface.
ipv6 unreachables	(VLAN interfaces only) Allow the interface to send ICMPv6 Destination Unreachable messages. The **no ipv6 unreachables** command suppresses the ICMPv6 unreachable messages for this interface.
exit	Exit the interface configuration mode.

Configuring IPv6 Neighbor Discovery

Use the following commands to configure IPv6 Neighbor Discovery settings.

Command	Purpose
ipv6 nd prefix *prefix/prefix-length* [{*valid-lifetime*| **infinite**} {*preferred-lifetime*| **infinite**}] [**no-autoconfig**] [**off-link**]	Configure parameters associated with network prefixes that the router advertises in its Neighbor Discovery advertisements. • ipv6-prefix—IPv6 network prefix. • *prefix-length*—IPv6 network prefix length. • *valid-lifetime*—Valid lifetime of the router in seconds. (Range: 0–4294967295 seconds.) • **infinite**—Indicates lifetime value is infinite. • *preferred-lifetime*—Preferred-lifetime of the router in seconds. (Range: 0–4294967295 seconds.) • **no-autoconfig**—Do not use the prefix for auto configuration. • **off-link**—Do not use the prefix for onlink determination.
ipv6 nd ra-interval *maximum minimum*	Set the transmission interval between router Neighbor Discovery advertisements. • *maximum* — The maximum interval duration (Range: 4–1800 seconds). • *minimum* — The minimum interval duration (Range: 3 – (0.75 * maximum) seconds).
ipv6 nd ra-lifetime *seconds*	Set the value that is placed in the Router Lifetime field of the router Neighbor Discovery advertisements sent from the interface. The *seconds* value must be zero, or it must be an integer between the value of the router advertisement transmission interval and 9000 seconds. A value of zero means this router is not to be used as the default router. (Range: 0-9000).
ipv6 nd suppress-ra	Suppress router advertisement transmission on an interface.
ipv6 nd dad attempts *value*	Set the number of duplicate address detection probes transmitted while doing Neighbor Discovery. The range for *value* is 0–600.

Command	Purpose
ipv6 nd ns-interval *milliseconds*	Set the interval between router advertisements for advertised neighbor solicitations. The range is 1000 to 4294967295 milliseconds.
ipv6 nd other-config-flag	Set the *other stateful configuration* flag in router advertisements sent from the interface.
ipv6 nd managed-config-flag	Set the *managed address configuration* flag in router advertisements. When the value is true, end nodes use DHCPv6. When the value is false, end nodes automatically configure addresses.
ipv6 nd reachable-time *milliseconds*	Set the router advertisement time to consider a neighbor reachable after neighbor discovery confirmation.

Configuring IPv6 Route Table Entries and Route Preferences

Beginning in Privileged EXEC mode, use the following commands to configure IPv6 Static Routes.

Command	Purpose
configure	Enter global configuration mode.
ipv6 route *ipv6-prefix/prefix-length* {*next-hop-address* | *interface-type interface-number next-hop-address* } [*preference*]	Configure a static route.Use the keyword **null** instead of the next hop router IP address to configure a static reject route. • *prefix/prefix-length*—The IPv6 network prefix and prefix length that is the destination of the static route. Use the ::/0 form (unspecified address and zero length prefix) to specify a default route. • *interface-type interface-number*—Must be specified when using a link-local address as the next hop. The interface-type can be **vlan** or **tunnel**. • *next-hop-address* —The IPv6 address of the next hop that can be used to reach the specified network. A link-local next hop address must have a prefix length of 128. The next hop address cannot be an unspecified address (all zeros), a multicast address, or a loopback address. If a link local next hop address is specified, the interface (VLAN or tunnel), must also be specified. • *preference*—Also known as Administrative Distance, a metric the router uses to compare this route with routes from other route sources that have the same network prefix. (Range: 1-255). Lower values have precedence over higher values. The default preference for static routes is 1. Routes with a preference of 255 are considered as "disabled" and will not be used for forwarding. Routes with a preference metric of 254 are used by the local router but will never be advertised to other neighboring routers.
ipv6 route *ipv6-prefix/prefix-length* **null** [*preference*]	Configure a static reject route. IPv6 packets matching the reject route will be silently discarded.

Command	Purpose
ipv6 route distance *integer*	Set the default distance (preference) for static IPv6 routes. Lower route preference values are preferred when determining the best route. The default distance (preference) for static routes is 1.
exit	Exit to Global Config mode.

IPv6 Show Commands

Use the following commands in Privileged EXEC mode to view IPv6 configuration status and related data.

Command	Purpose
show sdm prefer	Show the currently active SDM template.
show sdm prefer dual-ipv4-and-ipv6 default	Show parameters for the SDM template.
show ipv6 dhcp interface vlan *vlan-id*	View information about the DHCPv6 lease acquired by the specified interface.
show ipv6 interface {vlan | tunnel | loopback} *interface-id*	View the IP interface configuration information for the specified IPv6 routing interface.
show ipv6 brief	View the global IPv6 settings for the switch.
show ipv6 route [*ipv6-address* | *ipv6-prefix/prefix-length* | *protocol* | *interface-type interface-number*] [**best**]	View the routing table. • *ipv6-address*—Specifies an IPv6 address for which the best-matching route would be displayed. • *protocol*—Specifies the protocol that installed the routes. Is one of the following keywords: connected, ospf, static. • *ipv6-prefix/ prefix-length*—Specifies an IPv6 network for which the matching route would be displayed. • *interface-type interface-number*—Valid IPv6 interface. Specifies that the routes with next-hops on the selected interface be displayed. • **best**—Specifies that only the best routes are displayed. If the connected keyword is selected for protocol, the best option is not available because there are no best or non-best connected routes.
show ipv6 route summary	View summary information about the IPv6 routing table.
show ipv6 route preferences	View detailed information about the IPv6 route preferences.

IPv6 Static Reject and Discard Routes

A static configured route with a next-hop of "null" causes any packet matching the route to disappear or vanish from the network. This type of route is called a "Discard" route if the router returns an ICMP "network-unreachable" message, or is called a "Reject" route if no ICMP message is returned. The PowerConnect 8000-series and 8100-series switches support "Reject" routes, where any packets matching the route network prefix silently disappear.

A common use of a Reject route is to quickly discard packets that cannot be delivered because a valid route to the destination is not known. Without the Reject route, these undeliverable packets will continue to circulate through the network, following the default routes, until their TTL expires. Forwarding packets that cannot be delivered wastes bandwidth, particularly on expensive WAN connections. The Reject route will also suppress a type of "Denial of Service" (DoS) attack where an internal host sends large numbers of packets to unknown destinations, causing congestion of the WAN links.

- ipv6 route ::/0 null 254

  Use this in all routers except the ones with direct Internet connectivity. Routers with direct Internet connectivity should advertise a default route. The effect of this route is that when a router does not have connectivity to the Internet, the router will quickly discard packets that it cannot deliver.

  If the router learns a default route from another router, the learned route will have a lower distance metric and therefore a higher preference. Routes that are more specific (have more bits in the prefix) will have precedence over less specific routes. This will cause packets destined for non-existent networks to be quickly discarded. Also, because of the high distance metric (254), this route will never be advertised to any neighbor routers.

- ipv6 route fc00::/7 null 254

  This route covers the entire ULA (IPv6 private) address space. If you have networks configured in this address space, you will have more specific routes for those networks. The more specific routes (more bits of prefix) will have precedence over this route. Any destinations in this range not known via another, more specific route do not exist. The effect is that packets destined for private networks that do not exist in your network will be quickly discarded instead of being forwarded to the default route.

- ipv6 route 2001::/16 null 254
  ipv6 route 2002::/16 null 254

  These address ranges are reserved and not reachable in the Internet. If for some reason you have local networks in this range, a more specific route will have precedence.

Another use for the Reject route is to prevent internal hosts from communication with specific addresses or ranges of addresses. The effect is the same as an outgoing access-list with a “deny” statement. A route is generally more efficient than an access-list that performs the same function. If you need more fine-grained filtering, such as protocols or port numbers, use the access-list instead.

39

Configuring DHCPv6 Server and Relay Settings

This chapter describes how to configure the switch to dynamically assign network information to IPv6 hosts by using the Dynamic Host Configuration Protocol for IPv6 (DHCPv6).

The topics covered in this chapter include:

- DHCPv6 Overview
- Default DHCPv6 Server and Relay Values
- Configuring the DHCPv6 Server and Relay (Web)
- Configuring the DHCPv6 Server and Relay (CLI)
- DHCPv6 Configuration Examples

DHCPv6 Overview

DHCP is a protocol that is generally used between clients and servers for the purpose of assigning IP addresses, gateways, and other networking definitions such as Domain Name System (DNS) and Network Time Protocol (NTP) parameters. However, IPv6 natively provides IP address auto configuration through IPv6 Neighbor Discovery Protocol (NDP) and through the use of Router Advertisement messages. Thus, the role of DHCPv6 within the network is different than that of DHCPv4 because DHCPv6 is not the primary source for IP address assignment.

DHCPv6 server and client interactions are described by RFC 3315 [6]. There are many similarities between DHCPv6 and DHCPv4 interactions and options, but there are enough differences in the messages and option definitions that there is no DHCPv4 to DHCPv6 migration or interoperability.

What Is a DHCPv6 Pool?

DHCPv6 pools are used to specify information for DHCPv6 server to distribute to DHCPv6 clients. These pools are shared between multiple interfaces over which DHCPv6 server capabilities are configured.

What Is a Stateless Server?

DHCPv6 incorporates the notion of the *stateless* server, where DHCPv6 is not used for IP address assignment to a client; rather, it provides other networking information such as DNS or NTP information. The stateless server behavior is described by RFC 3736 [7], which simply contains descriptions of the portions of RFC 3315 that are necessary for stateless server behavior. In order for a router to drive a DHCPv6 client to utilize stateless DHCPv6, the *other stateful configuration* option must be configured for neighbor discovery on the corresponding IPv6 router interface. This, in turn, causes DHCPv6 clients to send the DHCPv6 Information Request message in response. A DHCPv6 server then responds by providing only networking definitions such as DNS domain name and server definitions, NTP server definitions, or SIP definitions.

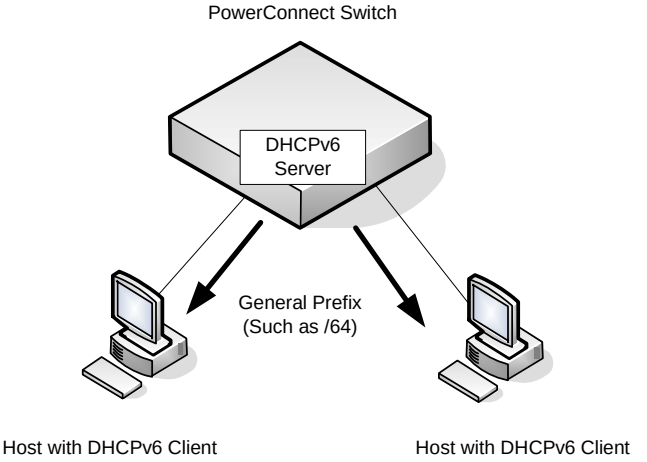
What Is the DHCPv6 Relay Agent Information Option?

The DHCPv6 Relay Agent Information Option allows for various sub-options to be attached to messages that are being relayed by the local router to a DHCPv6 server. The DHCPv6+ server may in turn use this information in determining an address to assign to a DHCPv6 client.RFC 3315 also describes DHCPv6 Relay Agent interactions, which are very much like DHCPv4 Relay Agents. Additionally, there is a DHCPv6 Relay Agent Option described in RFC 4649, which employs very similar capabilities as those described by the DHCPv4 Relay Agent Option in RFC 2132.

What Is a Prefix Delegation?

With the larger address space inherent to IPv6, addresses within a network can be allocated more effectively in a hierarchical fashion. DHCPv6 introduces the notion of prefix delegation as described in RFC 3633 as a way for routers to centralize and delegate IP address assignment. Figure 39-1 depicts a typical network scenario where prefix delegation is used.

Figure 39-1. DHCPv6 Prefix Delegation Scenario



In Figure 39-1, the PowerConnect acts as the Prefix Delegation (PD) server and defines one or more *general* prefixes to allocate and assign addresses to hosts that may be utilizing IPv6 auto-address configuration or acting as DHCPv6 clients.

DHCPv6 clients may request multiple IPv6 prefixes. Also, DHCPv6 clients may request specific IPv6 prefixes. If the configured DHCPv6 pool contains the specific prefix that a DHCPv6 client requests, then that prefix will be delegated to the client. Otherwise, the first available IPv6 prefix within the configured pool will be delegated to the client.

Default DHCPv6 Server and Relay Values

By default, the DHCPv6 server is disabled, and no address pools are configured. VLAN routing interfaces are not configured to perform DHCPv6 server or DHCPv6 relay functions.

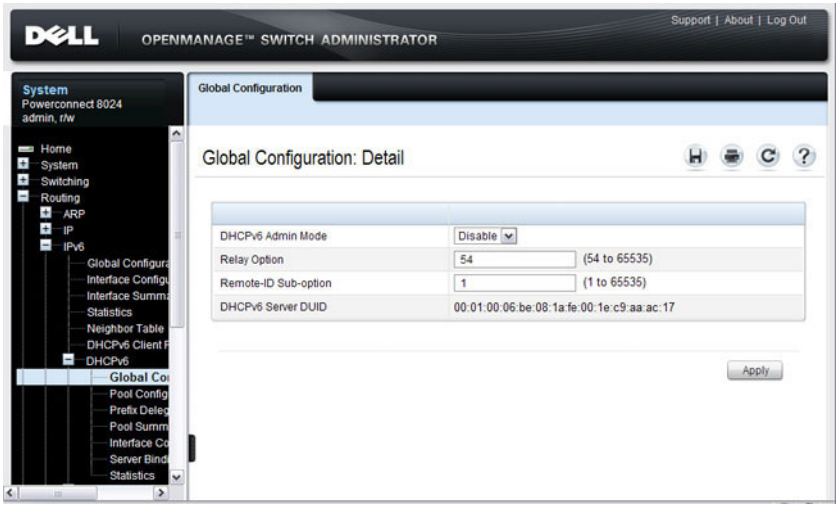
Configuring the DHCPv6 Server and Relay (Web)

This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring the DHCPv6 server on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

DHCPv6 Global Configuration

Use the **Global Configuration** page to configure DHCPv6 global parameters. To display the page, click **Routing → IPv6 → DHCPv6 → Global Configuration** in the navigation panel.

Figure 39-2. DHCPv6 Global Configuration

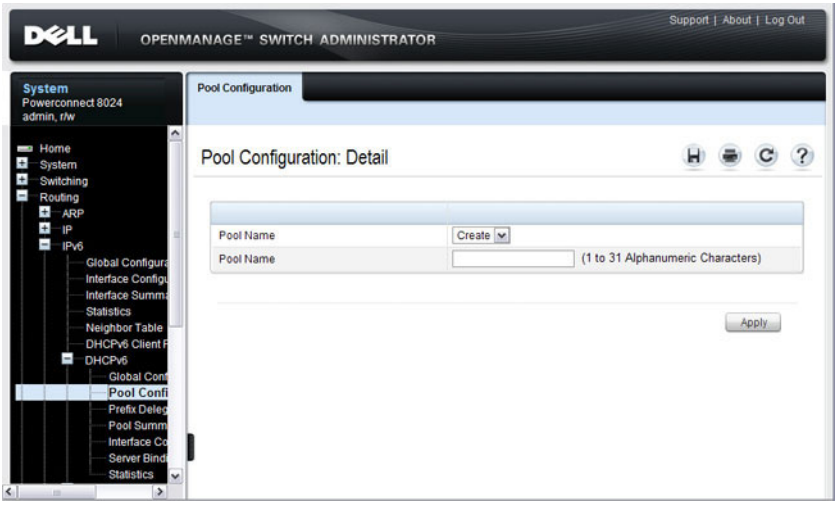


DHCPv6 Pool Configuration

Use the **Pool Configuration** page to set up a pool of DHCPv6 parameters for DHCPv6 clients. The pool is identified with a pool name and contains IPv6 addresses and domain names of DNS servers.

To display the page, click **Routing → IPv6 → DHCPv6 → Pool Configuration** in the navigation panel. Figure 39-3 shows the page when no pools have been created. After a pool has been created, additional fields display.

Figure 39-3. Pool Configuration

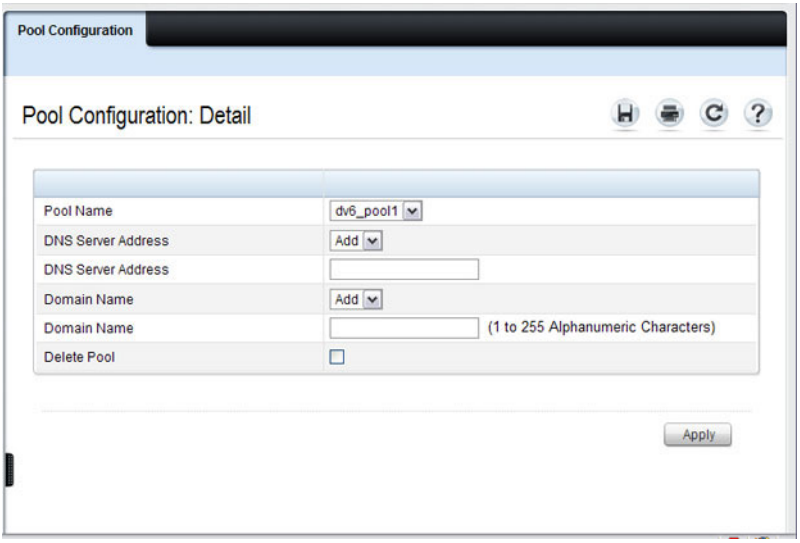


Configuring a DHCPv6 Pool

To configure the pool:

1. Open the **Pool Configuration** page.
2. Select **Create** from the **Pool Name** menu and type a name in the **Pool Name** text box.
3. Click **Apply**.

Figure 39-4. Pool Configuration



4 From the **DNS Server Address** menu, select an existing DNS Server Address to associate with this pool, or select **Add** and specify a new server to add.

5 From the **Domain Name** menu, select an existing domain name to associate with this pool, or select **Add** and specify a new domain name.

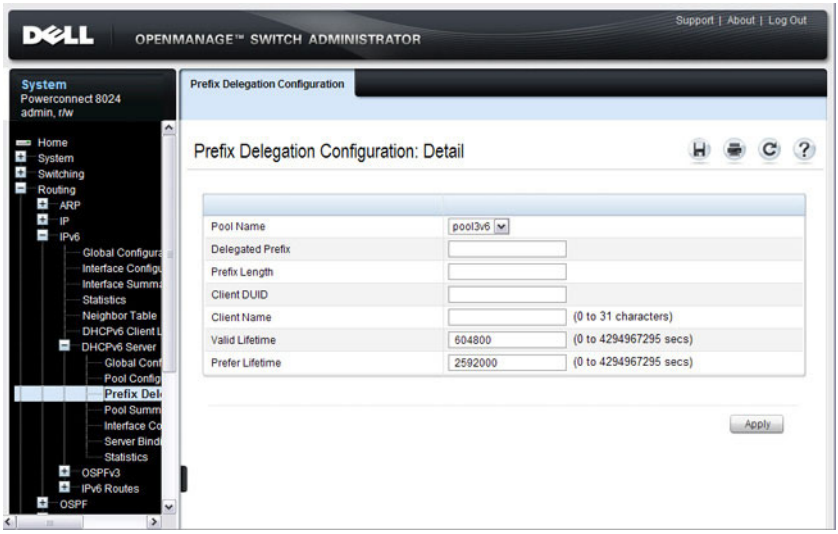
6 Click **Apply**.

Prefix Delegation Configuration

Use the **Prefix Delegation Configuration** page to configure a delegated prefix for a pool. At least one pool must be created using DHCPv6 Pool Configuration before a delegated prefix can be configured.

To display the page, click **Routing → IPv6 → DHCPv6 → Prefix Delegation Configuration** in the navigation panel.

Figure 39-5. Prefix Delegation Configuration

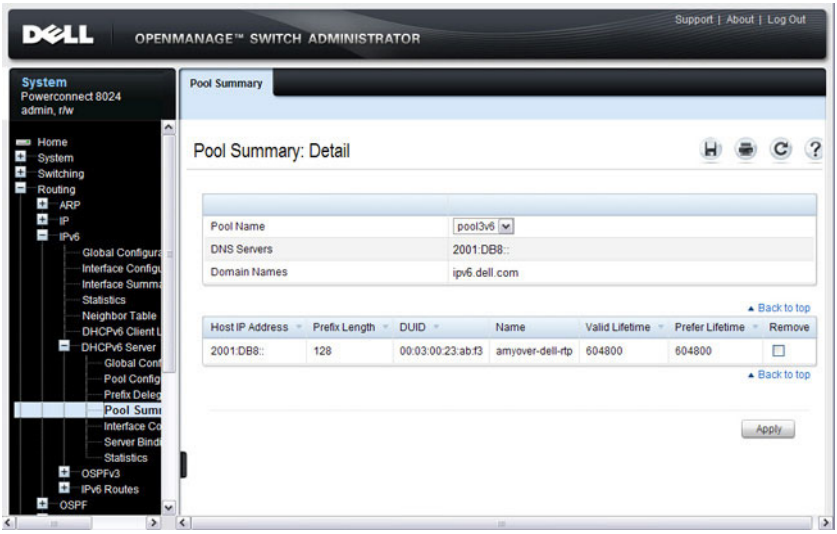


DHCPv6 Pool Summary

Use the **Pool Summary** page to display settings for all DHCPv6 Pools. At least one pool must be created using DHCPv6 Pool Configuration before the Pool Summary displays.

To display the page, click **Routing → IPv6 → DHCPv6 → Pool Summary** in the navigation panel.

Figure 39-6. Pool Summary



DHCPv6 Interface Configuration

Use the **DHCPv6 Interface Configuration** page to configure a DHCPv6 interface.

To display the page, click **Routing → IPv6 → DHCPv6 → Interface Configuration** in the navigation panel. The fields that display on the page depend on the selected interface mode.

Figure 39-7. DHCPv6 Interface Configuration

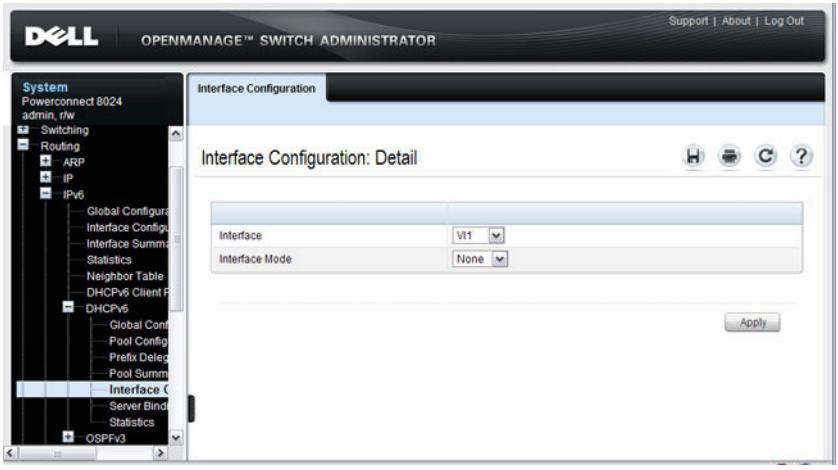


Figure 39-8 shows the screen when the selected interface mode is Server.

Figure 39-8. DHCPv6 Interface Configuration - Server Mode

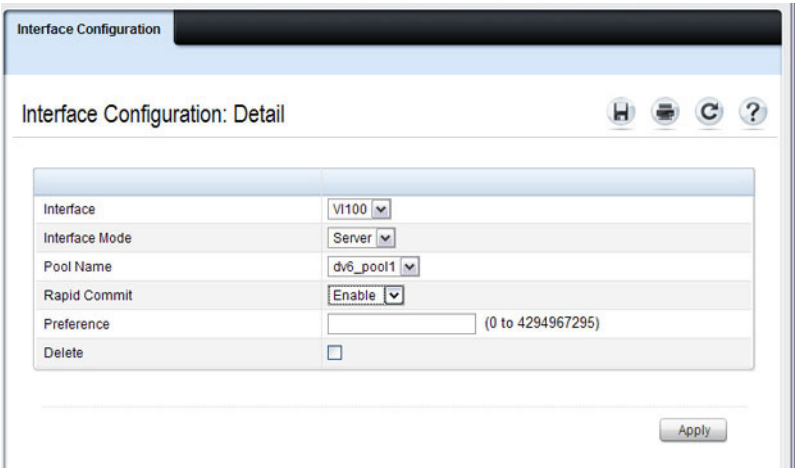
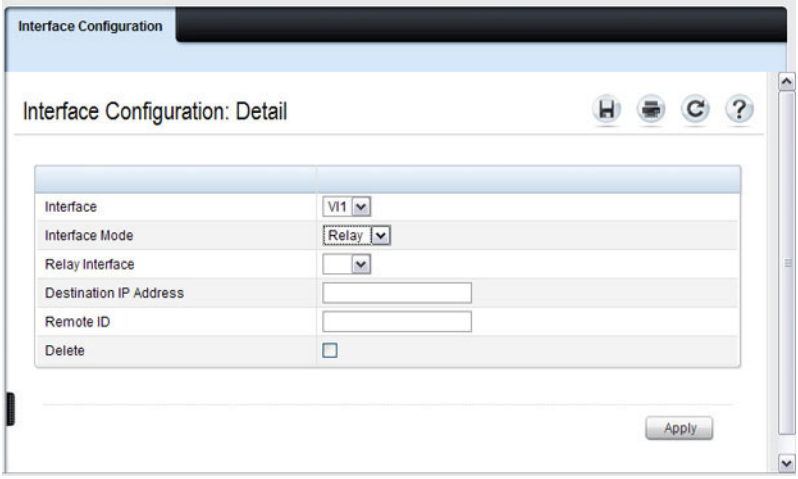


Figure 39-9 shows the screen when the selected interface mode is Relay.

Figure 39-9. DHCPv6 Interface Configuration - Relay Mode

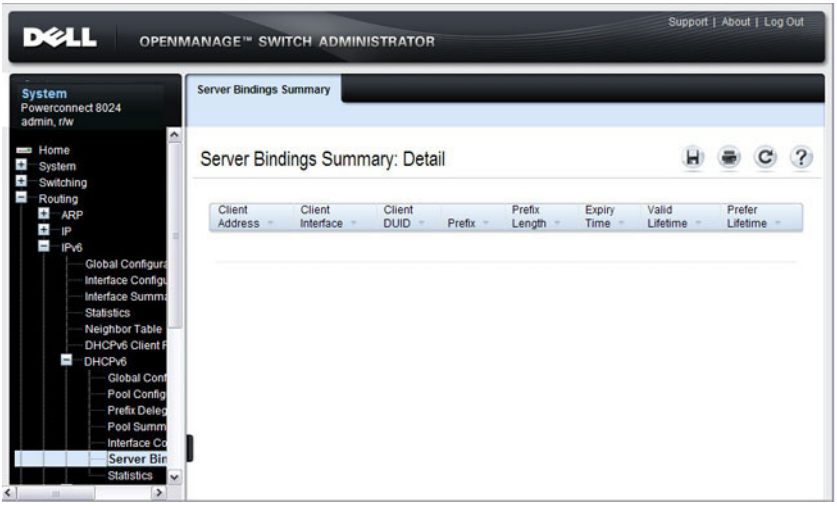


DHCPv6 Server Bindings Summary

Use the **Server Bindings Summary** page to display all DHCPv6 server bindings.

To display the page, click **Routing → IPv6 → DHCPv6 → Bindings Summary** in the navigation panel.

Figure 39-10. Server Bindings Summary

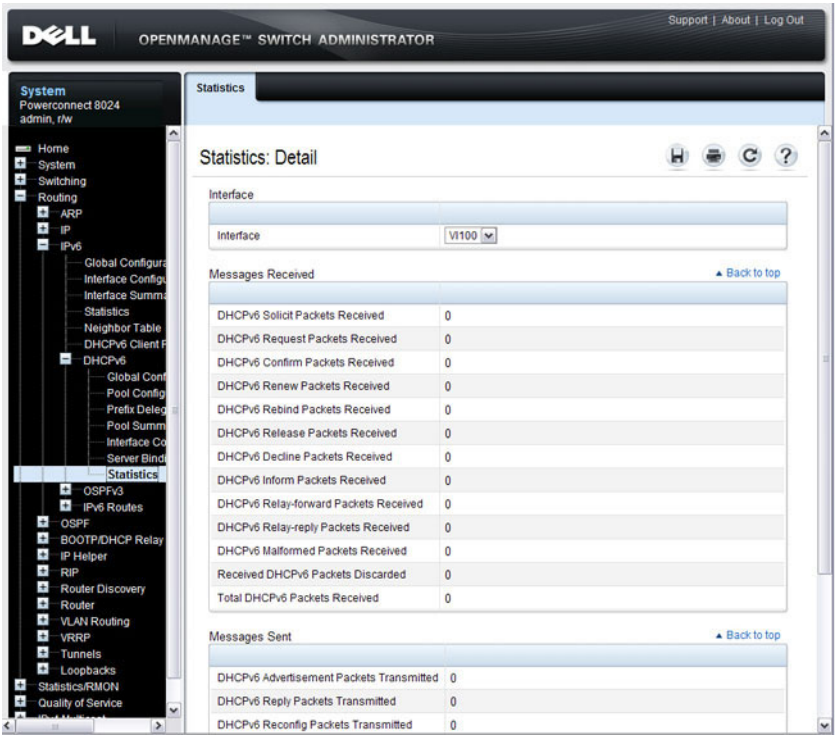


DHCPv6 Statistics

Use the **DHCPv6 Statistics** page to display DHCPv6 statistics for one or all interfaces.

To display the page, click **Routing → IPv6 → DHCPv6 → Statistics** in the navigation panel.

Figure 39-11. DHCPv6 Statistics



Configuring the DHCPv6 Server and Relay (CLI)

This section provides information about the commands you use to configure and monitor the DHCP server and address pools. For more information about the commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Configuring Global DHCP Server and Relay Agent Settings

Beginning in Privileged EXEC mode, use the following commands to configure settings for the DHCPv6 server.

Command	Purpose
configure	Enter Global Configuration mode.
service dhcpv6	Enable the DHCPv6 server.
ipv6 dhcp relay-agent-info-opt *option*	Configure a number to represent the DHCPv6 Relay Agent Information Option. The *option* parameter is an integer from 54–65535.
ipv6 dhcp relay-agent-info-remote-id-subopt *suboption*	Configure a number to represent the DHCPv6 remote-ID sub-option The *suboption* parameter is an integer from 1–65535.
exit	Exit to Privileged EXEC mode.
show ipv6 dhcp	Verify the global DHCPv6 server configuration.

Configuring a DHCPv6 Pool for Stateless Server Support

Beginning in Privileged EXEC mode, use the following commands to create a pool and configure pool parameters for DHCPv6 clients that obtain IPv6 network information dynamically.

Command	Purpose
configure	Enter Global Configuration mode.
ipv6 dhcp pool *name*	Create a DHCPv6 pool and enter DHCPv6 pool configuration mode.
dns-server *ipv6-address*	Set up to 8 IPv6 DNS server addresses to provide to a DHCPv6 client by the DHCPv6 server.

Command	Purpose
domain-name *domain*	Set up to 5 DNS domain names to provide to a DHCPv6 client by the DHCPv6 server.
CTRL + Z	Exit to Privileged EXEC mode.
show ipv6 dhcp pool [*name*]	View the settings for all DHCPv6 pools or for the specified pool.

Configuring a DHCPv6 Pool for Specific Hosts

Beginning in Privileged EXEC mode, use the following commands to create a pool and/or configure pool parameters for specific DHCPv6 clients.

Command	Purpose
configure	Enter Global Configuration mode.
ipv6 dhcp pool *name*	Create a DHCPv6 pool and enter DHCPv6 pool configuration mode.
prefix-delegation *ipv6-prefix/prefix-length client-DUID* [**name** *hostname*] [**valid-lifetime** {*valid-lifetime* | **infinite**}] [**preferred-lifetime** {*preferred-lifetime* | **infinite**}]	Define an IPv6 prefixes within a pool for distributing to specific DHCPv6 Prefix delegation clients. • *prefix/prefix-length*—Delegated IPv6 prefix. • *client-DUID*—DHCP Unique Identifier for the client (e.g. 00:01:00:09:f8:79:4e:00:04:76:73:43:76'). • *hostname*—Client hostname used for logging and tracing. (Range: 0-31 characters.) The command allows spaces in the host name. • *valid-lifetime*—Valid lifetime for delegated prefix. (Range: 0-4294967295 seconds) or use the keyword infinite. • *preferred-lifetime*—Preferred lifetime for delegated prefix. (Range: 0-4294967295 seconds) or use the keyword infinite.
CTRL + Z	Exit to Privileged EXEC mode.
show ipv6 dhcp pool	View information about the DHCPv6 pools configured on the switch.

Configuring DHCPv6 Interface Information

Beginning in Privileged EXEC mode, use the following commands to configure an interface as a DHCPv6 server or a DHCPv6 relay agent. The server and relay functionality are mutually exclusive. In other words, a VLAN routing interface can be configured as a DHCPv6 server or a DHCPv6 relay agent, but not both.

Command	Purpose
configure	Enter Global Configuration mode.
interface {**tunnel** *tunnel-id* | **vlan** *vlan-id*}	Enter interface configuration mode for a tunnel or VLAN routing interface to configure as a DHCPv6 relay agent.
ipv6 dhcp relay {**destination** *relay-address* [**interface vlan** *vlan-id*] | **interface vlan** *vlan-id*} [**remote-id** {**duid-ifid** | *user-defined-string*}]	Configure the interface for DHCPv6 relay functionality. • **destination** — Keyword that sets the relay server IPv6 address. • *relay-address* — An IPv6 address of a DHCPv6 relay server. • **interface** — Sets the relay server interface. • *vlan-id* — A valid VLAN ID. • [remote-id {duid-ifid | *user-defined-string*}] — The Relay Agent Information Option "remote ID" sub-option to be added to relayed messages. This can either be the special keyword duid-ifid, which causes the "remote ID" to be derived from the DHCPv6 server DUID and the relay interface number, or it can be specified as a user-defined string.
exit	Exit to Global Configuration Mode
interface {**tunnel** *tunnel-id* | **vlan** *vlan-id*}	Enter interface configuration mode for a tunnel or VLAN routing interface to configure with DHCPv6 server functionality.

Command	Purpose
ipv6 dhcp server *pool-name* [**rapid-commit**] [**preference** *pref-value*]	Configure DHCPv6 server functionality on the interface. • *pool-name* — The name of the DHCPv6 pool containing stateless and/or prefix delegation parameters • **rapid-commit** — Is an option that allows for an abbreviated exchange between the client and server. • *pref-value* — Preference value—used by clients to determine preference between multiple DHCPv6 servers. (Range: 0-4294967295)
CTRL + Z	Exit to Privileged Exec Mode.
show ipv6 dhcp interface [**tunnel** *tunnel-id* | **vlan** *vlan-id*]	View DHCPv6 information for all interfaces or for the specified interface.

Monitoring DHCPv6 Information

Beginning in Privileged EXEC mode, use the following commands to view bindings, and statistics, and to clear the information.

Command	Purpose
show ipv6 dhcp binding [*address*]	View the current binding information in the DHCP server database. Specify the IP address to view a specific binding.
show ipv6 dhcp statistics	View DHCPv6 server and relay agent statistics.
clear ipv6 dhcp statistics	Reset all DHCPv6 server and relay agent statistics to zero.

DHCPv6 Configuration Examples

This section contains the following examples:

- Configuring a DHCPv6 Stateless Server
- Configuring the DHCPv6 Server for Prefix Delegation
- Configuring an Interface as a DHCPv6 Relay Agent

Configuring a DHCPv6 Stateless Server

This example configures a DHCPv6 pool that will provide information for the DHCPv6 server to distribute to DHCPv6 clients that are members of VLAN 100. To define stateless information for the DHCPv6 server to distribute, multiple DNS domain names and DNS server addresses are defined within the pool.

VLAN routing interface 100 is configured as a DHCPv6 server. Setting NDP on the interface to send the other-config-flag option allows the interface to prompt DHCPv6 clients to request only stateless server information.

To configure the switch:

1. Enable the DHCPv6 feature.

```
console#configure
console(config)#service dhcpv6
```

2. Create the DHCPv6 pool and configure stateless information.

```
console(config)#ipv6 dhcp pool my-pool
console(config-dhcp6s-pool)#domain-name
pengo.dell.com
console(config-dhcp6s-pool)#domain-name dell.com
console(config-dhcp6s-pool)#dns-server
2001:DB8:A328:22C::1
console(config-dhcp6s-pool)#dns-server
2001:DB8:A328:22C::2
```

3. Configure VLAN 100 as a routing interface and assign an IPv6 address.

```
console(config)#interface vlan 100
console(config-if-vlan100)#ipv6 address
2001:DB8:A328:34B::/32
```

4 Configure the DHCPv6 server functionality on VLAN 100. Clients can use the preference value to determine which DHCPv6 server to use when multiple servers exist.

```
console(config-if-vlan100)#ipv6 dhcp server my-
pool preference 10
console(config-if-vlan100)#ipv6 nd other-config-
flag
console(config-if-vlan100)#exit
```

Configuring the DHCPv6 Server for Prefix Delegation

In this example, VLAN routing interface 200 is configured to delegate specific prefixes to certain DHCPv6 clients. The prefix-to-DUID mapping is defined within the DHCPv6 pool.

To configure the switch:

1 Create the DHCPv6 pool and specify the domain name and DNS server information.

```
console(config)#ipv6 dhcp pool my-pool2
console(config-dhcp6s-pool)#domain-name dell.com
console(config-dhcp6s-pool)#dns-server
2001:DB8:A328:22C::1
```

2 Specify the prefix delegations for specific clients. The first two commands provide multiple prefixes to the same client.

```
console(config-dhcp6s-pool)#prefix-delegation
2001:DB8:1000::/32
00:01:00:09:f8:79:4e:00:04:76:73:43:76 valid-
lifetime 600 preferred-lifetime 400

console(config-dhcp6s-pool)#prefix-delegation
2001:DB8:1001::/32
00:01:00:09:f8:79:4e:00:04:76:73:43:76 valid-
lifetime 600 preferred-lifetime 400
```

```
console(config-dhcp6s-pool)#prefix-delegation
2001:DB8:1002::/32
00:01:00:09:f8:79:4e:00:04:76:73:43:76 valid-
lifetime 600 preferred-lifetime 400

console(config-dhcp6s-pool)#exit
```

3 Configure the DHCPv6 server functionality on VLAN 200 and specify the pool to use for DHCPv6 clients.

```
console(config)#interface vlan 200
console(config-if-vlan200)#ipv6 dhcp server my-
pool2 preference 20
```

Configuring an Interface as a DHCPv6 Relay Agent

This example configures a VLAN routing interface as a DHCPv6 Relay. The command defines the destination address of the relay server and the interface used for reachability to the relay server.

To configure the switch:

1 Create VLAN 300 and define its IPv6 address.

```
console(config)#interface vlan 300
console(config-if-vlan300)#ipv6 address
2001:DB8:03a::/64
```

2 Configure the interface as a DHCPv6 relay agent and specify the IPv6 address of the relay server. The command also specifies that the route to the server is through the VLAN 100 routing interface.

```
console(config-if-vlan300)#ipv6 dhcp relay
destination FE80::250:A2FF:FEBF:A056 interface
vlan 100
console(config-if-vlan300)#exit
console(config)#exit
```

3 View the DHCPv6 configuration for VLAN 300.

```
console#show ipv6 dhcp interface vlan 300

IPv6 Interface.......................... ...Vl300
Mode.................................... Relay
Relay Address........... FE80::250:A2FF:FEBF:A056
```

```
Relay Interface Number.................. ...Vl100
Relay Remote ID.........................
Option Flags............................
```

40

Configuring Differentiated Services

This chapter describes how to configure the Differentiated Services (DiffServ) feature. DiffServ enables traffic to be classified into streams and given certain QoS treatment in accordance with defined per-hop behaviors.

The topics covered in this chapter include:

- DiffServ Overview
- Default DiffServ Values
- Configuring DiffServ (Web)
- Configuring DiffServ (CLI)
- DiffServ Configuration Examples

DiffServ Overview

Standard IP-based networks are designed to provide "best effort" data delivery service. Best effort service implies that the network delivers the data in a timely fashion, although there is no guarantee that it will. During times of congestion, packets may be delayed, sent sporadically, or dropped. For typical Internet applications, such as email and file transfer, a slight degradation in service is acceptable and in many cases unnoticeable. Conversely, any degradation of service has undesirable effects on applications with strict timing requirements, such as voice or multimedia.

How Does DiffServ Functionality Vary Based on the Role of the Switch?

How you configure DiffServ support in PowerConnect 8000-series and 8100-series switches software varies depending on the role of the switch in your network:

- **Edge device**: An edge device handles ingress traffic, flowing towards the core of the network, and egress traffic, flowing away from the core. An edge device segregates inbound traffic into a small set of traffic classes, and is responsible for determining a packet's classification. Classification is primarily based on the contents of the Layer 3 and Layer 4 headers, and is recorded in the Differentiated Services Code Point (DSCP) added to a packet's IP header.
- **Interior node**: A switch in the core of the network is responsible for forwarding packets, rather than for classifying them. It decodes the DSCP in an incoming packet, and provides buffering and forwarding services using the appropriate queue management algorithms.

Before configuring DiffServ on PowerConnect 8000-series and 8100-series switches, you must determine the QoS requirements for the network as a whole. The requirements are expressed in terms of rules, which are used to classify inbound or outbound traffic on a particular interface.

What Are the Elements of DiffServ Configuration?

During configuration, you define DiffServ rules in terms of classes, policies, and services:

- **Class**: A class consists of a set of rules that identify which packets belong to the class. Inbound traffic is separated into traffic classes based on Layer 2, Layer 3, and Layer 4 header data. The class type **All** is supported; this specifies that every match criterion defined for the class must be true for a match to occur.
- **Policy**: A policy defines the QoS attributes for one or more traffic classes. An attribute identifies the action taken when a packet matches a class rule. An example of an attribute is to mark a packet. The switch supports the ability to assign traffic classes to output CoS queues, and to mirror incoming packets in a traffic stream to a specific egress interface (physical port or LAG).

PowerConnect 8000-series and 8100-series switches software supports the **Traffic Conditioning Policy** type which is associated with an inbound traffic class and specifies the actions to be performed on packets meeting the class rules:

  – Marking the packet with a given DSCP, IP precedence, or CoS value. Traffic to be processed by the DiffServ feature requires an IP header if the system uses IP Precedence or IP DSCP marking.
  – Policing packets by dropping or re-marking those that exceed the class's assigned data rate.
  – Counting the traffic within the class.

- **Service:** Assigns a policy to an interface for inbound traffic.

NOTE: You can use an 802.1X authenticator or RADIUS server to dynamically assign DiffServ filters to ports when a host connects to a port and authenticates by using 802.1X. For more information, see "How Does the Authentication Server Assign DiffServ Filters?" on page 465

Default DiffServ Values

Table 40-1 shows the global default values for DiffServ.

Table 40-1. DiffServ Global Defaults

Parameter	Default Value
DiffServ	Enabled
Classes	None configured
Policies	None configured
Services	None configured

Configuring DiffServ (Web)

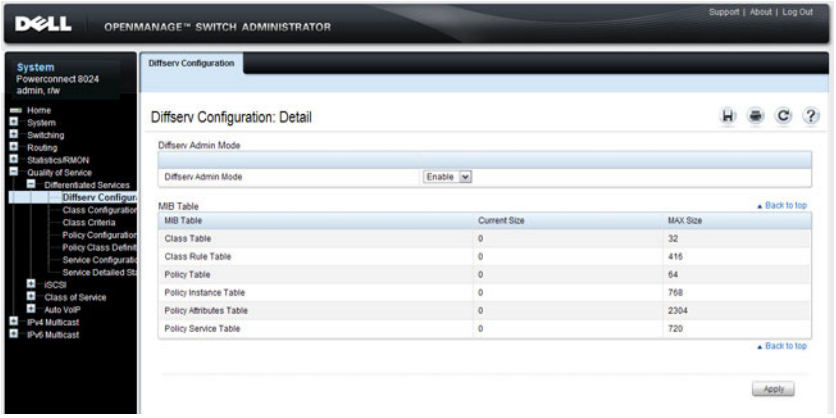
This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring DiffServ features on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

DiffServ Configuration

Use the **DiffServ Configuration** page to display the DiffServ administrative mode setting as well as the current and maximum number of rows in each of the main DiffServ private MIB tables.

To display the page, click **Quality of Service → Differentiated Services → DiffServ Configuration** in the navigation panel.

Figure 40-1. DiffServ Configuration

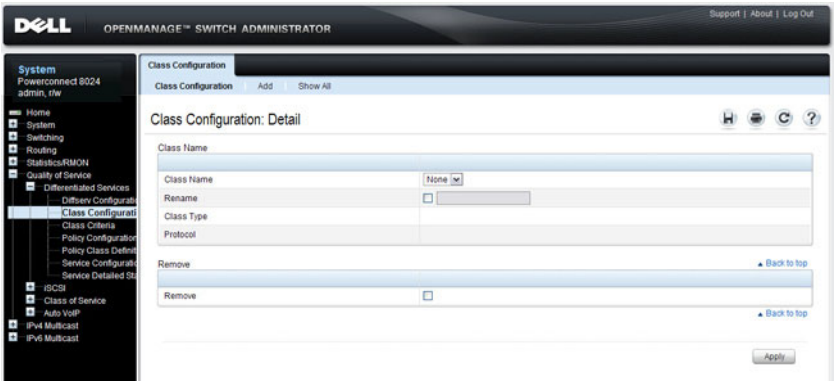


Class Configuration

Use the **DiffServ Class Configuration** page to add a new DiffServ class name, or to rename or delete an existing class.

To display the page, click **Quality of Service → Differentiated Services → Class Configuration** in the navigation panel.

Figure 40-2. DiffServ Class Configuration

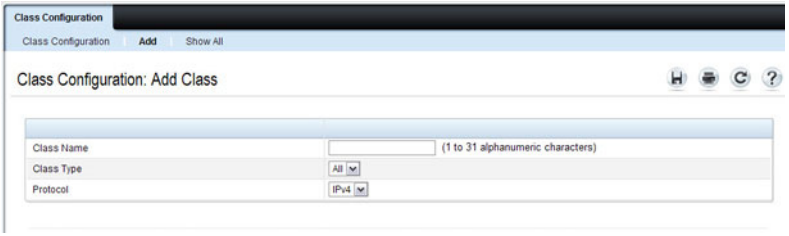


Adding a DiffServ Class

To add a DiffServ class:

1. From the **DiffServ Class Configuration** page, click **Add** to display the **Add Class** page.

Figure 40-3. Add DiffServ Class

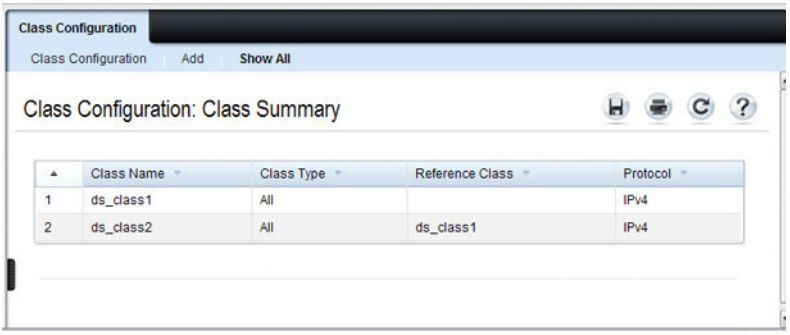


2. Enter a name for the class and select the protocol to use for class match criteria.

3 Click **Apply** to add the new class.

4 To view a summary of the classes configured on the switch, click **Show All**.

Figure 40-4. View DiffServ Class Summary

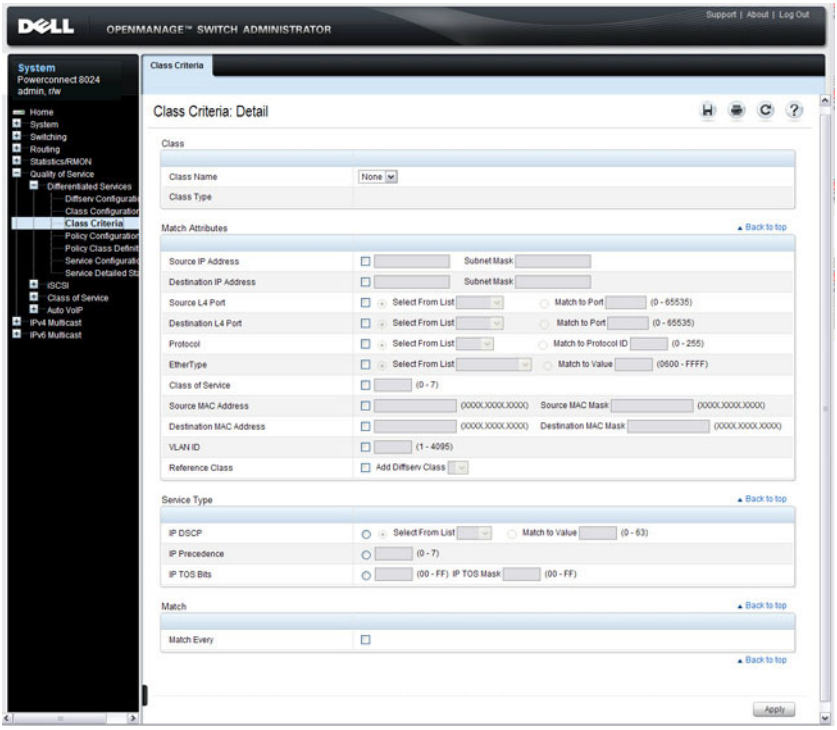


Class Criteria

Use the **DiffServ Class Criteria** page to define the criteria to associate with a DiffServ class. As packets are received, these DiffServ classes are used to identify packets.

To display the page, click **Quality of Service → Differentiated Services → Class Criteria** in the navigation panel.

Figure 40-5. DiffServ Class Criteria

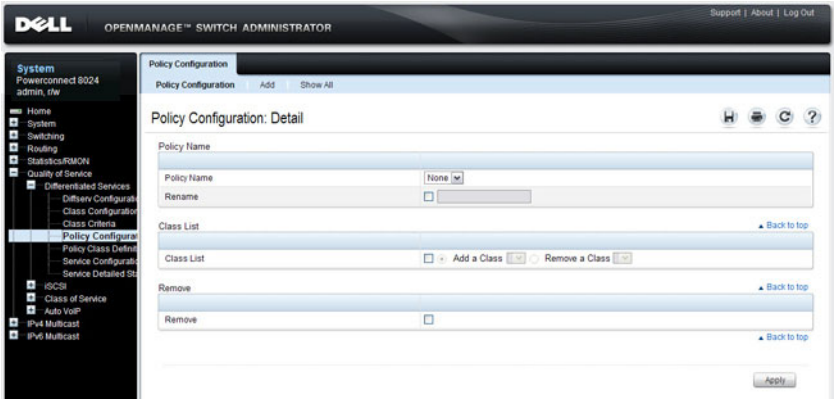


Policy Configuration

Use the **DiffServ Policy Configuration** page to associate a collection of classes with one or more policy statements.

To display the page, click **Quality of Service** → **Differentiated Services** → **Policy Configuration** in the navigation panel.

Figure 40-6. DiffServ Policy Configuration

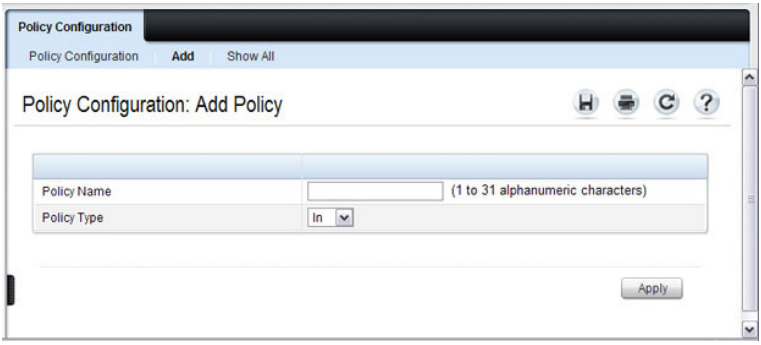


Adding a New Policy Name

To add a policy:

1 From the **DiffServ Policy Configuration** page, click **Add** to display the **Add Policy** page.

Figure 40-7. Add DiffServ Policy

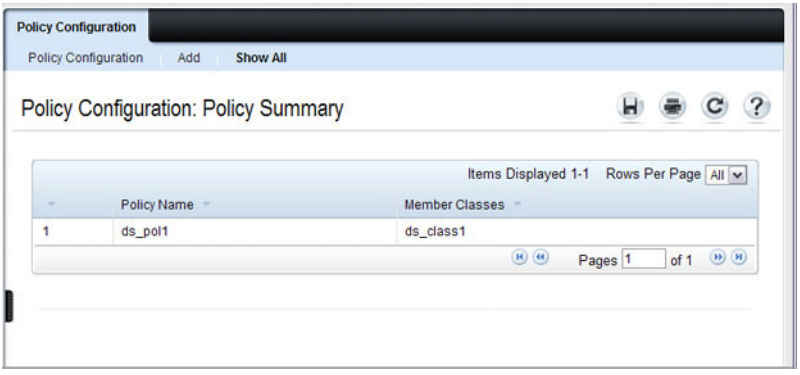


2 Enter the new **Policy Name**.

3 Click **Apply** to save the new policy.

4 To view a summary of the policies configured on the switch, click **Show All**.

Figure 40-8. View DiffServ Policies

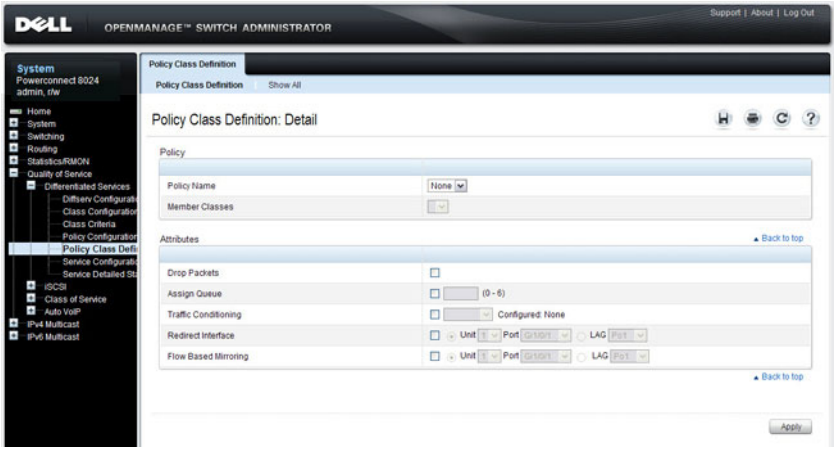


Policy Class Definition

Use the **DiffServ Policy Class Definition** page to associate a class to a policy, and to define attributes for that policy-class instance.

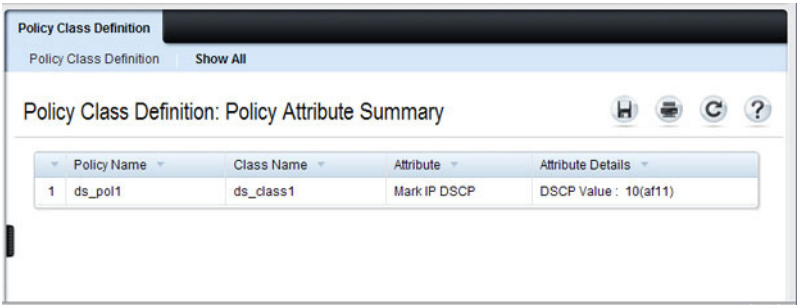
To display the page, click **Quality of Service → Differentiated Services → Policy Class Definition** in the navigation panel.

Figure 40-9. DiffServ Policy Class Definition



To view a summary of the policy attributes, click **Show All**.

Figure 40-10. Policy Attribute Summary



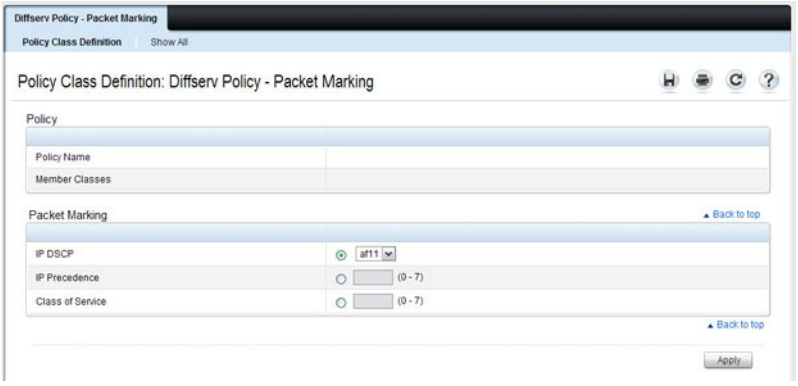
Packet Marking Traffic Condition

Follow these steps to have packets that match the class criteria for this policy marked with a marked with either an IP DSCP, IP precedence, or CoS value:

1 Select Marking from the **Traffic Conditioning** drop-down menu on the **DiffServ Policy Class Definition** page.

The **Packet Marking** page displays.

Figure 40-11. Policy Class Definition - Packet Marking



2 Select **IP DSCP**, **IP Precedence**, or **Class of Service** to mark for this policy-class.

3 Select or enter a value for this field.

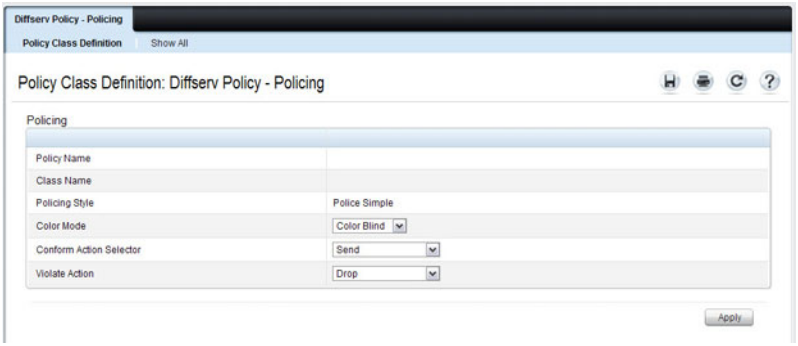
4 Click **Apply** to define the policy-class.

Policing Traffic Condition

Follow these steps to perform policing on the packets that match this policy class:

1 Select **Policing** from the **Traffic Conditioning** drop-down menu on the **DiffServ Policy Class Definition** page to display the **DiffServ Policy - Policing** page.

Figure 40-12. Policy Class Definition - Policing



The **DiffServ Policy - Policing** page displays the **Policy Name**, **Class Name**, and **Policing Style**.

Select a value for the following fields:

- **Color Mode** — The type of color policing used: Color Blind or Color Aware.
- **Conform Action Selector** — The action taken on packets that are considered conforming (below the police rate). Options are Send, Drop, Mark CoS, Mark IP DSCP, Mark IP Precedence.
- **Violate Action** — The action taken on packets that are considered non-conforming (above the police rate). Options are Send, Drop, Mark CoS, Mark IP DSCP, Mark IP Precedence.

2 Click **Apply**.

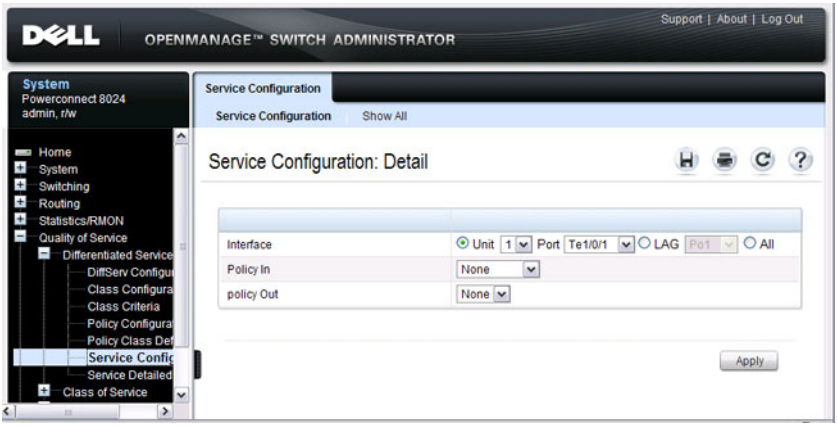
The policy-class is defined, and the device is updated.

Service Configuration

Use the **DiffServ Service Configuration** page to activate a policy on a port.

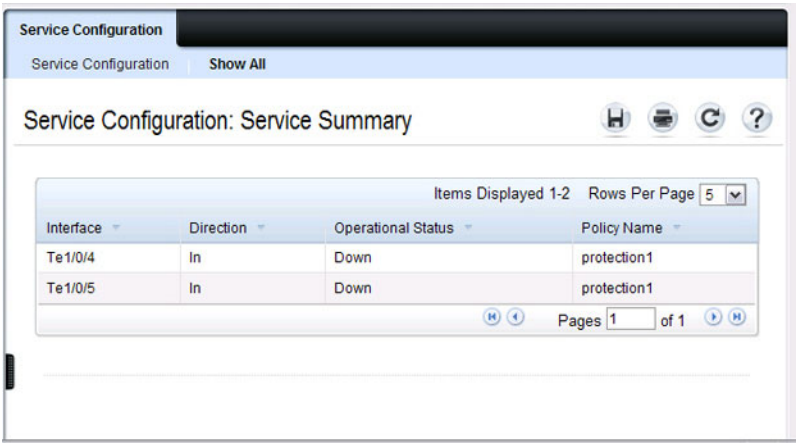
To display the page, click **Quality of Service → Differentiated Services → Service Configuration** in the navigation panel.

Figure 40-13. DiffServ Service Configuration



To view a summary of the services configured on the switch, click **Show All**.

Figure 40-14. DiffServ Service Summary

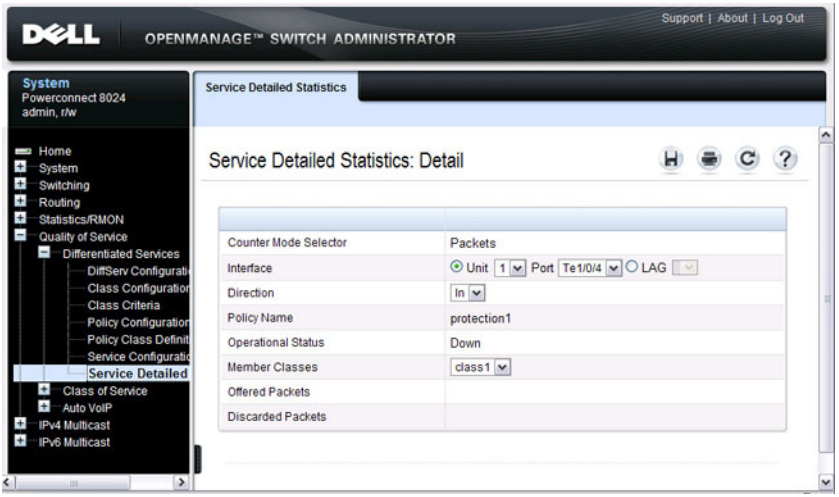


Service Detailed Statistics

Use the **DiffServ Service Detailed Statistics** page to display packet details for a particular port and class.

To display the page, click **Quality of Service → Differentiated Services → Service Detailed Statistics** in the navigation panel.

Figure 40-15. DiffServ Service Detailed Statistics



Flow-Based Mirroring

Use the **Flow-Based Mirroring** page to create a mirroring session in which the traffic that matches the specified policy and member class is mirrored to a destination port.

To display the **Flow-Based Mirroring** page, click **Switching** → **Ports** → **Traffic Mirroring** → **Flow-Based Mirroring** in the navigation panel.

Figure 40-16. Flow-Based Mirroring



Configuring DiffServ (CLI)

This section provides information about the commands you use to configure DiffServ settings on the switch. For more information about the commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

DiffServ Configuration (Global)

Beginning in Privileged Exec mode, use the following commands in to configure the global DiffServ mode and view related settings.

CLI Command	Description
configure	Enter global configuration mode.
diffserv	Set the DiffServ operational mode to active.
exit	Exit to Privileged EXEC mode.
show diffserv	Display the DiffServ general information, which includes the current administrative mode setting as well as the current and maximum number of DiffServ components.

DiffServ Class Configuration for IPv4

Beginning in Privileged Exec mode, use the following commands to configure DiffServ classes for IPv4 and view related information.

CLI Command	Description
configure	Enter global configuration mode.
class-map match-all *class-map-name*	Define a new DiffServ class and enter Class-Map Configuration mode for the specified class. **NOTE:** To enter Class-Map Configuration mode for a class that has already been created, use the **class-map** *class-map-name* command.
match any	Configure a match condition for all the packets.
match class-map	Add to the specified class definition the set of match conditions defined for another class.

CLI Command	Description
match cos	Add to the specified class definition a match condition for the Class of Service value.
match destination-address mac	Add to the specified class definition a match condition based on the destination MAC address of a packet.
match dstip	Add to the specified class definition a match condition based on the destination IP address of a packet.
match dstl4port	Add to the specified class definition a match condition based on the destination layer 4 port of a packet using a single keyword, or a numeric notation.
match ethertype	Add to the specified class definition a match condition based on the value of the ethertype.
match ip dscp	Add to the specified class definition a match condition based on the value of the IP DiffServ Code Point (DSCP) field in a packet.
match ip precedence	Add to the specified class definition a match condition based on the value of the IP.
match ip tos	Add to the specified class definition a match condition based on the value of the IP TOS field in a packet.
match protocol	Add to the specified class definition a match condition based on the value of the IP Protocol field in a packet using a single keyword notation or a numeric value notation.
match secondary-cos	Configure a match condition based on a secondary COS value.
match secondary-vlan	Configure a match condition based on a secondary VLAN value.
match source-address mac	Add to the specified class definition a match condition based on the source MAC address of the packet.

CLI Command	Description
match srcip	Add to the specified class definition a match condition based on the source IP address of a packet.
match srcl4port	Add to the specified class definition a match condition based on the source layer 4 port of a packet using a single keyword, a numeric notation, or a numeric range notation.
match vlan	Add to the specified class definition a match condition based on the value of the layer 2 VLAN Identifier field.

DiffServ Class Configuration for IPv6

Beginning in Privileged Exec mode, use the following commands to configure DiffServ classes for IPv6 and view related information.

CLI Command	Description
configure	Enter global configuration mode.
class-map match-all *class-map-name* **ipv6**	Define a new DiffServ class.
match any	Configure a match condition for all the packets.
match class-map	Add to the specified class definition the set of match conditions defined for another class.
match dstip6	Add to the specified class definition a match condition based on the destination IPv6 address of a packet.
match dstl4port	Add to the specified class definition a match condition based on the destination layer 4 port of a packet using a single keyword, or a numeric notation.
match ip6flowlbl	Add to the specified class definition a match condition based on the IPv6 flow label of a packet.
match ip dscp	Add to the specified class definition a match condition based on the value of the IP DiffServ Code Point (DSCP) field in a packet.

CLI Command	Description
match protocol	Add to the specified class definition a match condition based on the value of the IP Protocol field in a packet using a single keyword notation or a numeric value notation.
match source-address mac	Add to the specified class definition a match condition based on the source MAC address of the packet.
match srcip6	Add to the specified class definition a match condition based on the source IPv6 address of a packet.
match srcl4port	Add to the specified class definition a match condition based on the source layer 4 port of a packet using a single keyword, a numeric notation, or a numeric range notation.

DiffServ Policy Creation

Beginning in Privileged Exec mode, use the following commands to configure DiffServ policies and view related information.

CLI Command	Description
configure	Enter global configuration mode.
policy-map *policy-name* **in**	Create a new DiffServ policy for ingress traffic and enter Policy Map Configuration mode for the policy.
exit	Exit to Privilege Exec mode.
show policy-map	Displays all configuration information for the specified policy.
show policy-map *interface* **in**	Displays policy-oriented statistics information for the specified interface.

DiffServ Policy Attributes Configuration

Beginning in Privilege Exec mode, use the following commands to configure policy attributes and view related information.

CLI Command	Description
configure	Enter global configuration mode.
policy-map *policy-map-name*	Enter Policy Map Configuration mode for the specified policy.
class *class-name*	Create an instance of a class definition within the specified policy for the purpose of defining treatment of the traffic class through subsequent policy attribute statements. Also enters Policy-Class-Map Configuration mode for the policy-class-map instance.
assign-queue *queue-id*	Modify the queue ID (range: 0–6) to which the associated traffic stream is assigned.
police-simple {*datarate* *burstsize* **conform-action** {**drop** | **set-cos-transmit** *cos* | **set-prectransmit** *cos* | **set-dscp-transmit** *dscpval* | **transmit**} [**violateaction** {**drop** | **set-cos-transmit** *cos* | **set-prec-transmit** *cos* | **set-dscp-transmit** *dscpval* | **transmit**}]}	Establish the traffic policing style for the specified class. The simple form of the police command uses a single data rate and burst size, resulting in two outcomes: conform and nonconform. • *datarate* — Data rate in kilobits per second (kbps). (Range: 1–4294967295) • *burstsize* — Burst size in Kbps (Range: 1–128) • **conform action** — Indicates what happens when the packet is conforming to the policing rule: it could be dropped, it could have its COS modified, it could have its IP precedence modified, or it could have its DSCP modified. The same actions are available for packets that do not conform to the policing rule. • *cos* — Class of Service value. (Range: 0–7) • *dscpval* — DSCP value. (Range: 0–63 or a keyword from this list, **af11**, **af12**, **af13**, **af21**, **af22**, **af23**, **af31**, **af32**, **af33**, **af41**, **af42**, **af43**, **be**, **cs0**, **cs1**, **cs2**, **cs3**, **cs4**, **cs5**, **cs6**, **cs7**, **ef**)

CLI Command	Description
conform-color *class-map-name* [**exceed-color** *class-map-name*]	Specify the color class for color-aware policing. The action for the policy-class-map instance must be set to **police-simple** before issuing the **conform-color** command.
drop	Specify that all packets for the associated traffic stream are to be dropped at ingress.
mark cos *cos-value*	Mark all packets for the associated traffic stream with the specified class of service value (range: 0–7) in the priority field of the 802.1p header.
mark ip-dscp *dscp-value*	Mark all packets for the associated traffic stream with the specified IP DSCP value.
mark ip-precedence *value*	Mark all packets for the associated traffic stream with the specified IP precedence value (range: 0–7).
mirror *interface* | **redirect** *interface*	Use **mirror** to mirror all packets for the associated traffic stream that matches the defined class to the specified destination port or LAG. Use **redirect** to specify that all incoming packets for the associated traffic stream are redirected to the specified destination port or LAG.
exit	Exit to Policy-Map Config mode.
exit	Exit to Global Config mode.
exit	Exit to Privilege Exec mode.
show policy-map *policy-map-name*	Displays configuration information for the specified policy.

DiffServ Service Configuration

Beginning Privilege Exec mode, use the following commands to associate a policy with an interface and view related information.

CLI Command	Description
configure	Enter Global Configuration mode.
service-policy {in | out} *policy-map-name*	Attach a policy to an interface in the inbound or outbound direction. This command can be used in either Global Configuration mode (for all system interfaces) or Interface Configuration mode (for a specific interface).
exit	Exit to Privilege Exec mode.
show diffserv service brief [in | out]	Display all interfaces in the system to which a DiffServ policy has been attached.
show diffserv service interface *interface* **{in | out}**	Display policy service information for the specified interface, where *interface* is replaced by **tengigabitethernet** *unit/slot/port* or **port-channel** *port-channel number*.
show service-policy {in | out}	Display a summary of policy-oriented statistics information for all interfaces.

DiffServ Configuration Examples

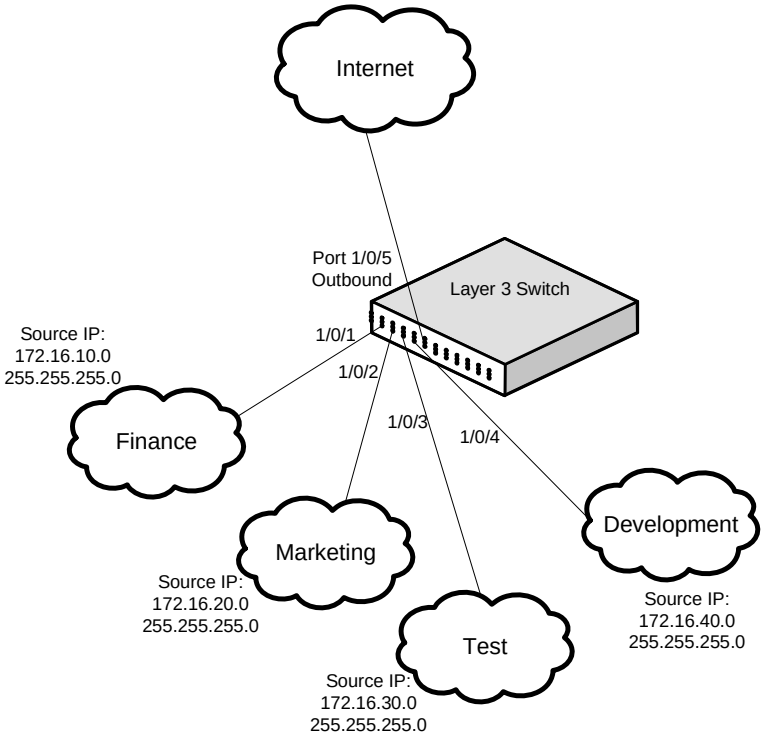
This section contains the following examples:

- Providing Subnets Equal Access to External Network
- DiffServ for VoIP

Providing Subnets Equal Access to External Network

This example shows how a network administrator can provide equal access to the Internet (or other external network) to different departments within a company. Each of four departments has its own Class B subnet that is allocated 25% of the available bandwidth on the port accessing the Internet.

Figure 40-17. DiffServ Internet Access Example Network Diagram



The following commands show how to configure the DiffServ example depicted in Figure 40-17.

1. Enable DiffServ operation for the switch.

```
console#config
console(config)#diffserv
```

2. Create a DiffServ class of type *all* for each of the departments, and name them. Also, define the match criteria—Source IP address—for the new classes.

```
console(config)#class-map match-all finance_dept
console(config-classmap)#match srcip 172.16.10.0
255.255.255.0
console(config-classmap)#exit

console(config)#class-map match-all marketing_dept
console(config-classmap)#match srcip 172.16.20.0
255.255.255.0
console(config-classmap)#exit

console(config)#class-map match-all test_dept
console(config-classmap)#match srcip 172.16.30.0
255.255.255.0
console(config-classmap)#exit

console(config)#class-map match-all
development_dept
console(config-classmap)#match srcip 172.16.40.0
255.255.255.0
console(config-classmap)#exit
```

3. Create a DiffServ policy for inbound traffic named *internet_access,* adding the previously created department classes as instances within this policy. This policy uses the assign-queue attribute to put each department's traffic on a different egress queue. This is how the DiffServ inbound policy connects to the CoS queue settings established below.

```
console(config)#policy-map internet_access in
console(config-policy-map)#class finance_dept
```

```
console(config-policy-classmap)#assign-queue 1
console(config-policy-classmap)#exit

console(config-policy-map)#class marketing_dept
console(config-policy-classmap)#assign-queue 2
console(config-policy-classmap)#exit

console(config-policy-map)#class test_dept
console(config-policy-classmap)#assign-queue 3
console(config-policy-classmap)#exit

console(config-policy-map)#class development_dept
console(config-policy-classmap)#assign-queue 4
console(config-policy-classmap)#exit
console(config-policy-map)#exit
```

4 Attach the defined policy to 10-Gigabit Ethernet interfaces 1/0/1 through 1/0/4 in the inbound direction

```
console(config)#interface tengigabitethernet 1/0/1
console(config-if-Te1/0/1)#service-policy in
internet_access
console(config-if-Te1/0/1)#exit

console(config)#interface tengigabitethernet 1/0/2
console(config-if-Te1/0/2)#service-policy in
internet_access
console(config-if-Te1/0/2)#exit

console(config)#interface tengigabitethernet 1/0/3
console(config-if-Te1/0/3)#service-policy in
internet_access
console(config-if-Te1/0/3)#exit

console(config)#interface tengigabitethernet 1/0/4
console(config-if-Te1/0/4)#service-policy in
internet_access
console(config-if-Te1/0/4)#exit
```

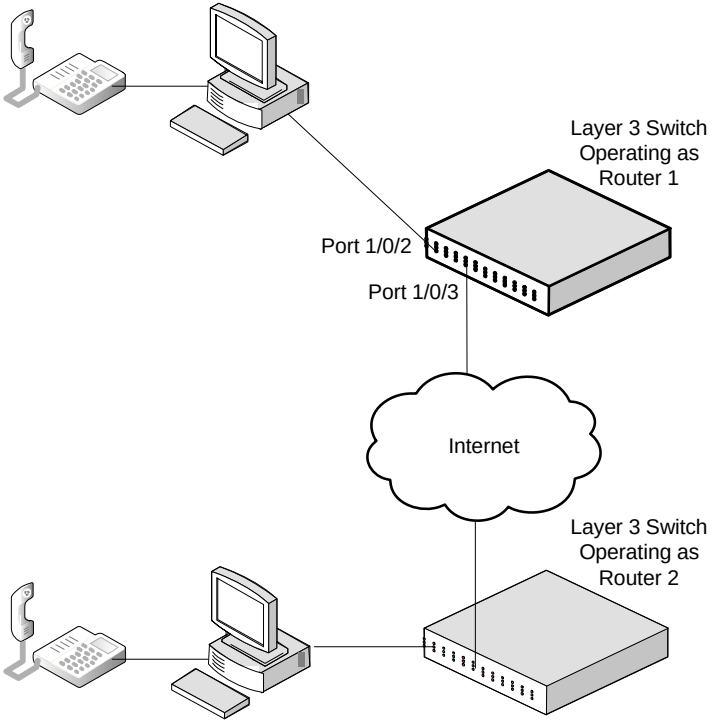
5 Set the CoS queue configuration for the (presumed) egress 10-Gigabit Ethernet interface 1/0/1 such that each of queues 1, 2, 3 and 4 get a minimum guaranteed bandwidth of 25%. All queues for this interface use weighted round robin scheduling by default. The DiffServ inbound policy designates that these queues are to be used for the departmental traffic through the assign-queue attribute. It is presumed that the switch will forward this traffic to 10-Gigabit Ethernet interface 1/0/1 based on a normal destination address lookup for internet traffic.

```
console(config)#interface tengigabitethernet 1/0/5
console(config-if-Te1/0/5)#cos-queue min-
bandwidth 0 25 25 25 25 0 0
console(config-if-Te1/0/5)#exit
console(config)#exit
```

DiffServ for VoIP

One of the most valuable uses of DiffServ is to support Voice over IP (VoIP). VoIP traffic is inherently time-sensitive: for a network to provide acceptable service, a guaranteed transmission rate is vital. This example shows one way to provide the necessary quality of service: how to set up a class for UDP traffic, have that traffic marked on the inbound side, and then expedite the traffic on the outbound side. The configuration script is for Router 1 in the accompanying diagram: a similar script should be applied to Router 2.

Figure 40-18. DiffServ VoIP Example Network Diagram



The following commands show how to configure the DiffServ example depicted in Figure 40-18.

1. Set queue 6 on all ports to use strict priority mode. This queue shall be used for all VoIP packets. Activate DiffServ for the switch.

```
console#config
console(config)#cos-queue strict 6
console(config)#diffserv
```

2. Create a DiffServ classifier named *class_voip* and define a single match criterion to detect UDP packets. The class type *match-all* indicates that all match criteria defined for the class must be satisfied in order for a packet to be considered a match.

```
console(config)#class-map match-all class_voip
console(config-classmap)#match protocol udp
console(config-classmap)#exit
```

3. Create a second DiffServ classifier named class_ef and define a single match criterion to detect a DiffServ code point (DSCP) of EF (expedited forwarding). This handles incoming traffic that was previously marked as expedited elsewhere in the network.

```
console(config)#class-map match-all class_ef
console(config-classmap)#match ip dscp ef
console(config-classmap)#exit
```

4. Create a DiffServ policy for inbound traffic named pol_voip, then add the previously created classes 'class_ef' and 'class_voip' as instances within this policy. This policy handles incoming packets already marked with a DSCP value of EF (per class_ef definition), or marks UDP packets (per the class_voip definition) with a DSCP value of EF. In each case, the matching packets are assigned internally to use queue 6 of the egress port to which they are forwarded.

```
console(config)#policy-map pol_voip in
console(config-policy-map)#class class_ef
console(config-policy-classmap)#assign-queue 6
console(config-policy-classmap)#exit

console(config-policy-map)#class class_voip
console(config-policy-classmap)#mark ip-dscp ef
console(config-policy-classmap)#assign-queue 6
```

```
console(config-policy-classmap)#exit
console(config-policy-map)#exit
```

5 Attach the defined policy to an inbound service interface.

```
console(config)#interface tengigabitethernet 1/0/1
console(config-if-Te1/0/1)#service-policy in
pol_voip
console(config-if-Te1/0/1)#exit
console(config)#exit
```

41

Configuring Class-of-Service

This chapter describes how to configure the Class-of-Service (CoS) feature. The CoS queueing feature lets you directly configure certain aspects of switch queueing. This provides the desired QoS behavior for different types of network traffic when the complexities of DiffServ are not required. The priority of a packet arriving at an interface can be used to steer the packet to the appropriate outbound CoS queue through a mapping table. CoS queue characteristics that affect queue mapping, such as minimum guaranteed bandwidth, transmission rate shaping, etc., are user-configurable at the queue (or port) level.

The topics covered in this chapter include:

- CoS Overview
- Default CoS Values
- Configuring CoS (Web)
- Configuring CoS (CLI)
- CoS Configuration Example

CoS Overview

The CoS feature lets you give preferential treatment to certain types of traffic over others. To set up this preferential treatment, you can configure the ingress ports, the egress ports, and individual queues on the egress ports to provide customization that suits your environment.

The level of service is determined by the egress port queue to which the traffic is assigned. When traffic is queued for transmission, the rate at which it is serviced depends on how the queue is configured and possibly the amount of traffic present in other queues for that port.

Some traffic is classified for service (i.e., packet marking) before it arrives at the switch. If you decide to use these classifications, you can map this traffic to egress queues by setting up a CoS Mapping table.

Each ingress port on the switch has a default priority value (set by configuring VLAN Port Priority in the Switching sub-menu) that determines the egress queue its traffic gets forwarded to. Packets that arrive without a VLAN user priority, or packets from ports you've identified as "untrusted," get forwarded according to this default.

What Are Trusted and Untrusted Port Modes?

Ports can be configured in "trusted" mode or "untrusted" mode with respect to ingress traffic.

Ports in Trusted Mode

When a port is configured in trusted mode, the system accepts at face value a priority designation encoded within packets arriving on the port. You can configure ports to trust priority designations based on one of the following fields in the packet header:

- 802.1 Priority: values 0–7
- IP DSCP: values 0–63

A mapping table associates the designated field values in the incoming packet headers with a traffic class priority (actually a CoS traffic queue).

Ports in Untrusted Mode

If you configure an ingress port in untrusted mode, the system ignores any priority designations encoded in incoming packets, and instead sends the packets to a traffic queue based on the ingress port's default priority.

How Is Traffic Shaping Used on Egress Traffic?

For unit/slot/port interfaces, you can specify a traffic shaping rate for the port (in Kbps) for egress traffic. The traffic shaping rate specifies an upper limit of the transmission bandwidth used.

How Are Traffic Queues Defined?

For each queue, you can specify:

- Minimum bandwidth guarantee—A percentage of the port's maximum negotiated bandwidth reserved for the queue. Unused reserved bandwidth can be utilized by lower-priority queues.
- Scheduler type—strict/weighted:
  - Strict priority scheduling gives an absolute priority based on CoS queue number, with traffic in the highest numbered queue sent first, then the next lowest numbered queue, and so on. Weighted queues are serviced after all strict priority queues have been serviced.
  - Weighted scheduling selects packets for transmission with a fixed weighting equal to the CoS queue number plus one. For example, if CoS queues 0, 1, and 2 have an equal offered load toward a congested output port, CoS queue 2 will receive 3/6 of the bandwidth, CoS queue 1 will receive 2/6 of the bandwidth, and CoS queue 0 will receive 1/6 of the bandwidth.

The minimum bandwidth setting can be used to override the strict priority and weighted settings. The highest numbered strict priority queue will receive no more bandwidth than 100 percent minus the sum of the minimum bandwidths percentages assigned to the other queues. If used, it is recommended that minimum bandwidth percentages only be high enough to ensure a minimum level of service for any queue; i.e., the sum of the minimum bandwidth percentages is a fraction of 100%. This ensures that the system can respond to bursts in traffic. Setting the minimum bandwidth percentages such that they sum to 100% effectively sets the scheduler such that no queue can exceed its minimum bandwidth percentage when congestion occurs.

Which Queue Management Methods Are Supported?

The switch supports the following methods, configurable per-interface-queue, for determining which packets are dropped when the queue is full:

- Taildrop—Any packet forwarded to a full queue is dropped regardless of its importance.

- Weighted Random Early Detection (WRED)—Drops packets selectively based their drop precedence level. For each of four drop precedence levels on each WRED-enabled interface queue, you can configure the following parameters:
  - Minimum Threshold: A percentage of the total queue size below which no packets of the selected drop precedence level are dropped.
  - Maximum Threshold: A percentage of the total queue size above which all packets of the selected drop precedence level are dropped.
  - Drop Probability: When the queue depth is between the minimum and maximum thresholds, this value provides a scaling factor for increasing the number of packets of the selected drop precedence level that are dropped as the queue depth increases.

CoS Queue Usage

CoS queue 7 is reserved by the system and is not assignable. It is generally recommended that the administrator utilize CoS queues 0 to 3, as CoS queues 4-6 may be used by the system for other types of system traffic, for example routing protocol PDU handling.

Default CoS Values

Table 41-1 shows the global default values for CoS.

Table 41-1. CoS Global Defaults

Parameter	Default Value	
Trust Mode	802.1p	
802.1p CoS value to queue mapping	**802.1p CoS**	**Queue**
	0, 3	1
	1, 2	0
	4, 5	2
	6, 7	3

Table 41-1. CoS Global Defaults

Parameter	Default Value	
IP DSCP value to queue mapping	**IP DSCP**	**Queue**
	0–7, 24–31	1
	8–23	0
	32–47	2
	48–63	3
Interface Shaping Rate	0 Kbps	
Minimum Bandwidth	0%	
Scheduler Type	Weighted	
Queue Management Type	Taildrop	
Drop Precedence Level	1	
WRED Decay Exponent	9	
WRED Minimum Threshold	40	
WRED Maximum Threshold	100	
WRED Drop Probability Scale	10	

Configuring CoS (Web)

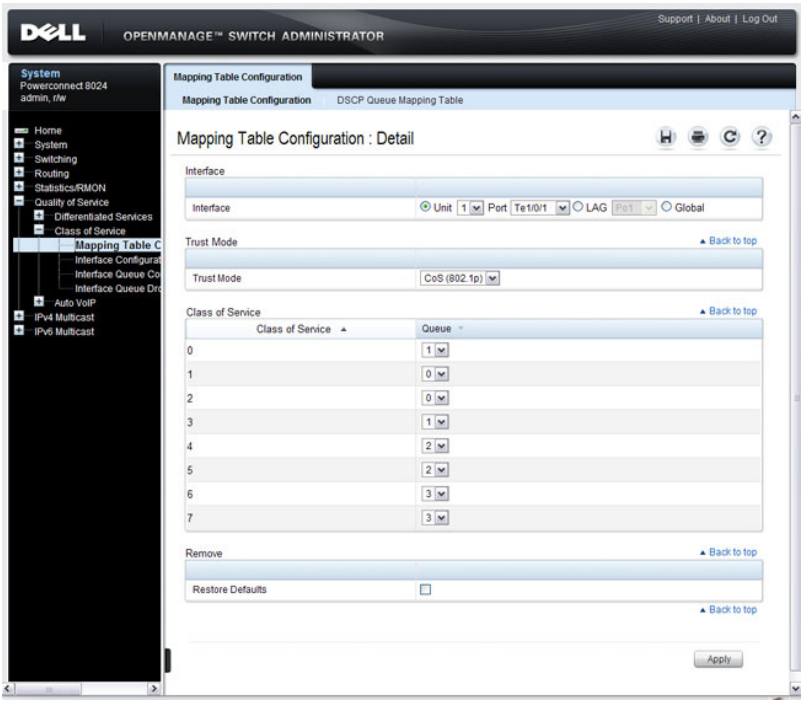
This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring CoS features on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

Mapping Table Configuration

Use the **Mapping Table Configuration** page to define how class of service is assigned to a packet.

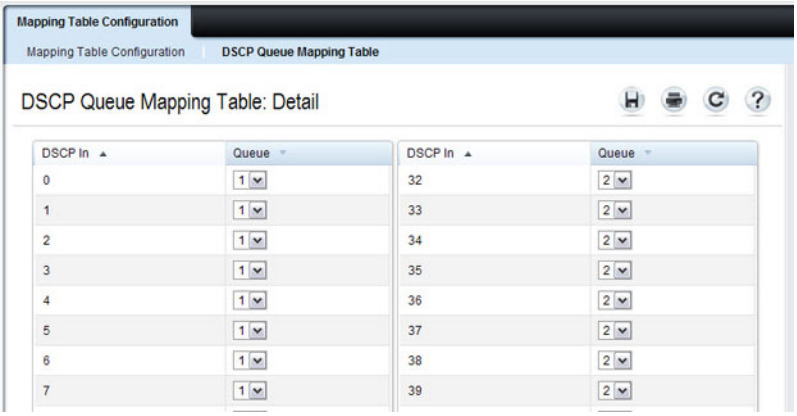
To display the page, click **Quality of Service → Class of Service → Mapping Table Configuration** in the navigation panel. CoS(802.1P) is the default mode, so this is the page that displays when **Mapping Table Configuration** is selected from the **Class of Service** menu page.

Figure 41-1. Mapping Table Configuration — CoS (802.1P)



To access the **DSCP Queue Mapping Table**, click the **DSCP Queue Mapping Table** link at the top of the page.

Figure 41-2. DSCP Queue Mapping Table



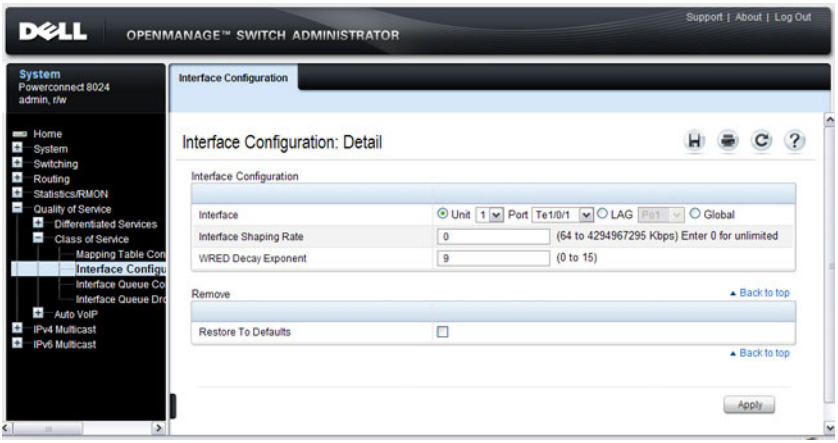
Interface Configuration

Use the **Interface Configuration** page to define the interface shaping rate for egress packets on an interface and the decay exponent for WRED queues defined on the interface.

Each interface CoS parameter can be configured globally or per-port. A global configuration change is applied to all interfaces in the system.

To display the Interface Configuration page, click **Quality of Service → Class of Service → Interface Configuration** in the navigation panel.

Figure 41-3. Interface Configuration



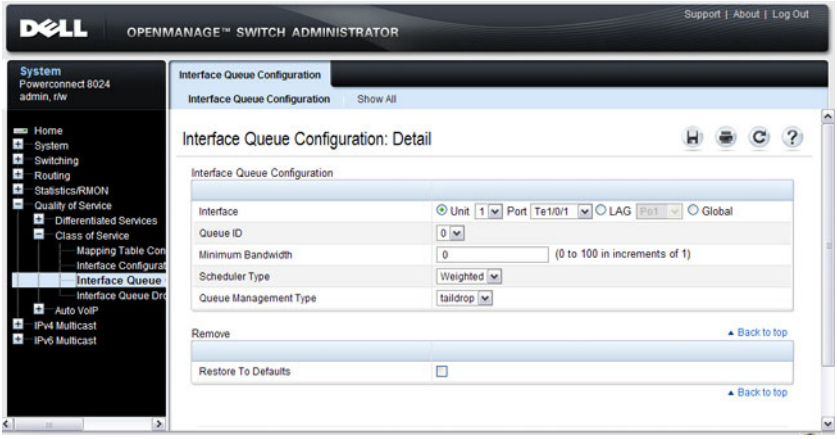
Interface Queue Configuration

Use the **Interface Queue Configuration** page to configure egress queues on interfaces. The settings you configure control the amount of bandwidth the queue uses, the scheduling method, and the queue management method.

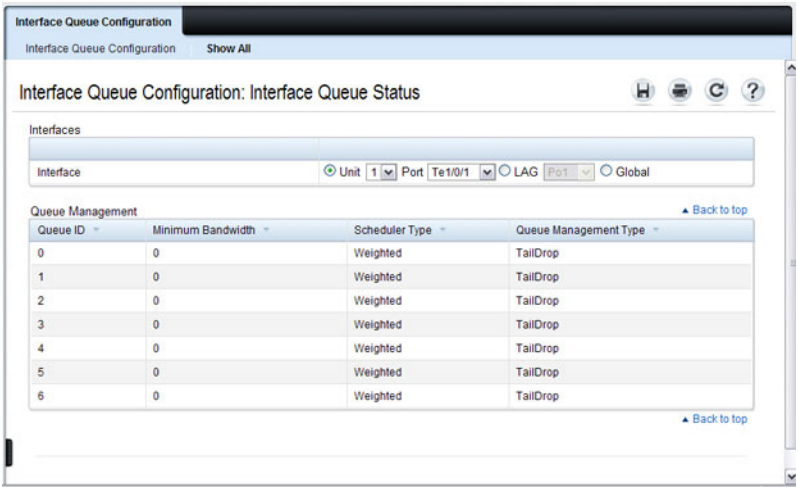
The configuration process is simplified by allowing each CoS queue parameter to be configured globally or per-port. A global configuration change is applied to the same queue ID on all ports in the system.

To display the **Interface Queue Configuration** page, click **Quality of Service → Class of Service → Interface Queue Configuration** in the navigation panel.

Figure 41-4. Interface Queue Configuration



To access the **Interface Queue Status** page, click the **Show All** link at the top of the page.



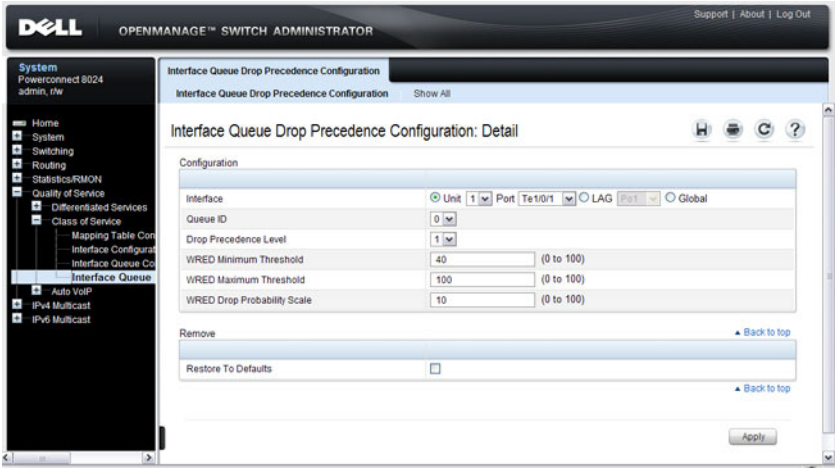
Interface Queue Drop Precedence Configuration

Use the **Interface Queue Drop Precedence Configuration** page to configure thresholds and scaling values for each of four drop precedence levels on a WRED-enabled interface queue. The settings you configure control the minimum and maximum thresholds and a drop probability scaling factor for the selected drop precedence level.

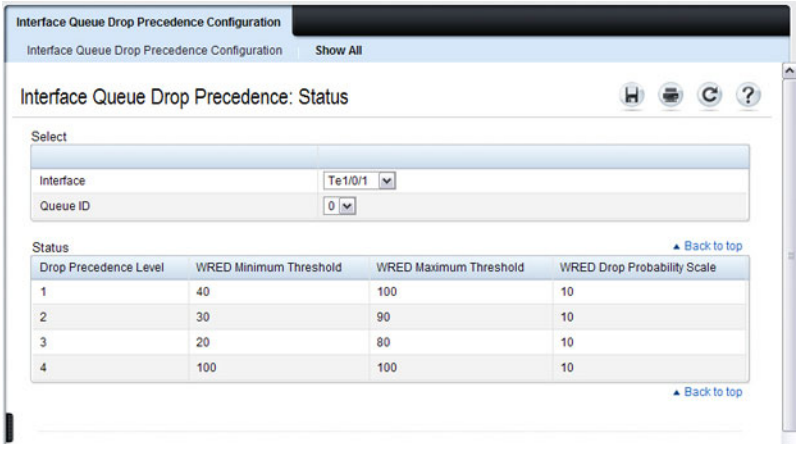
These parameters can be applied to each drop precedence level on a per-interface-queue basis, or can be set globally for the same drop precedence level and queue ID on all interfaces.

To display the **Interface Queue Drop Precedence Configuration** page, click **Quality of Service → Class of Service → Interface Queue Drop Precedence Configuration** in the navigation panel.

Figure 41-5. Interface Queue Drop Precedence Configuration



To access the **Interface Queue Drop Precedence Status** page, click the **Show All** link at the top of the page.



Configuring CoS (CLI)

This section provides information about the commands you use to configure CoS settings on the switch. For more information about the commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Mapping Table Configuration

Beginning in Privileged Exec mode, use the following commands in to configure the CoS mapping tables.

CLI Command	Description
configure	Enter Global Configuration mode.
interface *interface*	Enter Interface Configuration mode, where *interface* is replaced by **tengigabitethernet** *unit/slot/port* or **port-channel** *port-channel number.*
classofservice dotlp-mapping *priority*	Map an 802.1p priority to an internal traffic class for a switch. You can also use this command in Global Configuration mode to configure the same mappings on all interfaces.
classofservice trust {dot1p | ip-dscp | untrusted}	Set the class of service trust mode of an interface.
exit	Exit to Global Config mode.
exit	Exit to Privilege Exec mode.
show classofservice dotlp-mapping	Display the current Dot1p (802.1p) priority mapping to internal traffic classes for a specific interface.
show classofservice ip-dscp-mapping	Display the current IP DSCP mapping to internal traffic classes for a specific interface.
show classofservice trust	Display the current trust mode setting for a specific interface.

CoS Interface Configuration Commands

Beginning in Privileged Exec mode, use the following commands in to configure the traffic shaping and WRED exponent values for an interface.

CLI Command	Description
configure	Enter Global Configuration mode.
interface *interface*	Enter Interface Configuration mode, where *interface* is replaced by **tengigabitethernet** *unit/slot/port* or **port-channel** *port-channel number*.
traffic-shape *bw* **kbps**	Sets the upper limit on how much traffic can leave a port. The *bw* variable represents the shaping bandwidth value from 64 to 4294967295 kbps.
random-detect exponential-weighting-constant *exponent*	Configure the WRED decay exponent (range: 0–15) for the interface.

Interface Queue Configuration

Beginning in Privileged Exec mode, use the following commands in to configure and view CoS interface queue settings.

CLI Command	Description
configure	Enter Global Configuration mode.
interface *interface*	Enter Interface Configuration mode, where *interface* is replaced by **tengigabitethernet** *unit/slot/port*. or **port-channel** *port-channel number*.
cos-queue min-bandwidth *bw*	Specify the minimum transmission bandwidth (range: 0-100% in 1% increments) for each interface queue.
cos-queue strict *queue-id*	Activate the strict priority scheduler mode for each specified queue. The *queue-id* value ranges from 0 to 6.
cos-queue random-detect *queue-id*	Set the queue management type for the specified queue to WRED. The **no** version of this command resets the value to taildrop.
exit	Exit to Global Config mode.

CLI Command	Description
exit	Exit to Privilege Exec mode.
show interfaces cos-queue	Display the class-of-service queue configuration for a specified interface or all interfaces.

Configuring Interface Queue Drop Precedence

Beginning in Privileged Exec mode, use the following commands in to configure characteristics of the drop precedence levels and view related settings.

CLI Command	Description
configure	Enter Global Configuration mode.
interface *interface*	Enter Interface Configuration mode, where *interface* is replaced by **tengigabitethernet** *unit/slot/port* or **port-channel** *port-channel number*.
random-detect queue-parms *queue-id* [*queue-id...*] **min-thresh** *min1 min2 min3 min4* **max-thresh** *max1 max2 max3 max4* **drop-prob** *prob1 prob2 prob3 prob4*	Configure the maximum and minimum thresholds for one or more queue IDs on a WRED-enabled interface queue. You can also use this command in Global Configuration mode to configure the same parameters for one or more queues on all interfaces.
exit	Exit to Global Config mode.
exit	Exit to Privilege Exec mode.
show interfaces random-detect	Display WRED parameters for an interface or all interfaces.

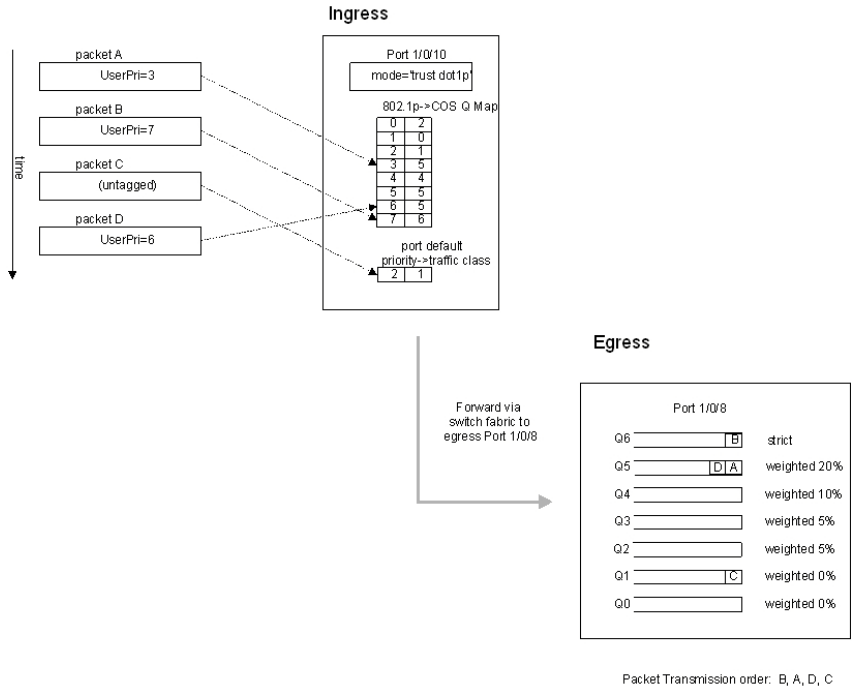
CoS Configuration Example

Figure 41-6 illustrates the network operation as it relates to CoS mapping and queue configuration.

Four packets arrive at the ingress port te1/0/10 in the order A, B, C, and D. port te1/0/10 is configured to trust the 802.1p field of the packet, which serves to direct packets A, B, and D to their respective queues on the egress port. These three packets utilize the 802.1p to CoS Mapping Table for port te1/0/10.

In this example, the 802.1p user priority 3 is configured to send the packet to queue 5 instead of the default queue 3. Since packet C does not contain a VLAN tag, the 802.1p user priority does not exist, so port te1/0/10 relies on its default port priority (2) to direct packet C to egress queue 1.

Figure 41-6. CoS Mapping and Queue Configuration



Continuing this example, the egress port te1/0/8 is configured for strict priority on queue 6, and a weighted scheduling scheme is configured for queues 5-0. Assuming queue 5 has a higher minimum bandwidth than queue 1 (relative bandwidth values are shown as a percentage, with 0% indicating the bandwidth is not guaranteed), the queue service order is 6 followed by 5 followed by 1. Assuming each queue transmits all packets shown in the diagram, the packet transmission order as seen on the network out of port te1/0/8 is B, A, D, C. Thus, packet B, with its strict priority scheduling, is transmitted ahead of the other packets at the egress port.

The following commands configure port 10 (ingress interface) and port 8 (egress interface).

1. Configure the Trust mode for port 10.

```
console#config
console(config)#interface tengigabitethernet
1/0/10
console(config-if-Te1/0/10)#classofservice trust
dot1p
```

2. For port 10, configure the 802.1p user priority 3 to send the packet to queue 5 instead of the default queue (queue 3).

```
console(config-if-Te1/0/10)#classofservice dot1p-
mapping 3 5
```

3. For port 10, specify that untagged VLAN packets should have a default priority of 2.

```
console(config-if-Te1/0/10)#vlan priority 2
console(config-if-Te1/0/10)#exit
```

4. For port 8, the egress port, configure a weighted scheduling scheme for queues 5-0.

```
console(config)#interface tengigabitethernet 1/0/8
console(config-if-Te1/0/8)#cos-queue min-
bandwidth 0 0 5 5 10 20 40
```

5. Configure port 8 to have strict priority on queue 6:

```
console(config-if-Te1/0/8)#cos-queue strict 6
```

To configure the CoS queues for lossless traffic when transporting iSCSI traffic, set the lossless traffic class to have a one-to-one mapping with the priority value. The following example illustrates how to change the dot1p

mapping from the switch defaults to support lossless transport of frames on CoS queue 4, with a 50% minimum bandwidth guarantee. Lossless traffic classes generally use the default WRED queue discipline, as opposed to strict priority, to avoid starving other traffic.

```
classofservice dot1p-mapping 4 4
cos-queue min-bandwidth 0 0 0 0 50 0 0
```

42

Configuring Auto VoIP

Voice over Internet Protocol (VoIP) allows you to make telephone calls using a computer network over a data network like the Internet. With the increased prominence of delay-sensitive applications (voice, video, and other multimedia applications) deployed in networks today, proper QoS configuration will ensure high-quality application performance. The Auto VoIP feature is intended to provide an easy classification mechanism for voice packets so that they can be prioritized above data packets in order to provide better QoS. Because Auto VoIP is limited to 16 sessions, Voice VLAN is the preferred solution for enterprises wishing to deploy a large scale voice service.

The topics covered in this chapter include:

- Auto VoIP Overview
- Default Auto VoIP Values
- Configuring Auto VoIP (Web)
- Configuring Auto VoIP (CLI)

Auto VoIP Overview

The Auto VoIP feature explicitly matches VoIP streams in Ethernet switches and provides them with a better class of service than ordinary traffic. If you enable the Auto VoIP feature on an interface, the interface scans incoming traffic for the following call-control protocols:

- Session Initiation Protocol (SIP)
- H.323
- Skinny Client Control Protocol (SCCP)

When a call-control protocol is detected the switch assigns the traffic in that session to the highest CoS queue, which is generally used for time-sensitive traffic.

Auto-VoIP is limited to 16 sessions and makes use of the switch CPU to classify traffic. It is preferable to use the Voice VLAN feature in larger enterprise environments as it uses the switching silicon to classify voice traffic onto a VLAN.

How Does Auto-VoIP Use ACLs?

Auto-VoIP borrows ACL lists from the global system pool. ACL lists allocated by Auto-VoIP reduce the total number of ACLs available for use by the network operator. Enabling Auto-VoIP uses one ACL list to monitor for VoIP sessions. Each monitored VoIP session utilizes two rules from an additional ACL list. This means that the maximum number of ACL lists allocated by Auto-VoIP is two.

Default Auto VoIP Values

Table 42-1 shows the global default value for Auto VoIP.

Table 42-1. Auto VoIP Global Defaults

Parameter	Default Value
Auto VoIP	Disabled

Configuring Auto VoIP (Web)

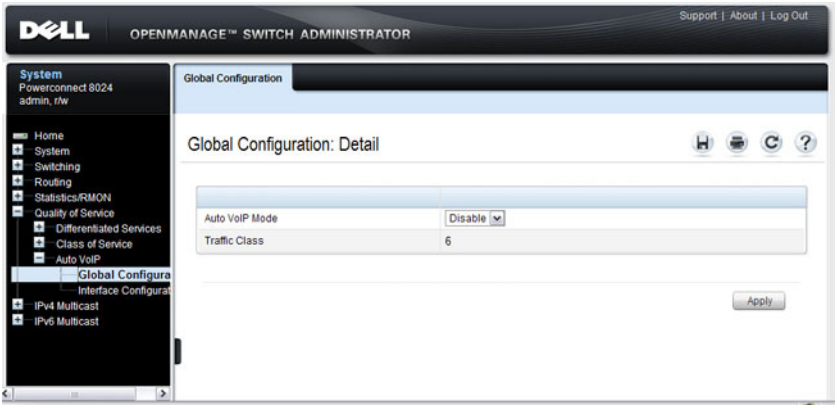
This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring Auto VoIP features on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click at the top of the page.

Auto VoIP Global Configuration

Use the **Global Configuration** page to enable or disable Auto VoIP on all interfaces.

To display the Auto VoIP **Global Configuration** page, click **Quality of Service → Auto VoIP → Global Configuration** in the navigation menu.

Figure 42-1. Auto VoIP Global Configuration

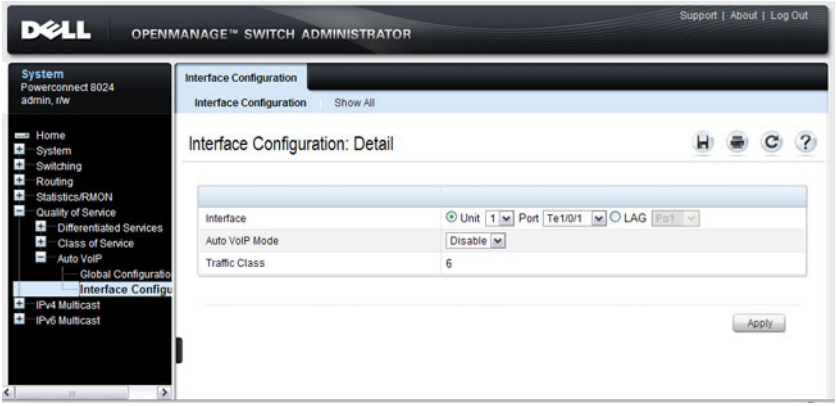


Auto VoIP Interface Configuration

Use the **Interface Configuration** page to enable or disable Auto VoIP on a particular interface.

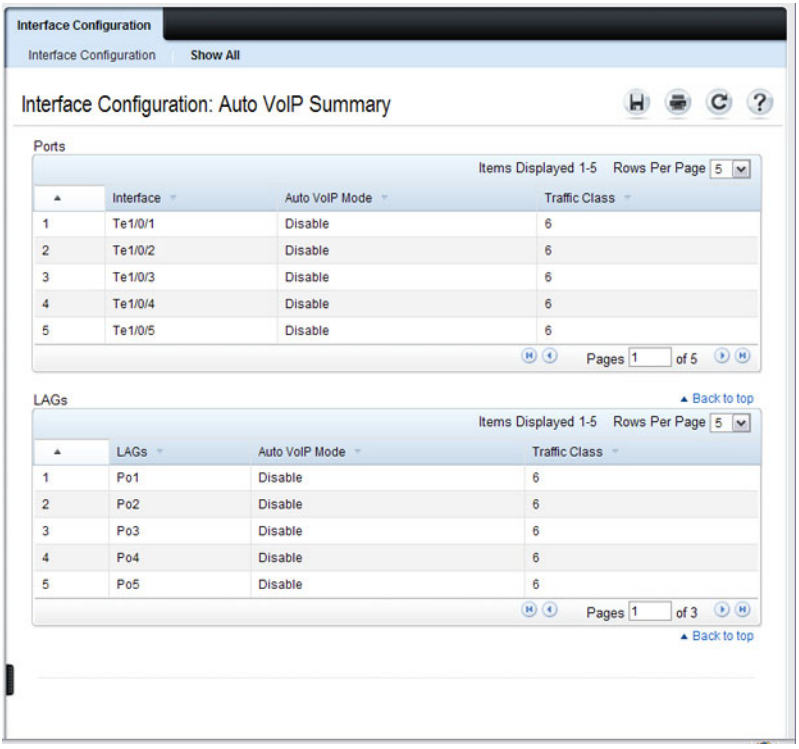
To display the **Interface Configuration** page, click **Quality of Service → Auto VoIP → Interface Configuration** in the navigation menu.

Figure 42-2. Auto VoIP Interface Configuration



To display summary Auto VoIP configuration information for all interfaces, click the **Show All** link at the top of the page.

Figure 42-3. Auto VoIP



Configuring Auto VoIP (CLI)

This section provides information about the commands you use to configure Auto VoIP settings on the switch. For more information about the commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Mapping Table Configuration

Beginning in Privileged Exec mode, use the following commands in to enable Auto VoIP and view its configuration.

CLI Command	Description
configure	Enter Global Configuration mode.
switchport voice detect auto	Enable the VoIP Profile on all the interfaces of the switch. You can also enter Interface Configuration mode and use the same command to enable it on a specific interface.
exit	Exit to Global Configuration Exec mode.
exit	Exit to Privilege Exec mode.
show switchport voice	Show the status of Auto-VoIP on all interfaces or on an interface, if one is specified.

43

Managing IPv4 and IPv6 Multicast

This chapter describes how to configure and monitor layer 3 multicast features for IPv4 and IPv6, including global IP and IPv6 multicast features as well as multicast protocols, including IGMP, DVMRP, and PIM for IPv4 and MLD and PIM for IPv6.

The topics covered in this chapter include:

- L3 Multicast Overview
- Default L3 Multicast Values
- Configuring General IPv4 Multicast Features (Web)
- Configuring IPv6 Multicast Features (Web)
- Configuring IGMP and IGMP Proxy (Web)
- Configuring MLD and MLD Proxy (Web)
- Configuring PIM for IPv4 and IPv6 (Web)
- Configuring DVMRP (Web)
- Configuring L3 Multicast Features (CLI)
- L3 Multicast Configuration Examples

L3 Multicast Overview

IP Multicasting enables a network host (or multiple hosts) to send an IP datagram to multiple destinations simultaneously. The initiating host sends each multicast datagram only once to a destination multicast group address, and multicast routers forward the datagram only to hosts who are members of the multicast group. Multicast enables efficient use of network bandwidth because each multicast datagram needs to be transmitted only once on each network link, regardless of the number of destination hosts. Multicasting contrasts with IP unicasting, which sends a separate datagram to each recipient host. The IP routing protocols can route multicast traffic, but the IP multicast protocols handle the multicast traffic more efficiently with better use of network bandwidth.

Applications that often send multicast traffic include video or audio conferencing, Whiteboard tools, stock distribution tickers, and IP-based television (IP/TV).

What Is IP Multicast Traffic?

IP multicast traffic is traffic that is destined to a host group. Host groups are identified by class D IP addresses, which range from 224.0.0.0 to 239.255.255.255. When a packet with a broadcast or multicast destination IP address is received, the switch will forward a copy into each of the remaining network segments in accordance with the IEEE MAC Bridge standard. Eventually, the packet is made accessible to all nodes connected to the network.

This approach works well for broadcast packets that are intended to be seen or processed by all connected nodes. In the case of multicast packets, however, this approach could lead to less efficient use of network bandwidth, particularly when the packet is intended for only a small number of nodes. Packets will be flooded into network segments where no node has any interest in receiving the packet. The L3 multicast features on the switch help to ensure that only the hosts in the multicast group receive the multicast traffic for that group.

Multicast applications send one copy of a packet, and address it to a group of receivers (Multicast Group Address) rather than to a single receiver (unicast address). Multicast depends on the network to forward the packets to only those networks and hosts that need to receive them.

What Multicast Protocols Does the Switch Support?

Multicast protocols are used to deliver Multicast packets from one source to multiple receivers. Table 43-1 summarizes the multicast protocols that the switch supports.

Table 43-1. Multicast Protocol Support Summary

Protocol	IPv4 or IPv6	For Communication Between
IGMP	IPv4	Host-to-L3 switch/router
MLD	IPv6	Host-to-L3 switch (router)
PIM-SM	IPv4 or IPv6	L3-switch/router-to-L3 switch/router
PIM-DM	IPv4 or IPv6	L3-switch/router-to-L3 switch/router
DVMRP	IPv4	L3-switch/router-to-L3 switch/router

What Are the Multicast Protocol Roles?

Hosts must have a way to identify their interest in joining any particular multicast group, and routers must have a way to collect and maintain group memberships. These functions are handled by the IGMP protocol in IPv4. In IPv6, multicast routers use the Multicast Listener Discover (MLD) protocol to maintain group membership information.

Multicast routers must also be able to construct a multicast distribution tree that enables forwarding multicast datagrams only on the links that are required to reach a destination group member. Protocols such as DVMRP, and PIM handle this function.

IGMP and MLD are multicast group discovery protocols that are used between the clients and the local multicast router. PIM-SM, PIM-DM, and DVMRP are multicast routing protocols that are used across different subnets, usually between the local multicast router and remote multicast router.

When Is L3 Multicast Required on the Switch?

You use the IPv4/IPv6 multicast feature on the PowerConnect 8000-series and 8100-series switches to route multicast traffic between VLANs on the switch. If all hosts connected to the switch are on the same subnet, there is no need to configure the IP/IPv6 multicast feature. If the switch does not handle L3 routing, you can use IGMP snooping or MLD snooping to manage port-based multicast group membership. For more information, see "What Is IGMP Snooping?" on page 683 and "What Is MLD Snooping?" on page 683. If the local network does not have a multicast router, you can configure the switch to act as the IGMP querier. For more information, see "IGMP Snooping Querier" on page 683.

If the switch is configured as a L3 switch and handles inter-VLAN routing through static routes, OSPF, or RIP, and multicast traffic is transmitted within the network, enabling and configuring L3 multicast routing on the switch is recommended.

Determining Which Multicast Protocols to Enable

IGMP is recommended on any switch that participates in IPv4 multicasting. MLD is recommended on any switch that participates in IPv6 multicasting. PIM-DM, PIM-SM, and DVMRP are multicast routing protocols that help determine the best route for IP (PIM and DVMRP) and IPv6 (PIM) multicast traffic. For more information about when to use PIM-DM, see "Using PIM-DM as the Multicast Routing Protocol" on page 1164. For more information about when to use PIM-SM, see "Using PIM-SM as the Multicast Routing Protocol" on page 1163. For more information about when to configure DVMRP, see "Using DVMRP as the Multicast Routing Protocol" on page 1166.

What Is the Multicast Routing Table?

Multicast capable/enabled routers forward multicast packets based on the routes in the Multicast Routing Information Base (MRIB). These routes are created in the MRIB during the process of building multicast distribution trees by the Multicast Protocols running on the router. Different IP Multicast routing protocols use different techniques to construct these multicast distribution trees.

What Is Multicast Tunneling?

If Multicast traffic is to be routed through a part of a network that does not support multicasting (routers which are not multicast capable) then the multicast packets are encapsulated in an IP datagram and sent as a unicast packet. When the multicast router at the remote end of the tunnel receives the packet, the router strips off the IP encapsulation and forwards the packet as an IP Multicast packet. This process of encapsulating multicast packets in IP is called tunneling.

What Is IGMP?

The Internet Group Management Protocol (IGMP) is used by IPv4 systems (hosts, L3 switches, and routers) to report their IP multicast group memberships to any neighboring multicast routers. The PowerConnect 8000-series and 8100-series switches performs the multicast router role of the IGMP protocol, which means it collects the membership information needed by the active multicast routing protocol.

The PowerConnect 8000-series and 8100-series switches supports IGMP Version 3. Version 3 adds support for source filtering, which is the ability for a system to report interest in receiving packets only from specific source addresses, as required to support Source-Specific Multicast [SSM], or from all but specific source addresses, sent to a particular multicast address. Version 3 is designed to be interoperable with Versions 1 and 2.

Understaning IGMP Proxy

IGMP proxy enables a multicast router to learn multicast group membership information and forward multicast packets based upon the group membership information. The IGMP Proxy is capable of functioning only in certain topologies that do not require Multicast Routing Protocols (i.e., DVMRP, PIM-DM, and PIM-SM) and have a tree-like topology, as there is no support for features like reverse path forwarding (RPF) to correct packet route loops.

The proxy contains many downstream interfaces and a unique upstream interface explicitly configured. It performs the host side of the IGMP protocol on its upstream interface and the router side of the IGMP protocol on its downstream interfaces.

The IGMP proxy offers a mechanism for multicast forwarding based only on IGMP membership information. The router must decide about forwarding packets on each of its interfaces based on the IGMP membership information. The proxy creates the forwarding entries based on the membership information and adds it to the multicast forwarding cache (MFC) in order not to make the forwarding decision for subsequent multicast packets with same combination of source and group.

What Is MLD?

Multicast Listener Discovery (MLD) protocol enables IPv6 routers to discover the presence of multicast listeners, the hosts that wish to receive the multicast data packets, on its directly-attached interfaces. The protocol specifically discovers which multicast addresses are of interest to its neighboring nodes and provides this information to the active multicast routing protocol that makes decisions on the flow of multicast data packets.

The Multicast router sends General Queries periodically to request multicast address listeners information from systems on an attached network. These queries are used to build and refresh the multicast address listener state on attached networks. Multicast listeners respond to these queries by reporting their multicast addresses listener state and their desired set of sources with Current-State Multicast address Records in the MLD2 Membership Reports. The Multicast router also processes unsolicited Filter-Mode-Change records and Source-List-Change Records from systems that want to indicate interest in receiving or not receiving traffic from particular sources.

The PowerConnect implementation of MLD v2 supports the multicast router portion of the protocol (i.e., not the listener portion). It is backward-compatible with MLD v1.

What Is PIM?

The Protocol Independent Multicast protocol is a simple, protocol-independent multicast routing protocol. PIM uses an existing unicast routing table and a Join/Prune/Graft mechanism to build a tree. PIM PowerConnect 8000-series and 8100-series switches support two types of PIM: sparse mode (PIM-SM) and dense mode (PIM-DM).

PIM-SM is most effective in networks with a sparse population of multicast receivers. In contrast, PIM-DM is most effective in networks with densely populated multicast receivers. In other words, PIM-DM can be used if the majority of network hosts request to receive a multicast stream, while PIM-SM might be a better choice in networks in which a small percentage of network hosts, located throughout the network, wish to receive the multicast stream.

Using PIM-SM as the Multicast Routing Protocol

PIM-SM is used to efficiently route multicast traffic to multicast groups that may span wide area networks where bandwidth is a constraint.

PIM-SM uses shared trees by default and implements source-based trees for efficiency; it assumes that no hosts want the multicast traffic unless they specifically ask for it. It creates a shared distribution tree centered on a defined rendezvous point (RP) from which source traffic is relayed to the receivers. Senders first send the multicast data to the RP, which in turn sends the data down the shared tree to the receivers.

Shared trees centered on an RP do not necessarily provide the shortest, most optimal path. In such cases, PIM-SM provides a means to switch to more efficient source-specific trees. A data threshold rate is configured to determine when to switch from shared-tree to source-tree.

PIM-SM uses a Bootstrap Router (BSR), which advertises information to other multicast routers about the RP. In a given network, a set of routers can be administratively enabled as candidate bootstrap routers. If it is not apparent which router should be the BSR, the candidates flood the domain with advertisements. The router with the highest priority is elected. If all the priorities are equal, then the candidate with the highest IP address becomes the BSR.

Only one RP address can be used at a time within a PIM domain. You can configure a static RP on the switch. However, if the PIM domain uses the BSR to dynamically learn the RP, configuring a static RP is not required. By default the RP advertised by the BSR is used, but you can specify that the static RP to override any dynamically learned RP from the BSR.

NOTE: Dell recommends configuring a first-hop router from the multicast video source as the RP.

If an interface on a switch configured with PIM-SM neighbors another PIM-SM domain, the PIM BSR messages should not flood into the neighboring PIM domain because the neighbor domain might not share the same set of RPs, candidate RPs, BSR, and candidate BSRs. The switch software allows you to configure an interface that borders the PIM boundary prevent transmission (sending and receiving) of PIM BSR messages. PIM-SM is defined in RFC 4601.

Using PIM-DM as the Multicast Routing Protocol

Unlike PIM-SM, PIM-DM creates source-based shortest-path distribution trees that make use of reverse-path forwarding (RPF). PIM-DM assumes that when a sender starts sending data, all downstream routers and hosts want to receive a multicast datagram. PIM-DM initially floods multicast traffic throughout the network. Routers that do not have any downstream neighbors prune back the unwanted traffic. In addition to PRUNE messages, PIM-DM makes use of graft and assert messages. Graft messages are used whenever a new host wants to join the group. Assert messages are used to shutoff duplicate flows on the same multi-access network.

There are two versions of PIM-DM. Version 2 does not use the IGMP message; instead, it uses a message that is encapsulated in IP package, with protocol number 103. In Version 2, a Hello message is introduced in place of a query message.

PIM-DM is appropriate for:

- Densely distributed receivers
- Few senders-to-many receivers (due to frequent flooding)
- High volume of multicast traffic
- Constant stream of traffic

To minimize the repeated flooding of datagrams and subsequent pruning associated with a particular source-group (S,G) pair, PIM-DM uses a State Refresh message. This message is sent by the router(s) directly connected to the source and is propagated throughout the network. When received by a router on its RPF interface, the State Refresh message causes an existing prune state to be refreshed. State Refresh messages are generated periodically by the router directly attached to the source.

What Is DVMRP?

DVMRP is an interior gateway protocol that is suitable for routing multicast traffic within an autonomous system (AS). DVMRP should not be used between different autonomous systems due to limitations with hop count and scalability.

NOTE: In addition to DVMRP, the switch supports the Protocol-Independent Multicast (PIM) sparse-mode (PIM-SM) and dense-mode (PIM-SM) routing protocol. Only one multicast routing protocol can be operational on the switch at any time. If you enable DVMRP, PIM must be disabled. Similarly, if PIM is enabled, DVMRP must be disabled.

DVMRP exchanges probe packets with all its DVMRP-enabled routers, it establishes two-way neighboring relationships, and it builds a neighbor table. DVMRP exchanges report packets and creates a unicast topology table, with which it builds the multicast routing table. This table is used to route the multicast packets. Since every DVMRP router uses the same unicast routing protocol, routing loops are avoided.

Understanding DVMRP Multicast Packet Routing

DVMRP is based on RIP; it forwards multicast datagrams to other routers in the AS and constructs a forwarding table based on information it learns in response. More specifically, it uses this sequence.

- A new multicast packet is forwarded to the entire multicast network, with respect to the time-to-live (TTL) of the packet.
- The TTL restricts the area to be flooded by the message.
- All routers that do not have members on directly-attached subnetworks send back *Prune messages* to the upstream router.

- The branches that transmit a prune message are deleted from the delivery tree.
- The delivery tree which is spanning to all the members in the multicast group, is constructed in the form of a DVMRP forwarding table.

Using DVMRP as the Multicast Routing Protocol

DVMRP is used to communicate multicast information between L3 switches or routers. If a PowerConnect 8000-series and 8100-series switches handles inter-VLAN routing for IP traffic, including IP multicast traffic, multicast routing might be required on the switch.

DVMRP is best suited for small networks where the majority of hosts request a given multicast traffic stream. DVMRP is similar to PIM-DM in that it floods multicast packets throughout the network and prunes branches where the multicast traffic is not desired. DVMRP was developed before PIM-DM, and it has several limitations that do not exist with PIM-DM.

You might use DVMRP as the multicast routing protocol if it has already been widely deployed within the network.

Microsoft Network Load Balancing

PowerConnect switches support Microsoft Network Load Balancing (NLB) in unicast mode only. When using Microsoft NLB, ensure that the Cluster Operation Mode is configured to the default value of Unicast.

Default L3 Multicast Values

IP and IPv6 multicast is disabled by default. Table 43-2 shows the default values for L3 multicast and the multicast protocols.

Table 43-2. L3 Multicast Defaults

Parameter	Default Value
IPv4 Multicast Defaults	
L3 Multicast Admin Mode	Disabled
Maximum Multicast Routing Table Entries	2048
Static Multicast Routes	None configured
Interface TTL Threshold	1
IGMP Defaults	
IGMP Admin Mode	Disabled globally and on all interfaces
IGMP Version	v3
IGMP Robustness	2
IGMP Query Interval	125 seconds
IGMP Query Max Response Time	100 seconds
IGMP Startup Query Interval	31 seconds
IGMP Startup Query Count	2
IGMP Last Member Query Interval	1 second
IGMP Last Member Query Count	2
IGMP Proxy Interface Mode	Disabled
IGMP Proxy Unsolicited Report Interval	1 second
MLD Defaults	
MLD Admin Mode	Disabled globally and on all interfaces
MLD Version	v2
MLD Query Interval	125 seconds
MLD Query Max Response Time	10,000 milliseconds

Table 43-2. L3 Multicast Defaults (Continued)

Parameter	Default Value
MLD Last Member Query Interval	1000 milliseconds
MLD Last Member Query Count	2
MLD Proxy Interface Mode	Disabled
MLD Proxy Unsolicited Report Interval	1 second
PIM Defaults	
PIM Protocol	Disabled globally and on all interfaces
PIM-SM Data Threshold Rate	0 Kpbs
PIM-SM Register Threshold Rate	0 Kbps
PIM Hello Interval	30 seconds (when enabled on an interface)
PIM-SM Join/Prune Interval	60 seconds (when enabled on an interface)
PIM-SM BSR Border	Disabled
PIM-SM DR Priority	1 (when enabled on an interface)
PIM Candidate Rendezvous Points (RPs)	None configured
PIM Static RP	None configured
PIM Source-Specific Multicast (SSM) Range	None configured. Default SSM group address is 232.0.0.0/8 for IPv4 multicast and ff3x::/32 for IPv6 multicast.
PIM BSR Candidate Hash Mask Length	30 (IPv4) 126 (IPv6)
PIM BSR Candidate Priority	0
DVMRP Defaults	
DVMRP Admin Mode	Disabled globally and on all interfaces
DVMRP Version	3
DVMRP Interface Metric	1

Configuring General IPv4 Multicast Features (Web)

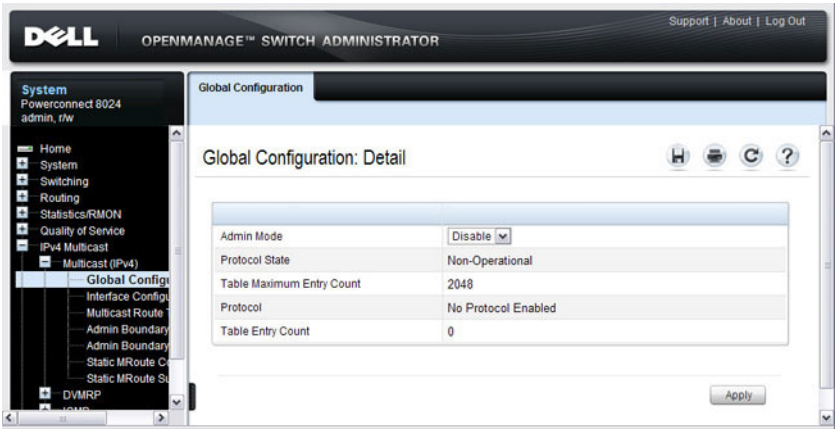
This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring the L3 multicast features that are not protocol-specific on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

Multicast Global Configuration

Use the **Global Configuration** page to configure the administrative status of Multicast Forwarding in the router, and to display global multicast parameters.

To display the page, click **IPv4 Multicast → Multicast → Global Configuration** in the navigation panel.

Figure 43-1. Multicast Global Configuration

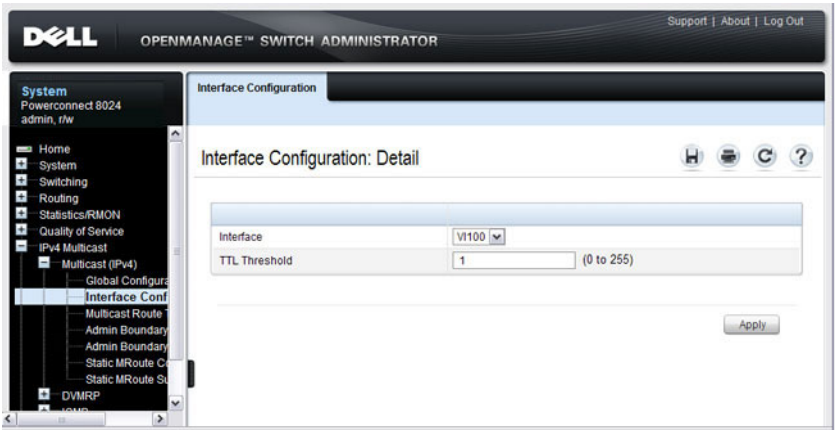


Multicast Interface Configuration

Use the **Interface Configuration** page to configure the TTL threshold of a multicast interface. At least one VLAN routing interface must be configured on the switch before fields display on this page.

To display the page, click **IPv4 Multicast → Multicast → Interface Configuration** in the navigation panel.

Figure 43-2. Multicast Interface Configuration

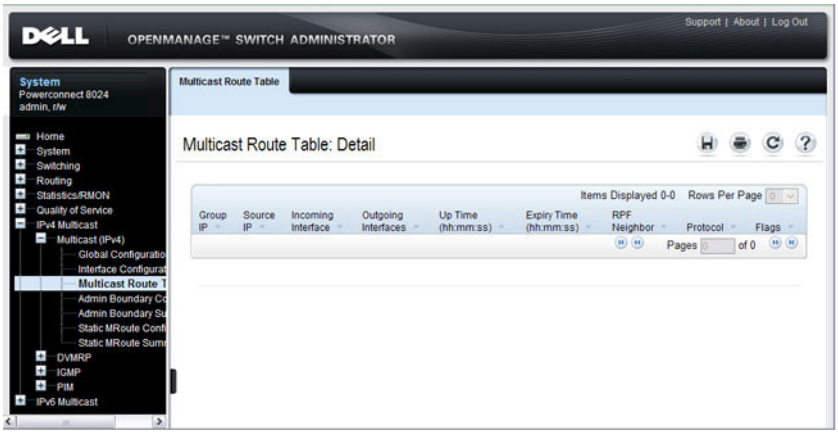


Multicast Route Table

Use the **Route Table** page to view information about the multicast routes in the IPv4 multicast routing table.

To display the page, click **IPv4 Multicast → Multicast → Multicast Route Table** Multicast Route Table

Figure 43-3. Multicast Route Table

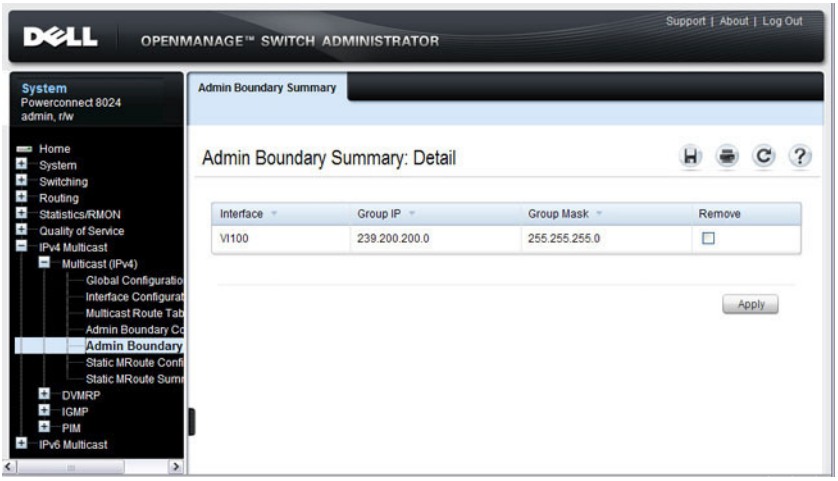


Multicast Admin Boundary Configuration

The definition of an administratively scoped boundary is a way to stop the ingress and egress of multicast traffic for a given range of multicast addresses on a given routing interface. Use the **Admin Boundary Configuration** page to configure a new or existing administratively scoped boundary. To see this page, you must have configured a valid routing interface and multicast.

To display the page, click **IPv4 Multicast → Multicast → Admin Boundary Configuration** in the navigation panel.

Figure 43-4. Multicast Admin Boundary Configuration

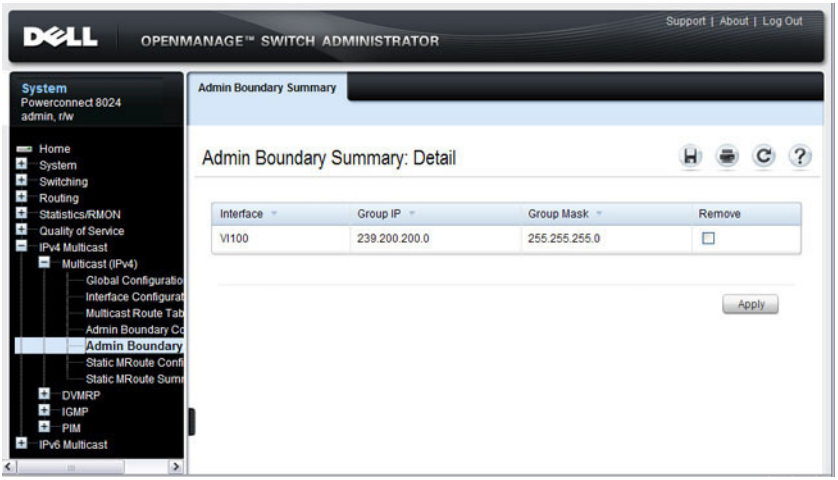


Multicast Admin Boundary Summary

Use the **Admin Boundary Summary** page to display existing administratively scoped boundaries.

To display the page, click **IPv4 Multicast → Multicast → Admin Boundary Summary** in the navigation panel.

Figure 43-5. Multicast Admin Boundary Summary

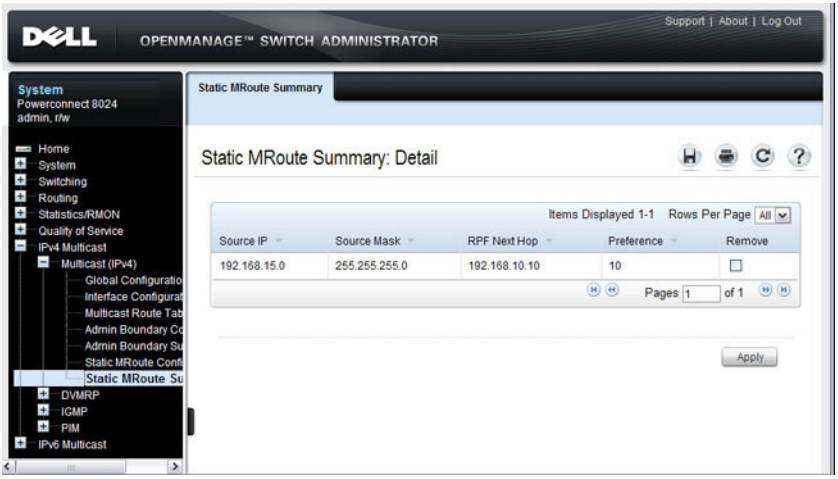


Multicast Static MRoute Configuration

Use the **Static MRoute Configuration** page to configure a new static entry in the Mroute table or to modify an existing entry.

To display the page, click **IPv4 Multicast → Multicast → Static MRoute Configuration** in the navigation panel.

Figure 43-6. Multicast Static MRoute Configuration

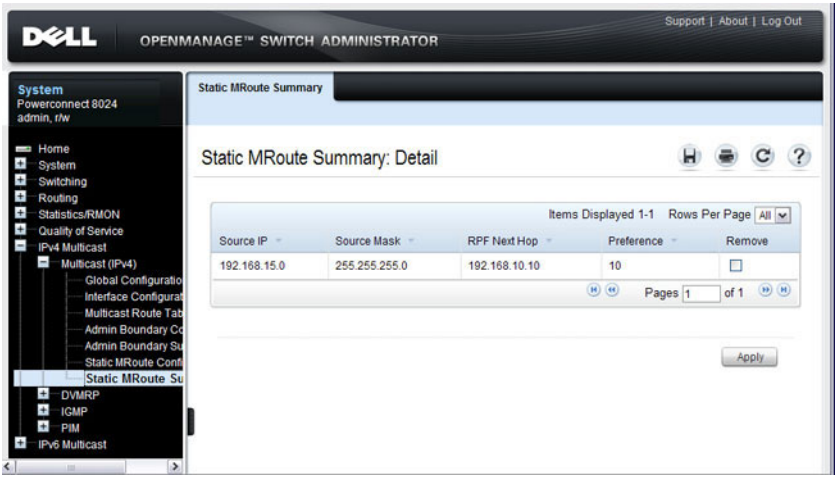


Multicast Static MRoute Summary

Use the **Static MRoute Summary** page to display static routes and their configurations.

To display the page, click **IPv4 Multicast → Multicast → Static MRoute Summary** in the navigation panel.

Figure 43-7. Multicast Static MRoute Summary



Configuring IPv6 Multicast Features (Web)

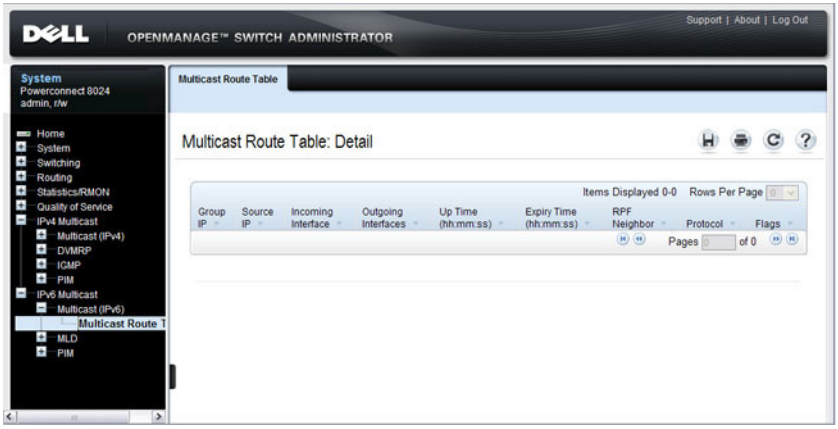
This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring the IPv6 multicast features that are not protocol-specific on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

IPv6 Multicast Route Table

Use the **Multicast Route Table** page to view information about the multicast routes in the IPv6 multicast routing table.

To display the page, click **IPv6 Multicast → Multicast → Multicast Route Table**.

Figure 43-8. IPv6 Multicast Route Table



Configuring IGMP and IGMP Proxy (Web)

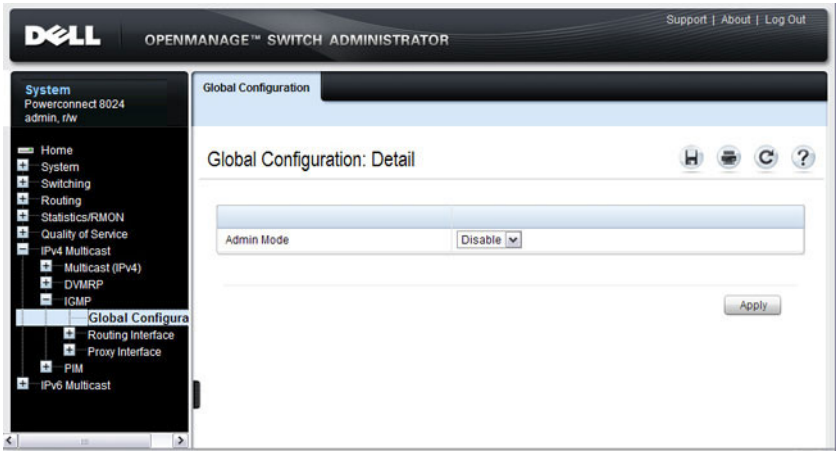
This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring the IGMP and IGMP proxy features on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

IGMP Global Configuration

Use the **Global Configuration** page to set IGMP on the system to active or inactive.

To display the page, click **IPv4 Multicast → IGMP → Global Configuration** in the navigation panel.

Figure 43-9. IGMP Global Configuration

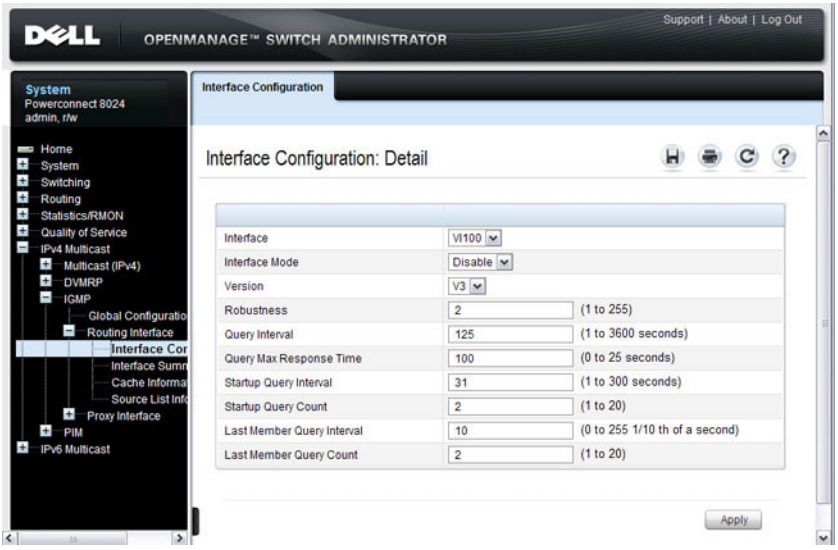


IGMP Interface Configuration

Use the **Interface Configuration** page to configure and/or display router interface parameters. You must configure at least one valid routing interface before you can access this page and configure IP Multicast IGMP.

To display the page, click **IPv4 Multicast → IGMP → Routing Interface → Interface Configuration** in the navigation panel.

Figure 43-10. IGMP Interface Configuration

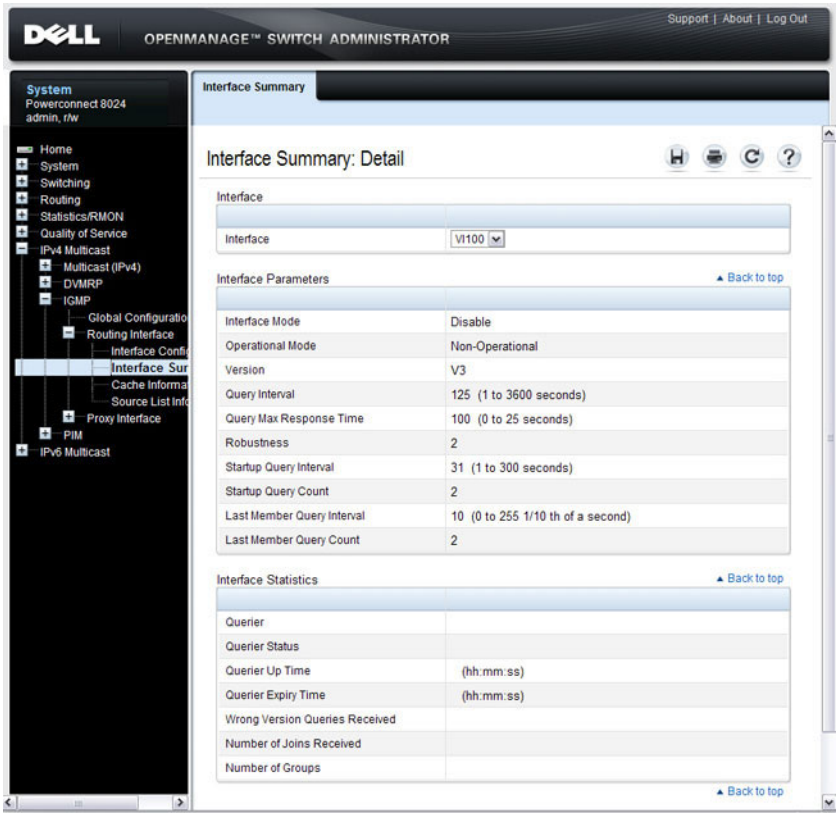


IGMP Interface Summary

Use the **Interface Summary** page to display IGMP routing parameters and data. You must configure at least one IGMP router interface to access this page.

To display the page, click **IPv4 Multicast → IGMP → Routing Interface → Interface Summary** in the navigation panel.

Figure 43-11. IGMP Interface Summary

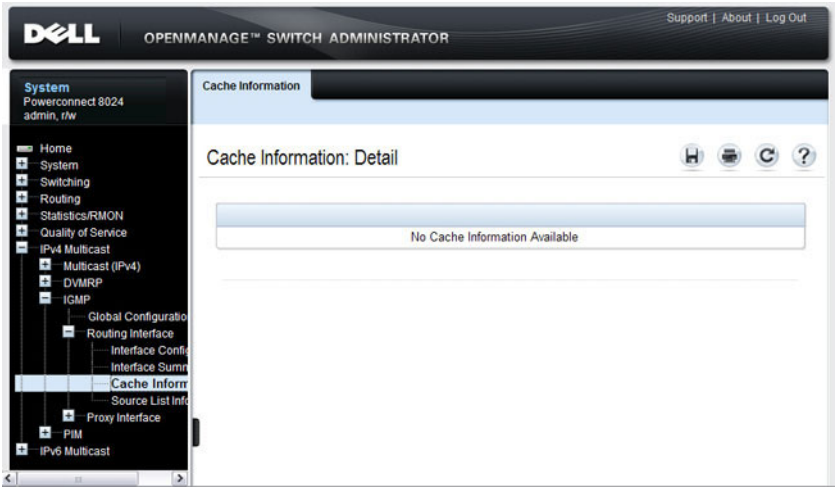


IGMP Cache Information

Use the **Cache Information** page to display cache parameters and data for an IP multicast group address. Group membership reports must have been received on the selected interface for data to display on the page.

To display the page, click **IPv4 Multicast → IGMP → Routing Interface → Cache Information** in the navigation panel.

Figure 43-12. IGMP Cache Information

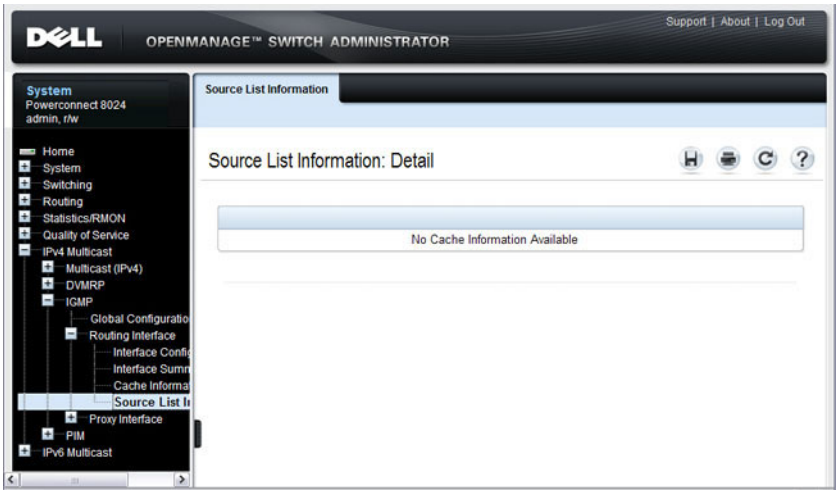


IGMP Interface Source List Information

Use the **Source List Information** page to display detailed membership information for an interface. Group membership reports must have been received on the selected interface for data to display information.

To display the page, click **IPv4 Multicast** → **IGMP** → **Routing Interface** → **Source List Information** in the navigation panel.

Figure 43-13. IGMP Interface Source List Information



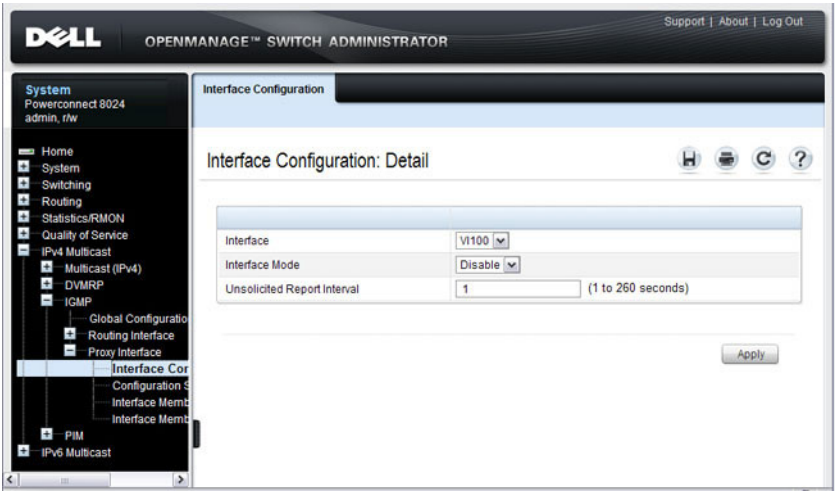
IGMP Proxy Interface Configuration

The IGMP Proxy is used by IGMP Router (IPv4 system) to enable the system to issue IGMP host messages on behalf of hosts that the system discovered through standard IGMP router interfaces. Thus, this feature acts as proxy to all hosts residing on its router interfaces.

Use the **Interface Configuration** page to configure IGMP proxy for an interface. You must have configured at least one router interface before configuring or displaying data for an IGMP proxy interface, and it should not be an IGMP routing interface.

To display the page, click **IPv4 Multicast → IGMP → Proxy Interface → Interface Configuration** in the navigation panel.

Figure 43-14. IGMP Proxy Interface Configuration

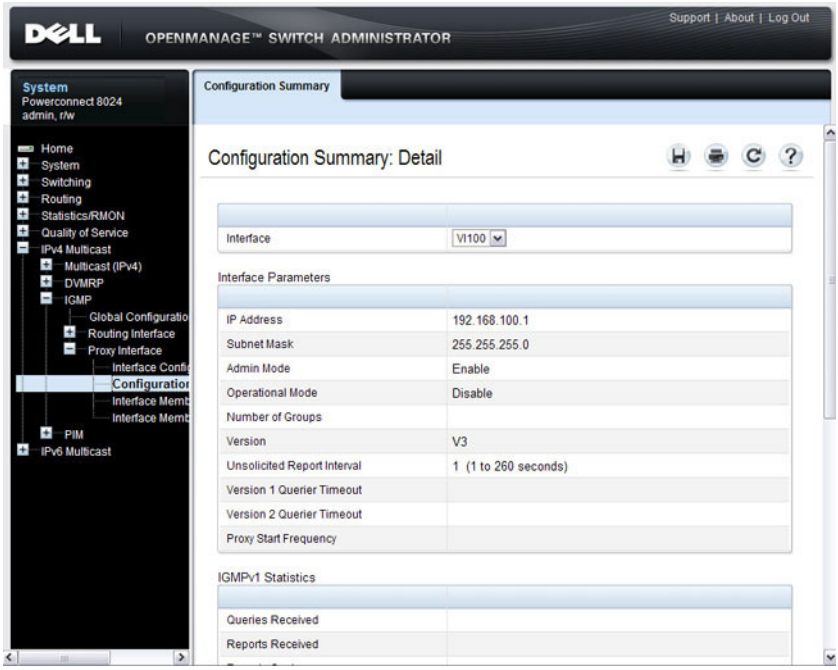


IGMP Proxy Configuration Summary

Use the **Configuration Summary** page to display proxy interface configurations by interface. You must have configured at least one router interface configured before data displays on this page.

To display the page, click **IPv4 Multicast → IGMP → Proxy Interface → Configuration Summary** in the navigation panel.

Figure 43-15. IGMP Proxy Configuration Summary

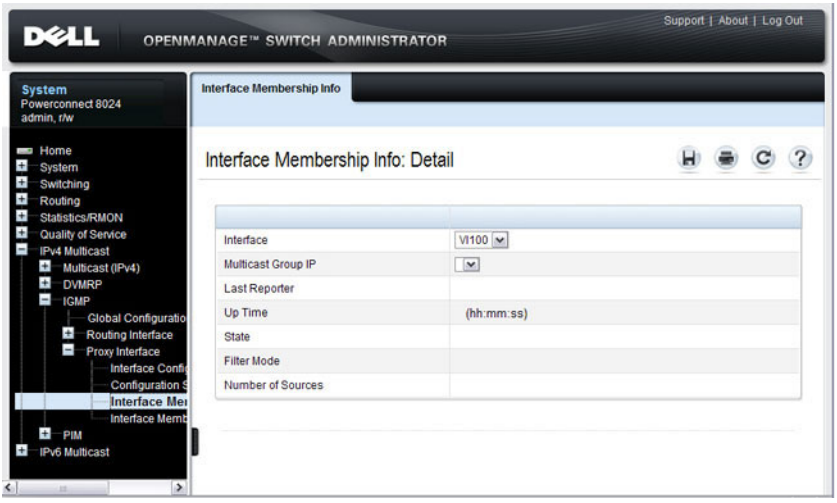


IGMP Proxy Interface Membership Info

Use the **Interface Membership Info** page to display interface membership data for a specific IP multicast group address. You must have configured at least one router interface before you can display interface membership information, and it should not be an IGMP routing interface. Also, if no group membership reports have been received on the selected interface, no data displays on this page.

To display the page, click **IPv4 Multicast** → **IGMP** → **Proxy Interface** → **Interface Membership Info** in the navigation panel.

Figure 43-16. IGMP Proxy Interface Membership Info

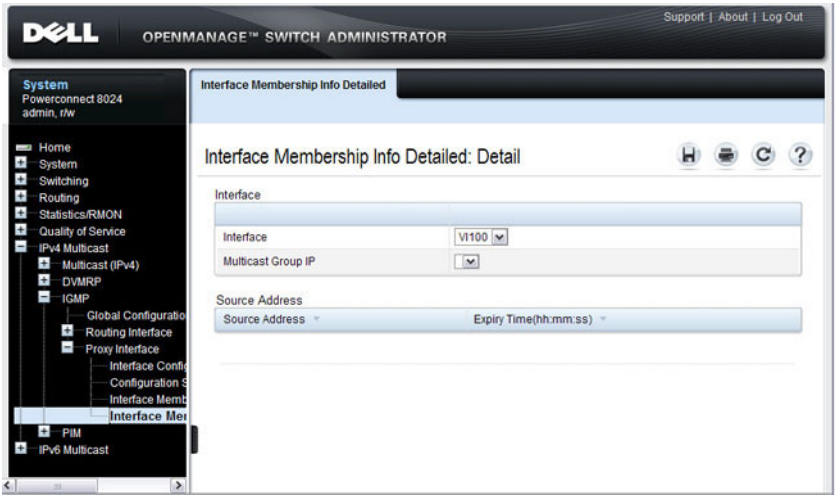


Detailed IGMP Proxy Interface Membership Information

Use the **Interface Membership Info Detailed** page to display detailed interface membership data. You must have configured at least one router interface before you can display detailed interface membership information, and it should not be an IGMP routing interface. Also, if no group membership reports have been received on the selected interface you cannot display data.

To display the page, click **IPv4 Multicast → IGMP → Proxy Interface → Interface Membership Info Detailed** in the navigation panel.

Figure 43-17. IGMP Proxy Interface Membership Info Detailed



Configuring MLD and MLD Proxy (Web)

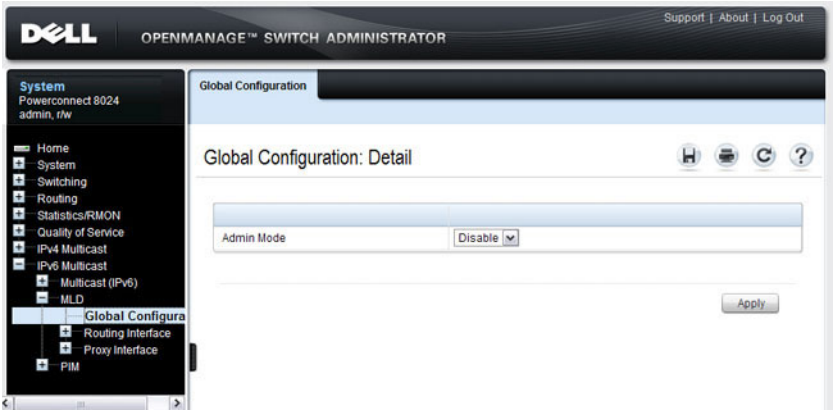
This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring the MLD and MLD proxy features on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

MLD Global Configuration

Use the **Global Configuration** page to administratively enable and disable the MLD service.

To display the page, click **IPv6 Multicast → MLD → Global Configuration** in the navigation panel.

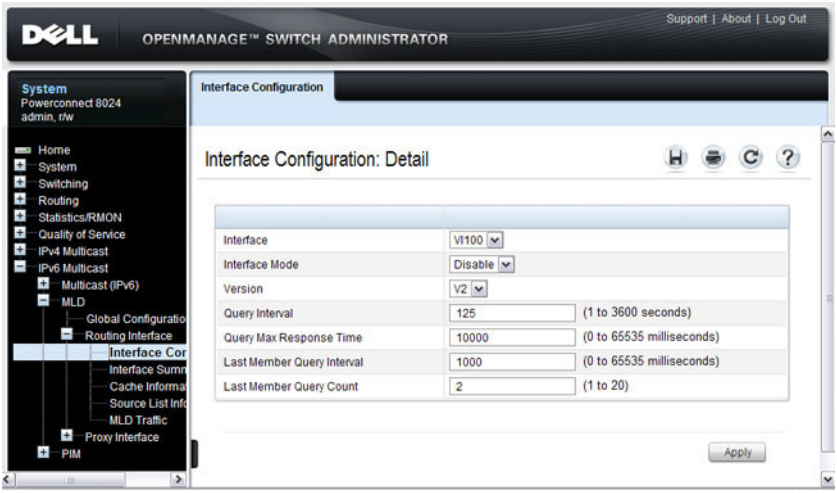
Figure 43-18. MLD Global Configuration



MLD Routing Interface Configuration

Use the **Interface Configuration** page to enable selected IPv6 router interfaces to discover the presence of multicast listeners, the nodes who wish to receive the multicast data packets, on its directly attached interfaces. To access this page, click **IPv6 Multicast → MLD → Routing Interface → Interface Configuration** in the navigation panel.

Figure 43-19. MLD Routing Interface Configuration

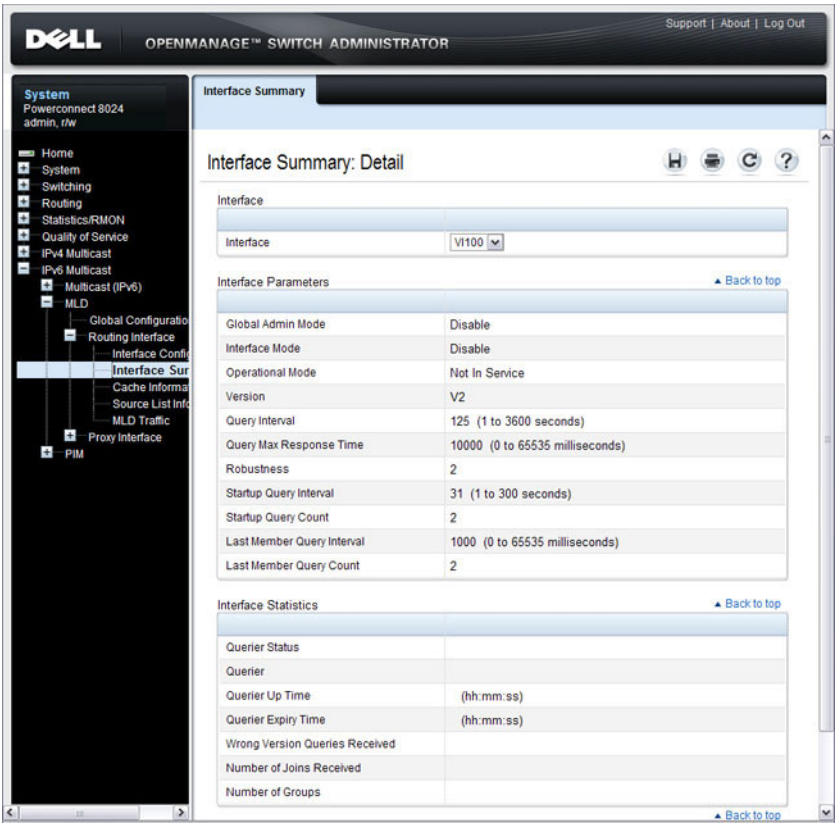


MLD Routing Interface Summary

Use the **Interface Summary** page to display information and statistics on a selected MLD-enabled interface. You must configure at least one IGMP router interface to access this page.

To access this page, click **IPv6 Multicast → MLD → Routing Interface → Interface Summary** in the navigation panel.

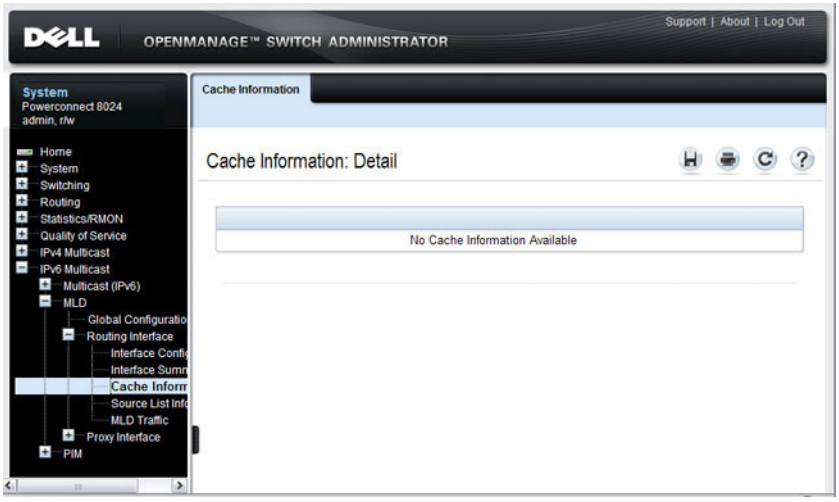
Figure 43-20. MLD Routing Interface Summary



MLD Routing Interface Cache Information

The **Interface Cache Information** page displays cache parameters and data for an IP multicast group address that has been reported to operational MLD routing interfaces. You must configure at least one MLD router interface to access this page. Also, group membership reports must have been received on the selected interface in order for data to be displayed here. To access this page, click **IPv6 Multicast → MLD → Routing Interface → Cache Information** in the navigation panel.

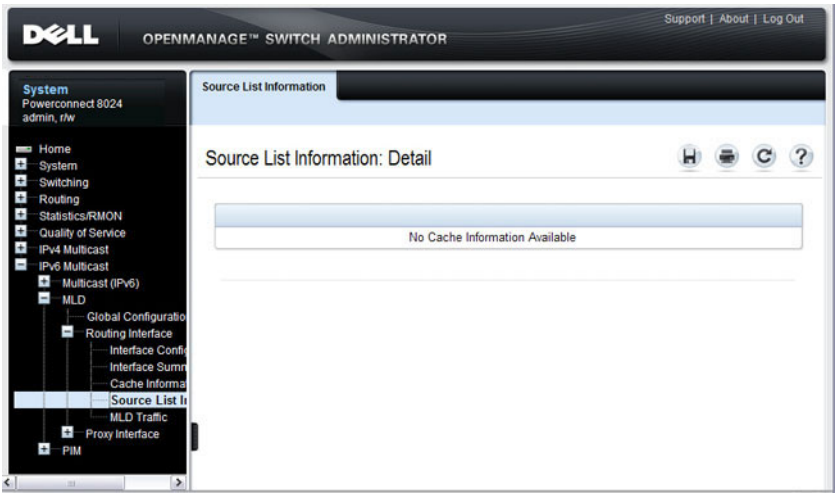
Figure 43-21. MLD Routing Interface Cache Information



MLD Routing Interface Source List Information

The **Interface Source List Information** page displays detailed membership information for an interface. You must configure at least one MLD router interface to access this page. Also, group membership reports must have been received on the selected interface in order for data to be displayed here. To access this page, click **IPv6 Multicast → MLD → Routing Interface → Source List Information** in the navigation panel.

Figure 43-22. MLD Routing Interface Source List Information

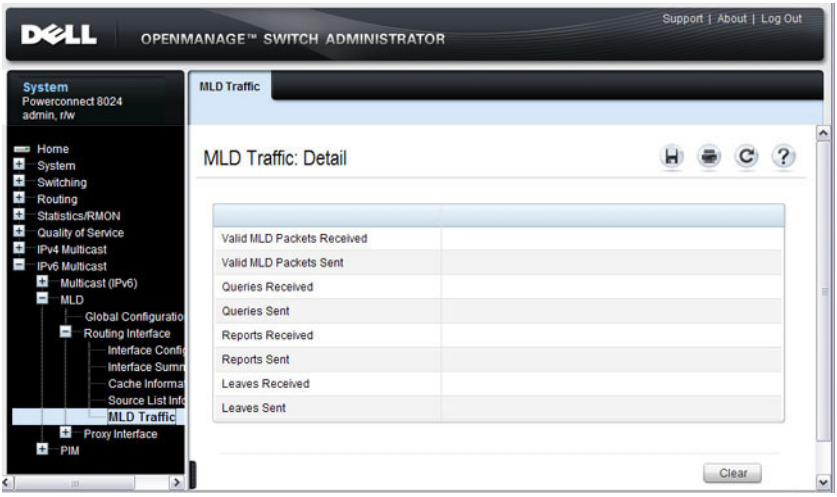


MLD Traffic

The **MLD Traffic** page displays summary statistics on the MLD messages sent to and from the router.

To access this page, click **IPv6 Multicast** → **MLD** → **Routing Interface** → **MLD Traffic** in the navigation panel.

Figure 43-23. MLD Traffic

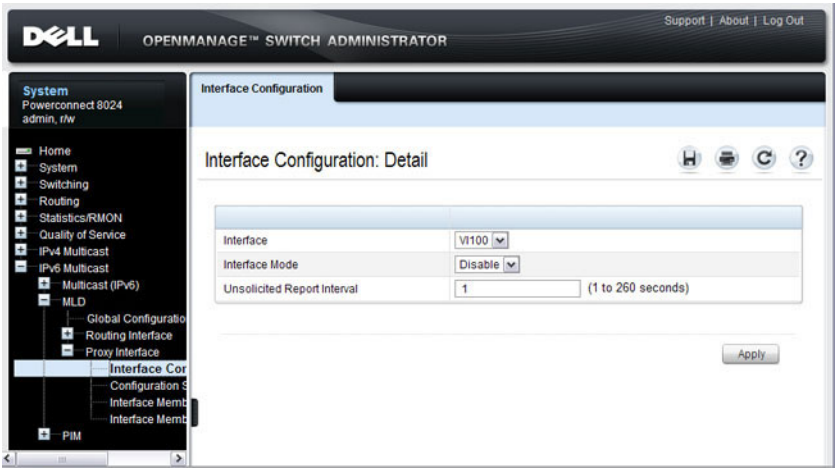


MLD Proxy Configuration

When you configure an interface in MLD proxy mode, it acts as a proxy multicast host that sends MLD membership reports on one interface for MLD Membership reports received on all other MLD-enabled router interfaces.

Use the **Interface Configuration** page to enable and disable ports as MLD proxy interfaces. To display this page, click **IPv6 Multicast → MLD → Proxy Interface → Interface Configuration** in the navigation panel.

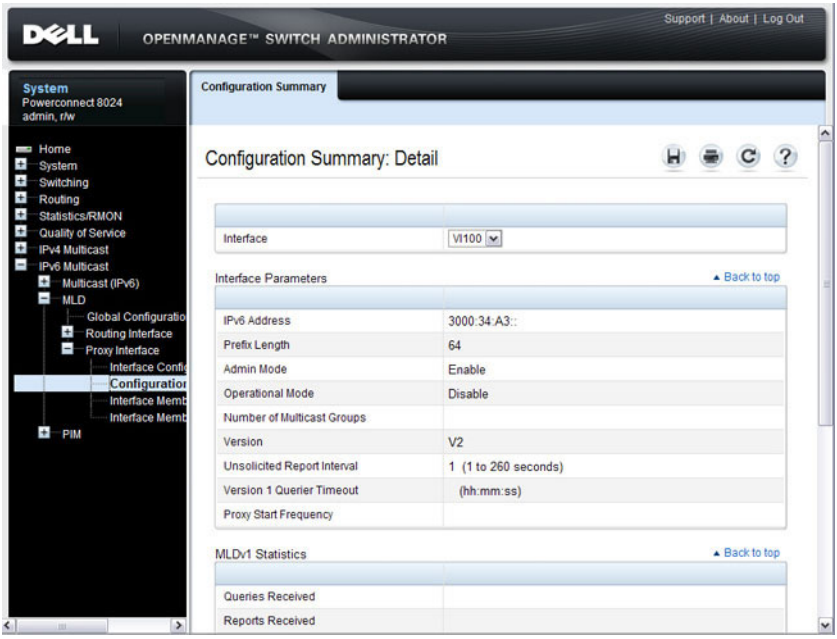
Figure 43-24. MLD Proxy Interface Configuration



MLD Proxy Configuration Summary

Use the **Configuration Summary** page to view configuration and statistics on MLD proxy-enabled interfaces. To display this page, click **IPv6 Multicast** → **MLD** → **Proxy Interface** → **Configuration Summary** in the navigation panel.

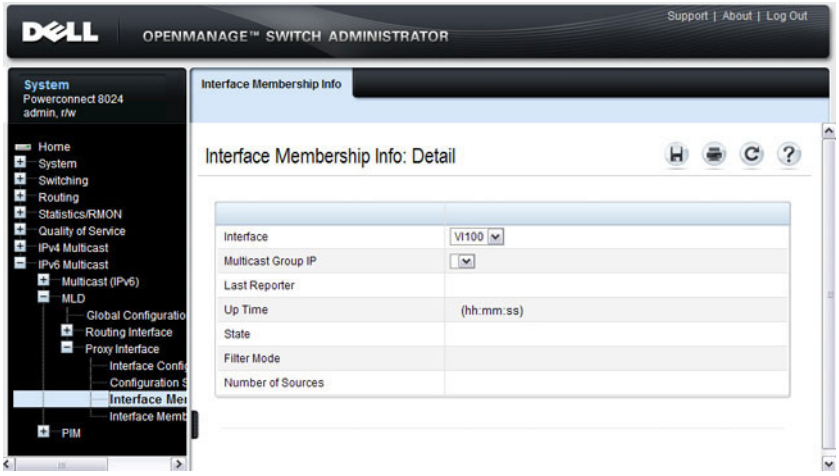
Figure 43-25. MLD Proxy Configuration Summary



MLD Proxy Interface Membership Information

The **Interface Membership Information** page lists each IP multicast group for which the MLD proxy interface has received membership reports. To display this page, click **IPv6 Multicast → MLD → Proxy interface → Interface Membership Info** in the navigation panel.

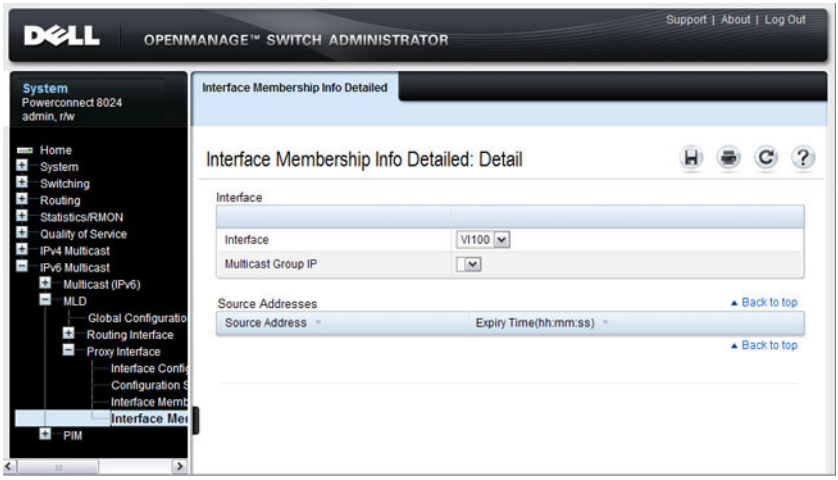
Figure 43-26. Interface Membership Information



Detailed MLD Proxy Interface Membership Information

The **Interface Membership Information Detailed** page provides additional information about the IP multicast groups for which the MLD proxy interface has received membership reports. To display this page, click **IPv6 Multicast → MLD → Proxy Interface → Interface Membership Info Detailed** in the navigation panel.

Figure 43-27. Interface Membership Information—Detailed



Configuring PIM for IPv4 and IPv6 (Web)

This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring PIM-SM and PIM-DM for IPv4 and IPv6 multicast routing on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

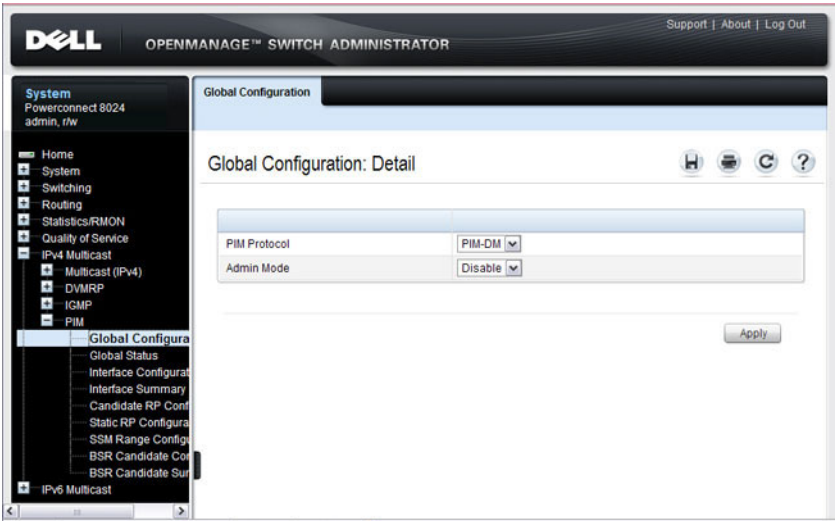
NOTE: The OpenManage Switch Administrator pages to configure IPv4 multicast routing and IPv6 multicast routing is very similar. The figures in this section show the IPv4 multicast configuration pages. To configure IPv6 multicast with PIM, use the pages available from the IPv6 Multicast → PIM menu.

PIM Global Configuration

Use the **Global Configuration** page to configure the administrative status of PIM-DM or PIM-SM on the switch.

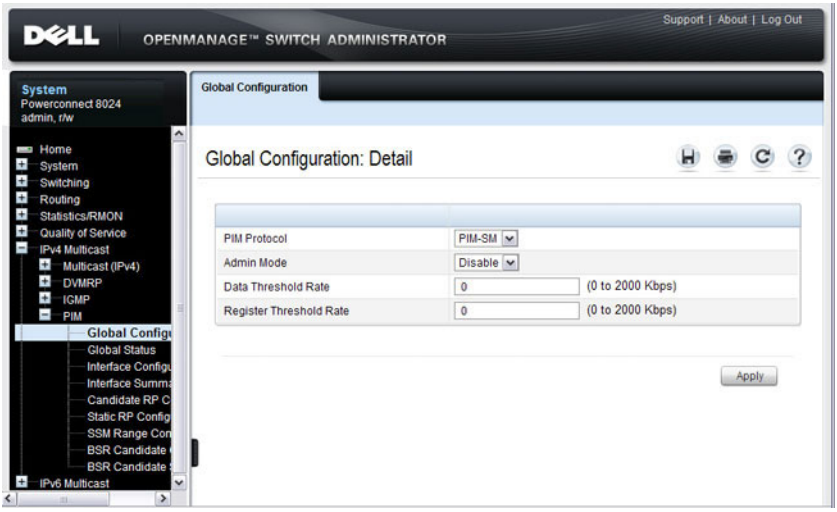
To display the page, click **IPv4 Multicast → PIM → Global Configuration** or **IPv6 Multicast → PIM → Global Configuration** in the navigation panel.

Figure 43-28. PIM-DM Global Configuration



If you select PIM-SM as the PIM protocol, additional fields appear, as Figure 43-29 shows.

Figure 43-29. PIM-SM Global Configuration

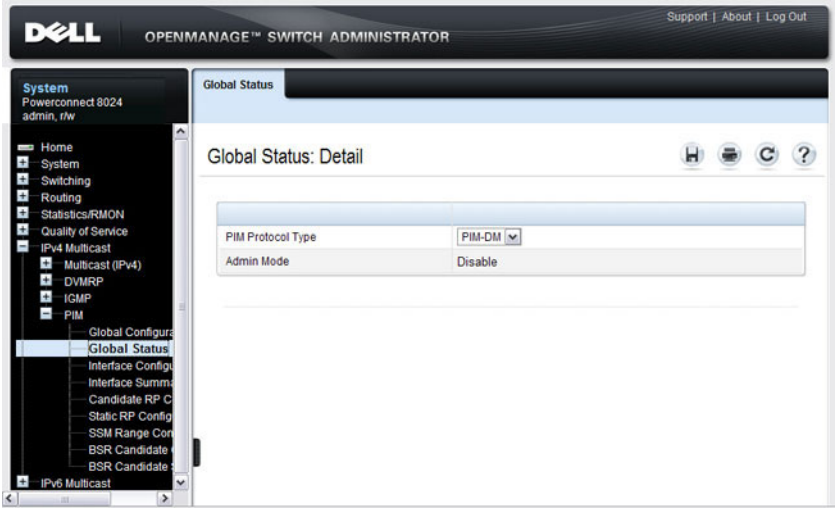


PIM Global Status

Use the **Global Status** page to view the administrative status of PIM-DM or PIM-SM on the switch.

To display the page, click **IPv4 Multicast → PIM → Global Status** or **IPv6 Multicast → PIM → Global Status** in the navigation panel.

Figure 43-30. PIM Global Status

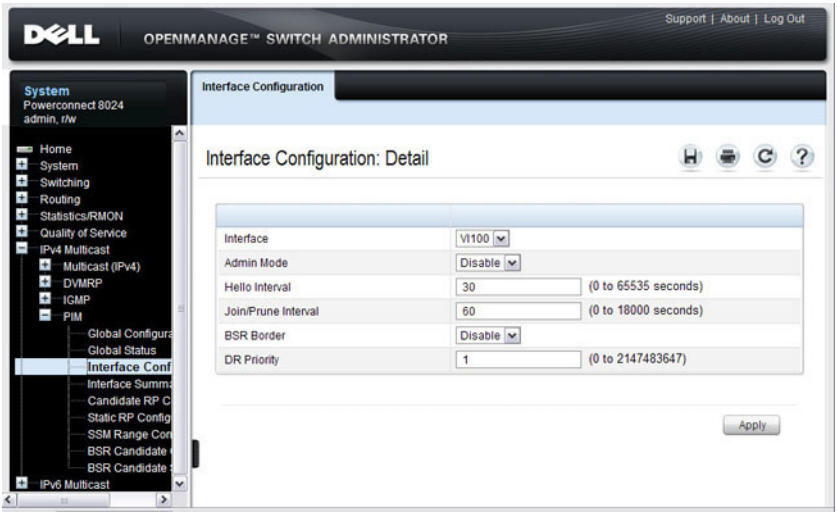


PIM Interface Configuration

Use the **Interface Configuration** page to configure specific interfaces with PIM.

To display the page, click **IPv4 Multicast → PIM → Interface Configuration** or **IPv6 Multicast → PIM → Interface Configuration** in the navigation panel.

Figure 43-31. PIM Interface Configuration

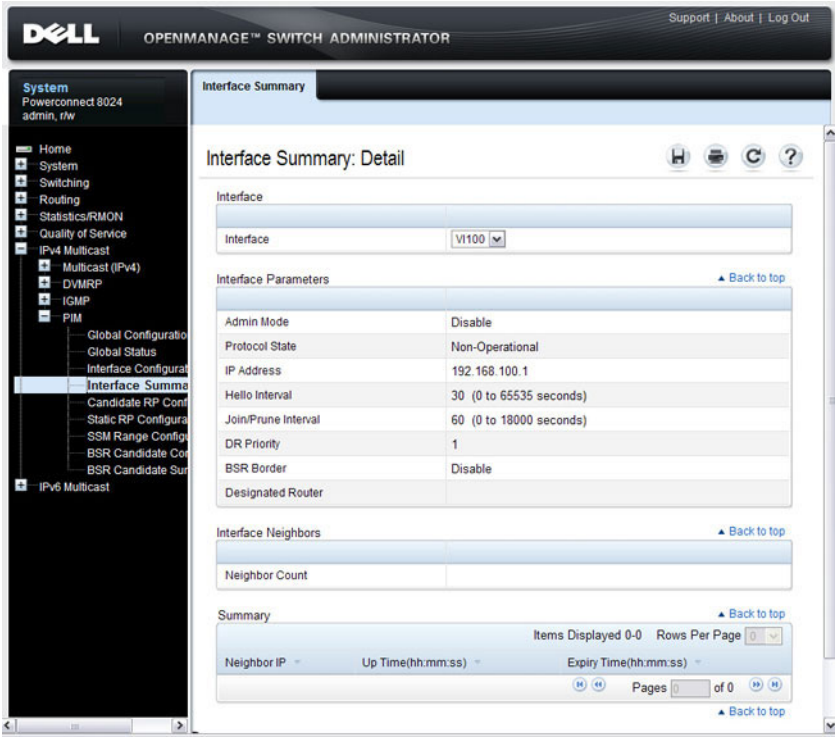


PIM Interface Summary

Use the **Interface Summary** page to display a PIM interface and its settings.

To display the page, click **IPv4 Multicast → PIM → Interface Summary** or **IPv6 Multicast → PIM → Interface Summary** in the navigation panel.

Figure 43-32. PIM Interface Summary

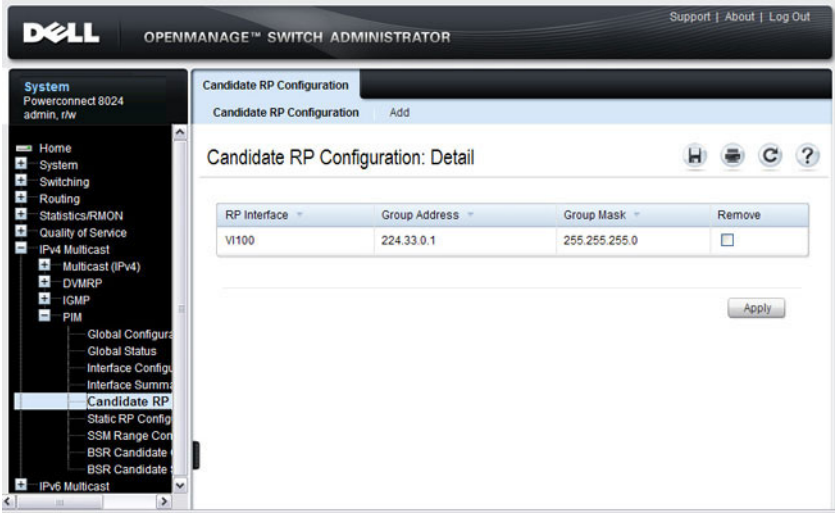


Candidate RP Configuration

The Candidate RP is configured on the **Add Candidate RP** page. Use the **Candidate RP Configuration** page to display and delete the configured rendezvous points (RPs) for each port using PIM.

To access the page, click **IPv4 Multicast → PIM → Candidate RP Configuration** or **IPv6 Multicast → PIM → Candidate RP Configuration**.

Figure 43-33. Candidate RP Configuration



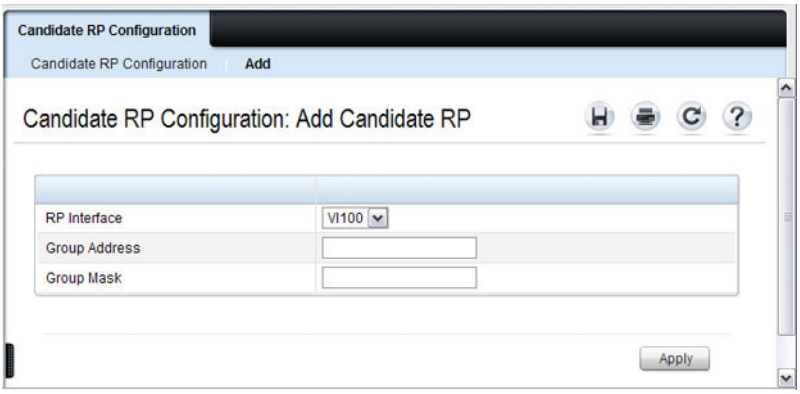
Adding a Candidate RP

To add PIM Candidate rendezvous points (RPs) for each IP multicast group:

1. Open the **Candidate RP Configuration** page.
2. Click **Add**.

   The **Add Candidate RP** page displays.

Figure 43-34. Add Candidate RP



3 Select the interface for which the Candidate RP is to be configured.

4 Enter the group address transmitted in Candidate-RP-Advertisements.

5 Enter the prefix length transmitted in Candidate-RP-Advertisements to fully identify the scope of the group which the router supports if elected as a Rendezvous Point.

6 Click **Apply Changes**.

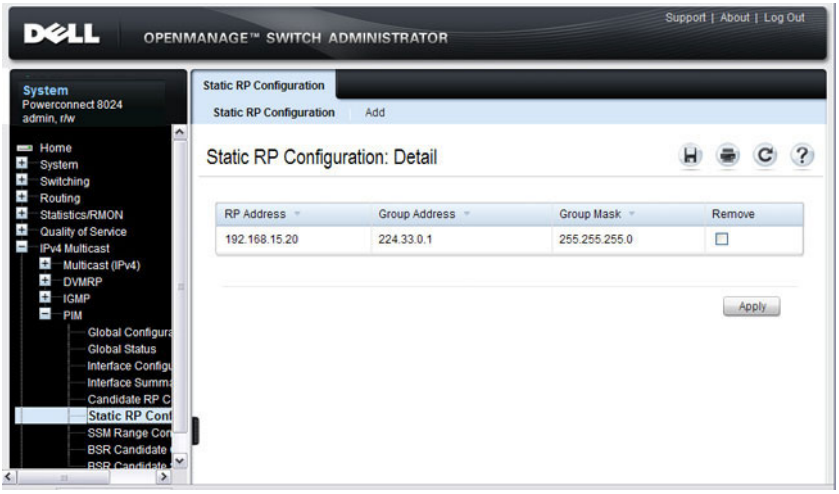
The new Candidate RP is added, and the device is updated.

Static RP Configuration

Use the **Static RP Configuration** page to display or remove the configured RP. The page also allows adding new static RPs by clicking the **Add** button. Only one RP address can be used at a time within a PIM domain. If the PIM domain uses the BSR to dynamically learn the RP, configuring a static RP is not required. However, you can configure the static RP to override any dynamically learned RP from the BSR.

To access the page, click **IPv4 Multicast → PIM → Static RP Configuration** or **IPv6 Multicast → PIM → Static RP Configuration**.

Figure 43-35. Static RP Configuration



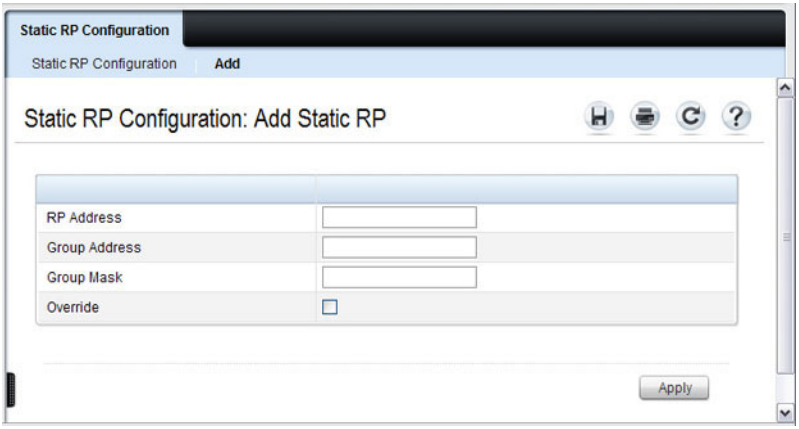
Adding a Static RP

To add a static RP for the PIM router.

1. Open the **Static RP Configuration** page.
2. Click **Add**.

   The **Add Static RP** page displays.

Figure 43-36. Add Static RP



3 Enter the IP address of the RP for the group range.

4 Enter the group address of the RP.

5 Enter the group mask of the RP.

6 Check the **Override** option to configure the static RP to override the dynamic (candidate) RPs learned for same group ranges.

7 Click **Apply**.

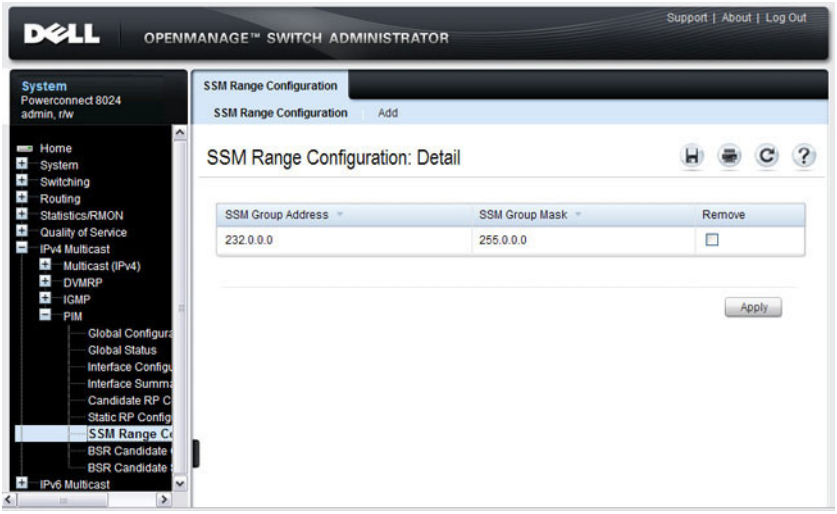
The new Static RP is added, and the device is updated.

SSM Range Configuration

Use this page to display or remove the Source Specific Multicast (SSM) group IP address and group mask for the PIM router.

To display the page, click **IPv4 Multicast → PIM → SSM Range Configuration** or **IPv6 Multicast → PIM → SSM Range Configuration**.

Figure 43-37. SSM Range Configuration



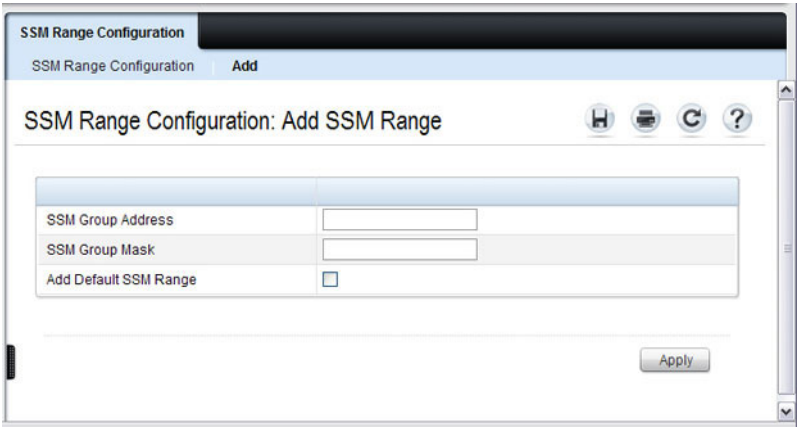
Adding an SSM Range

To add the Source-Specific Multicast (SSM) Group IP Address and Group Mask (IPv4) or Prefix Length (IPv6) for the PIM router:

1. Open the **SSM Range Configuration** page.
2. Click **Add**.

   The **Add SSM Range** page displays.

Figure 43-38. Add SSM Range



3 Click the Add Default SSM Range check box to add the default SSM Range. The default SSM Range is 232.0.0.0/8 for IPv4 multicast and ff3x::/32 for IPv6 multicast.

4 Enter the SSM Group IP Address.

5 Enter the SSM Group Mask (IPv4) or SSM Prefix Length (IPv6).

6 Click **Apply**.

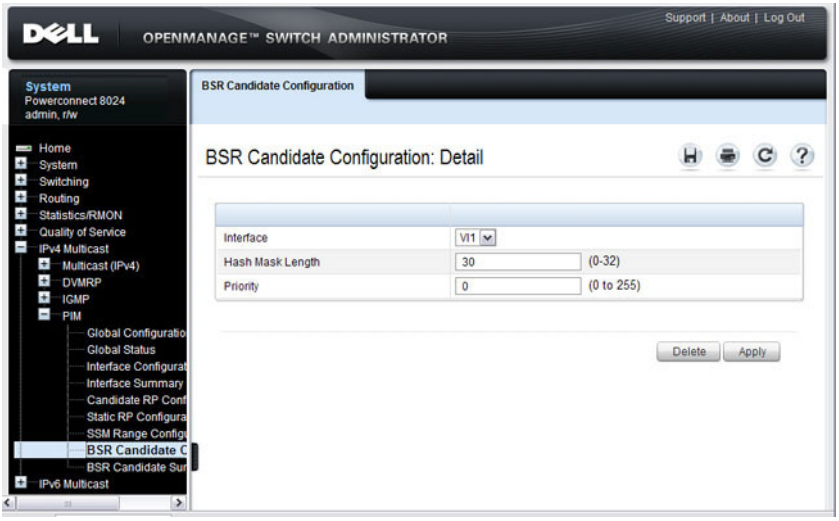
The new SSM Range is added, and the device is updated.

BSR Candidate Configuration

Use this page to configure information to be used if the interface is selected as a bootstrap router.

To display the page, click **IPv4 Multicast → PIM → BSR Candidate Configuration** or **IPv6 Multicast → PIM → BSR Candidate Configuration**.

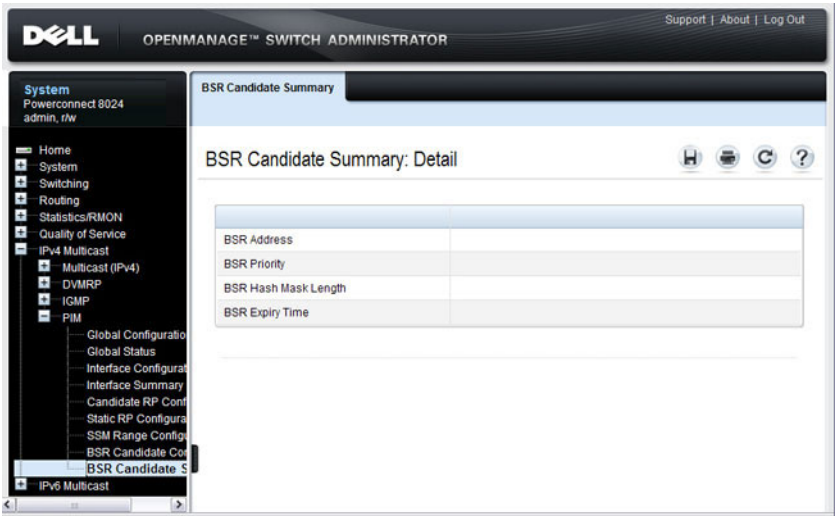
Figure 43-39. BSR Candidate Configuration



BSR Candidate Summary

Use this page to display information about the configured BSR candidates. To display this page, click **IPv4 Multicast → PIM → BSR Candidate Summary** or **IPv6 Multicast → PIM → BSR Candidate Summary**.

Figure 43-40. BSR Candidate Summary



Configuring DVMRP (Web)

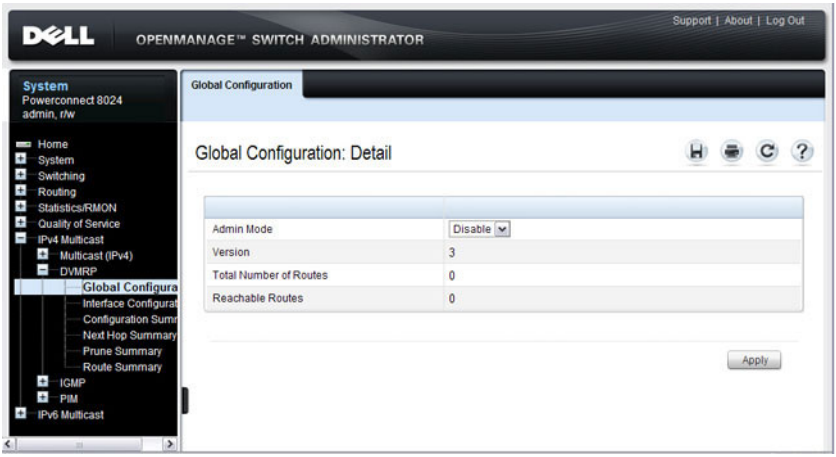
This section provides information about the OpenManage Switch Administrator pages for configuring and monitoring DVMRP on a PowerConnect 8000-series and 8100-series switches. For details about the fields on a page, click ? at the top of the page.

DVMRP Global Configuration

Use the **Global Configuration** page to configure global DVMRP settings.

To display the page, click **IPv4 Multicast → DVMRP → Global Configuration** in the navigation panel.

Figure 43-41. DVMRP Global Configuration

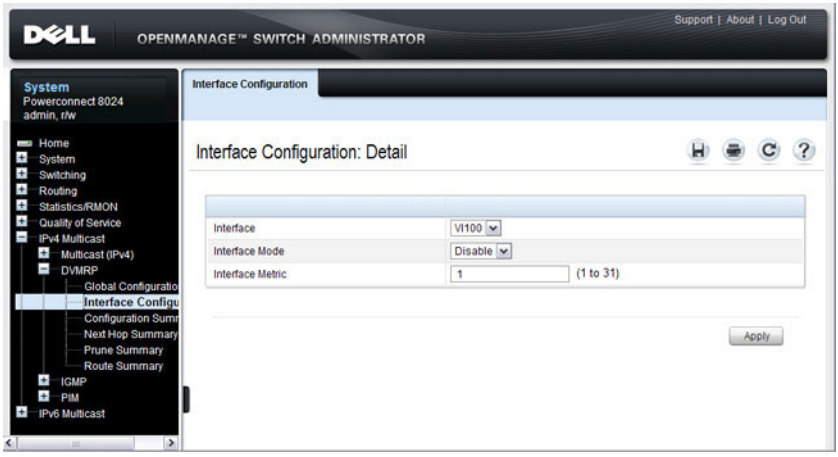


DVMRP Interface Configuration

Use the **Interface Configuration** page to configure a DVMRP interface. You must configure at least one router interface before you configure a DVMRP interface. Otherwise you see a message telling you that no router interfaces are available, and the configuration screen is not displayed.

To display the page, click **IPv4 Multicast → DVMRP → Interface Configuration** in the navigation panel.

Figure 43-42. DVMRP Interface Configuration

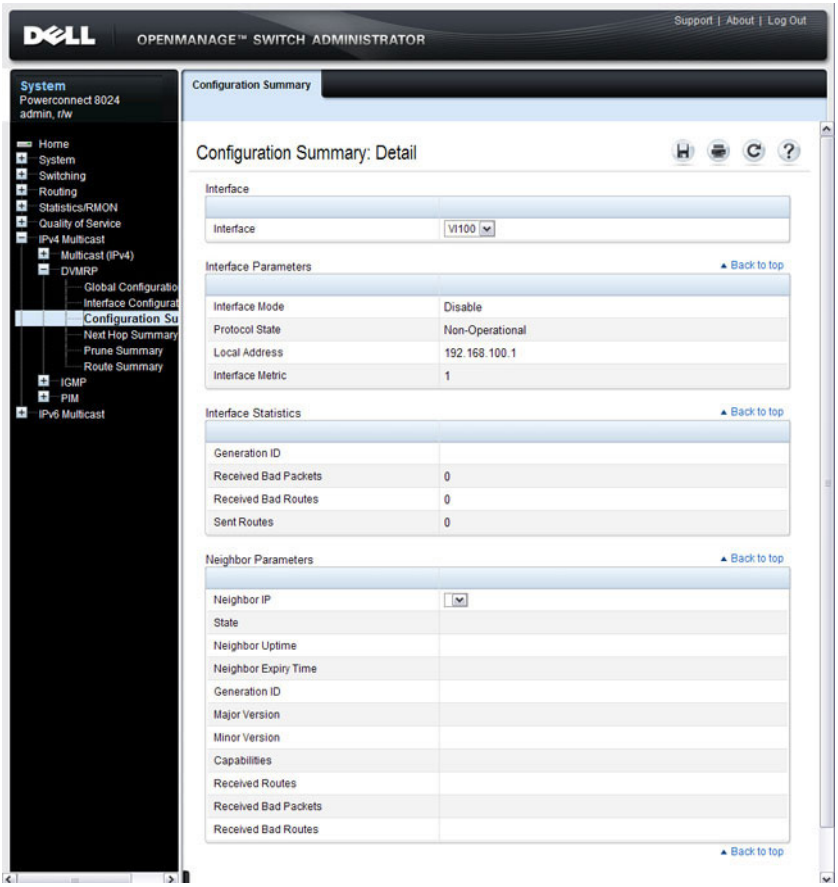


DVMRP Configuration Summary

Use the **Configuration Summary** page to display the DVMRP configuration and data for a selected interface. You must configure at least one router interface before you can display data for a DVMRP interface. Otherwise you see a message telling you that no router interfaces are available, and the configuration summary screen is not displayed.

To display the page, click **IPv4 Multicast → DVMRP → Configuration Summary** in the navigation panel.

Figure 43-43. DVMRP Configuration Summary

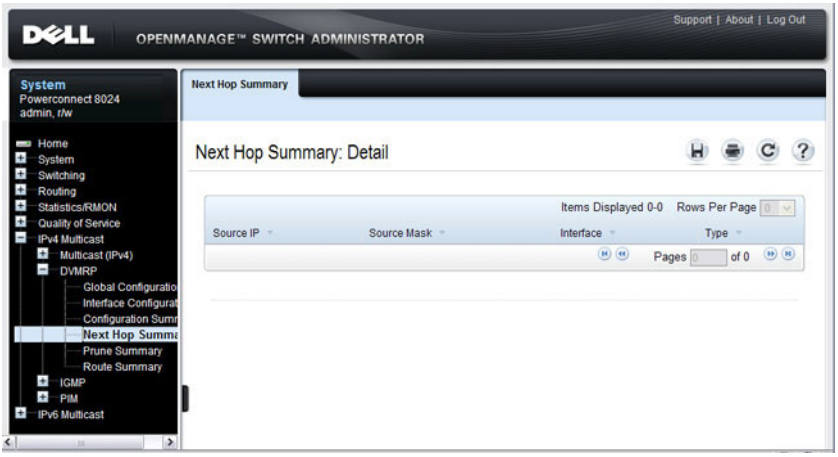


DVMRP Next Hop Summary

Use the **Next Hop Summary** page to display the next hop summary by Source IP.

To display the page, click **IPv4 Multicast → DVMRP → Next Hop Summary** in the navigation panel.

Figure 43-44. DVMRP Next Hop Summary

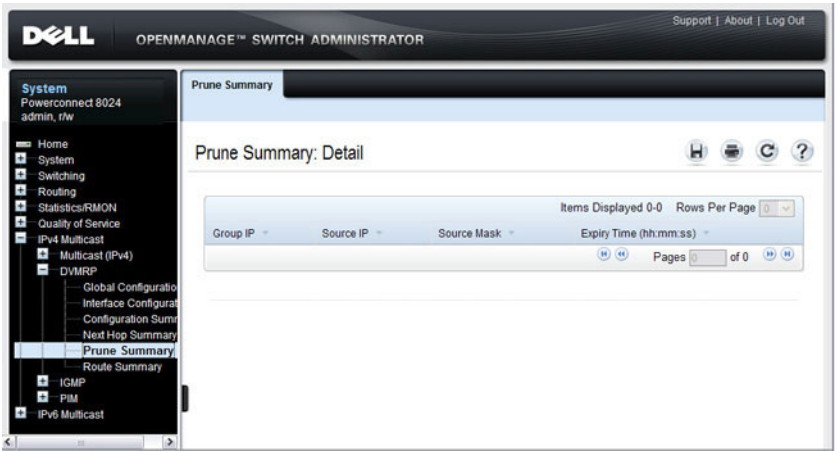


DVMRP Prune Summary

Use the **Prune Summary** page to display the prune summary by Group IP.

To display the page, click **IPv4 Multicast → DVMRP → Prune Summary** in the navigation panel.

Figure 43-45. DVMRP Prune Summary

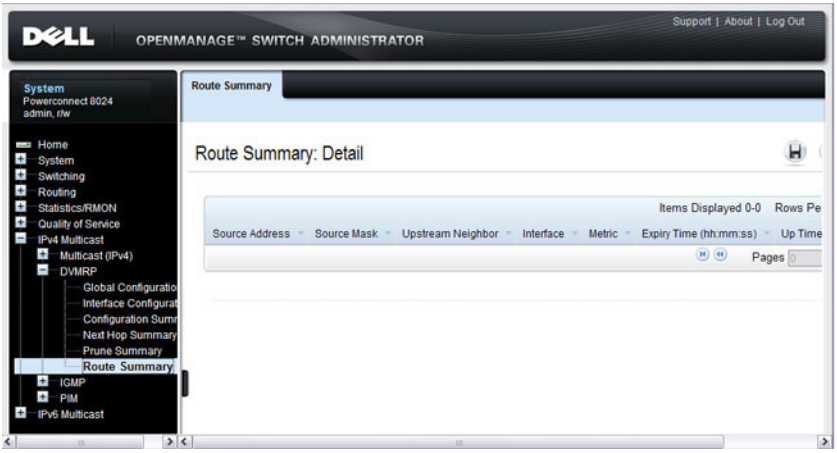


DVMRP Route Summary

Use the **Route Summary** page to display the DVMRP route summary.

To display the page, click **IPv4 Multicast → DVMRP → Route Summary** in the navigation panel.

Figure 43-46. DVMRP Route Summary



Configuring L3 Multicast Features (CLI)

This section provides information about the commands you use to configure general IPv4 multicast settings on the switch. For more information about the commands, see the *PowerConnect 8024/8024F/8132/8132F/8164/8164F CLI Reference Guide* at support.dell.com/manuals.

Configuring and Viewing IPv4 Multicast Information

Beginning in Privileged EXEC mode, use the following commands to enable IPv4 multicast on the switch and to view and configure other general multicast settings.

Command	Purpose
configure	Enter global configuration mode.
ip multicast	Enable IPv4 multicast on the switch.
ip mroute *source-address mask rpf-address preference*	Create a static multicast route for a source range. • *source-address* — The IP address of the multicast data source. • *mask* — The IP subnet mask of the multicast data source. • *rpf-address* — The IP address of the next hop towards the source. • *preference* — The cost of the route (Range: 1 - 255).
interface vlan *vlan-id*	Enter Interface Configuration mode for the specified VLAN.
ip mcast boundary *groupipaddr mask*	Add an administrative scope multicast boundary specified by the multicast group IP address (*groupipaddr*) and group IP subnet mask (*mask*) for which this multicast administrative boundary is applicable. The group IP address valid range is 239.0.0.0 to 239.255.255.255.
ip multicast ttl-threshold *ttlvalue*	Apply a Time to Live (TTL) value to the interface. The *ttlvalue* is the TTL threshold which is applied to the multicast data packets forwarded through the interface.

Command	Purpose
exit	Exit to Global Config mode.
exit	Exit to Privileged EXEC mode.
show ip multicast	View system-wide multicast information.
show ip mcast boundary {**vlan** *vlan-id* | **all**}	View all the configured administrative scoped multicast boundaries.
show ip mcast interface [**vlan** *vlan-id*]	View the multicast information for all interfaces or for the specified interface.
show ip mcast mroute {**detail** | **summary**}	View a summary or all the details of the multicast table.
show bridge multicast address-table count	View statistical information about the entries in the multicast address table.
show ip mcast mroute group *groupipaddr* {**detail** | **summary**}	View the multicast configuration settings such as flags, timer settings, incoming and outgoing interfaces, RPF neighboring routers, and expiration times of all the entries in the multicast mroute table containing the *groupipaddr* value.
show ip mcast mroute source *sourceipaddr* {**summary** | *groupipaddr*}	View the multicast configuration settings such as flags, timer settings, incoming and outgoing interfaces, RPF neighboring routers, and expiration times of all the entries in the multicast mroute table containing the *sourceipaddr* or *sourceipaddr* | *groupipaddr* pair value(s).
show ip mcast mroute static [*sourceipaddr*]	View all the static routes configured in the static mcast table if it is specified or display the static route associated with the particular *sourceipaddr.*

Configuring and Viewing IPv6 Multicast Route Information

Beginning in Privileged EXEC mode, use the following commands to configure static IPv6 multicast routes on the switch and to view IPv6 multicast table information.

Command	Purpose
configure	Enter global configuration mode.
ipv6 mroute *source-address/prefix-length rpf-address preference*	Create a static multicast route for a source range. • *source-address/prefix-length* — The IPv6 address of the multicast data source. • *rpf-address* — The IPv6 address of the next hop towards the source. • *preference* — The cost of the route (Range: 1 - 255).
exit	Exit to Privileged EXEC mode.
show ip mcast interface [**vlan** *vlan-id*]	View the multicast information for all interfaces or for the specified interface.
show ipv6 mroute {**detail** | **summary**}	View a summary or all the details of the multicast table.
show ipv6 mroute group *groupipaddr* {**detail** | **summary**}	View the multicast configuration settings such as flags, timer settings, incoming and outgoing interfaces, RPF neighboring routers, and expiration times of all the entries in the multicast mroute table containing the *groupipaddr* value.
show ipv6 mroute source *sourceipaddr* {**summary** | *groupipaddr*}	View the multicast configuration settings such as flags, timer settings, incoming and outgoing interfaces, RPF neighboring routers, and expiration times of all the entries in the multicast mroute table containing the *sourceipaddr* or *sourceipaddr* | *groupipaddr* pair value(s).
show ipv6 mroute static [*sourceipaddr*]	View all the static routes configured in the static mcast table if it is specified or display the static route associated with the particular *sourceipaddr.*

Configuring and Viewing IGMP

Beginning in Privileged EXEC mode, use the following commands to configure IGMP on the switch and on VLAN routing interfaces and to view IGMP information.

Command	Purpose
configure	Enter global configuration mode.
ip igmp	Enable IGMP on the switch.
interface vlan *vlan-id*	Enter Interface Configuration mode for the specified VLAN.
ip igmp	Enable IGMP on the interface.
ip igmp version *version*	Set the version of IGMP for an interface. The *version* variable can be 1, 2, or 3.
ip igmp robustness *robustness*	Configure the robustness that allows tuning of the interface, that is, tuning for the expected packet loss on a subnet. If a subnet is expected to have significant loss, the robustness variable may be increased for the interface. The range for *robustness* is 1–255.
ip igmp query-interval *seconds*	Configure the query interval for the specified interface. The query interval determines how fast IGMP Host-Query packets are transmitted on this interface. The range for *seconds* is 0–3600 seconds.
ip igmp query-max-response-time *seconds*	Configure the maximum response time interval for the specified interface. It is the maximum query response time advertised in IGMPv2 queries on this interface. The range for *seconds* is 0–25 seconds.
ip igmp startup-query-interval *seconds*	Set the interval between general queries sent at startup on the interface. The range for *seconds* is 0–300 seconds.
ip igmp startup-query-count *count*	Set the number of queries sent out on startup —at intervals equal to the startup query interval for the interface. The range for *count* is 1–20.

Command	Purpose
ip igmp last-member-query-interval *tenthsofseconds*	Configure the Maximum Response Time inserted in Group-Specific Queries which are sent in response to Leave Group messages. The range is 0–255 tenths of a second.
ip igmp last-member-query-count *count*	Set the number of Group-Specific Queries sent before the router assumes that there are no local members on the interface. The range for *count* is 1–20.
CTRL + Z	Exit to Privileged EXEC mode.
show ip igmp	View system-wide IGMP information.
show ip igmp interface [**vlan** *vlan-id*]	View IGMP information for all interfaces or for the specified interface.
show ip igmp interface stats [**vlan** *vlan-id*]	View IGMP statistics for all interfaces or for the specified interface.
show ip igmp groups [**interface vlan** *vlan-id*]	View the registered multicast groups on the interface.
show ip igmp membership	View the list of interfaces that have registered in any multicast group.

Configuring and Viewing IGMP Proxy

Beginning in Privileged EXEC mode, use the following commands to configure the upstream VLAN routing interface as an IGMP proxy. The IGMP proxy issues host messages on behalf of the hosts that have been discovered on IGMP-enabled interfaces. The upstream interface is the interface closest to the root multicast router, which should be running IGMP.

NOTE: Configure only the upstream interface as the IGMP proxy. IGMP should be enabled on all downstream interfaces. IP routing and IP multicast must be enabled on the switch for the IGMP proxy feature to operate.

Command	Purpose
configure	Enter global configuration mode.
interface vlan *vlan-id*	Enter Interface Configuration mode for the specified VLAN.
ip igmp-proxy	Configure the interface as an IGMP proxy interface.
ip igmp-proxy reset-status	(Optional) Reset the host interface status parameters of the IGMP Proxy.
ip igmp-proxy unsolicit-rprt-interval *seconds*	Configure the unsolicited report interval for the IGMP proxy interface. The range for *seconds* is 0–260 seconds.
CTRL + Z	Exit to Privileged EXEC mode.
show ip igmp-proxy	View a summary of the host interface status parameters.
show ip igmp-proxy interface	View a detailed list of the host interface status parameters. This command displays information only when IGMP Proxy is operational.
show ip igmp-proxy groups	View a table of information about multicast groups that IGMP Proxy reported. This command displays information only when IGMP Proxy is operational.

Configuring and Viewing MLD

Beginning in Privileged EXEC mode, use the following commands to configure MLD on the switch and on VLAN routing interfaces and to view IGMP information.

Command	Purpose
configure	Enter global configuration mode.
ipv6 mld router	Enable MLD on the switch.
interface vlan *vlan-id*	Enter Interface Configuration mode for the specified VLAN.
ipv6 mld router	Enable MLD on the interface.
ipv6 mld version *version*	Set the version of MLD for an interface.
	The *version* variable can be 1 or 2.
ipv6 mld query-interval *seconds*	Configure the query interval for the specified interface. The query interval determines how fast MLD Host-Query packets are transmitted on this interface.
	The range for *seconds* is 0–3600 seconds.
ipv6 mld query-max-response-time *seconds*	Configure the maximum response time interval for the specified interface. It is the maximum query response time advertised in MLD queries on this interface.
	The range for *seconds* is 0–25 seconds.
ipv6 mld last-member-query-interval *tenthsofseconds*	Set the last member query interval for the MLD interface, which is the value of the maximum response time parameter in the group-specific queries sent out of this interface.
	The range is 0–65535 milliseconds.
ipv6 mld last-member-query-count *count*	Set the number of listener-specific queries sent before the router assumes that there are no local members on the interface.
	The range for *count* is 1–20.
CTRL + Z	Exit to Privileged EXEC mode.

Command	Purpose
show ipv6 mld interface [**vlan** *vlan-id*]	View MLD information for all interfaces or for the specified interface.
show ip igmp interface stats [**vlan** *vlan-id*]	View MLD statistics for all interfaces or for the specified interface.
show ip igmp groups [**interface vlan** *vlan-id*]	View the registered multicast groups on the interface.
show ip igmp membership	View the list of interfaces that have registered in any multicast group.

Configuring and Viewing MLD Proxy

Beginning in Privileged EXEC mode, use the following commands to configure the upstream VLAN routing interface as an MLD proxy. The MLD proxy issues host messages on behalf of the hosts that have been discovered on the downstream MLD-enabled interfaces. The upstream interface is the interface closest to the root multicast router, which should be running IGMP.

NOTE: Configure only the upstream interface as the MLD proxy. MLD should be enabled on all downstream interfaces. IPv6 routing must be enabled on the switch for the MLD proxy feature to operate.

Command	Purpose
configure	Enter global configuration mode.
interface vlan *vlan-id*	Enter Interface Configuration mode for the specified VLAN.
ipv6 mld-proxy	Configure the interface as an MLD proxy interface.
ipv6 mld-proxy reset-status	(Optional) Reset the host interface status parameters of the MLD Proxy.
ipv6 igmp-proxy unsolicit-rprt-interval *seconds*	Configure the unsolicited report interval for the MLD proxy interface. The range for *seconds* is 0–260 seconds.
CTRL + Z	Exit to Privileged EXEC mode.

Command	Purpose
show ipv6 mld-proxy	View a summary of the host interface status parameters.
show ipv6 mld-proxy interface	View a detailed list of the host interface status parameters. This command displays information only when MLD Proxy is operational.
show ipv6 mld-proxy groups	View a table of information about multicast groups that MLD Proxy reported. This command displays information only when MLD Proxy is operational.

Configuring and Viewing PIM-DM for IPv4 Multicast Routing

Beginning in Privileged EXEC mode, use the following commands to configure PIM-DM for IPv4 multicast routing on the switch and on VLAN routing interfaces and to view PIM-DM information.

Command	Purpose
configure	Enter global configuration mode.
ip pim dense	Enable PIM-DM on the switch.
interface vlan *vlan-id*	Enter Interface Configuration mode for the specified VLAN.
ip pim	Enable PIM-DM on the interface.
ip pim hello-interval *seconds*	Specify the number of seconds (range: 0–65535) to wait between sending PIM hello messages on the interface.
exit	Exit to Privileged EXEC mode.
show ip pim	View system-wide PIM information.
show ip pim interface vlan *vlan-id*	View the PIM-DM information for the specified interface.
show ip pim neighbor [**interface vlan** *vlan-id* | **all**]	View a summary or all the details of the multicast table.

Configuring and Viewing PIM-DM for IPv6 Multicast Routing

Beginning in Privileged EXEC mode, use the following commands to configure PIM-DM for IPv6 multicast routing on the switch and on VLAN routing interfaces and to view PIM-DM information.

Command	Purpose
configure	Enter global configuration mode.
ipv6 pim dense	Enable PIM-DM on the switch.
interface vlan *vlan-id*	Enter Interface Configuration mode for the specified VLAN.
ipv6 pim	Enable PIM on the interface.
ipv6 pim hello-interval *seconds*	Specify the number of seconds (range: 0–65535) to wait between sending PIM hello messages on the interface.
exit	Exit to Privileged EXEC mode.
show ipv6 pim	View system-wide PIM information.
show ipv6 pim interface vlan *vlan-id*	View the PIM information for the specified interface.
show ipv6 pim neighbor [**interface vlan** *vlan-id* | **all**]	View a summary or all the details of the multicast table.

Configuring and Viewing PIM-SM for IPv4 Multicast Routing

Beginning in Privileged EXEC mode, use the following commands to configure PIM-SM for IPv4 multicast routing on the switch and on VLAN routing interfaces and to view PIM-SM information.

Command	Purpose
configure	Enter global configuration mode.
ip pim sparse	Enable PIM-SM as the multicast routing protocol on the switch.
ip pim spt-threshold *threshold*	Set the Data Threshold rate, in Kbps (range: 0–2000) for the last-hop (or leaf) router to switch to the shortest path.
ip pim register-threshold *threshold*	Set the rate, in Kbps (range: 0–2000), above which the Rendezvous Point router will switch to a source-specific shortest path tree.
ip pim bsr-candidate vlan *vlan-id hash-mask-length* [*priority*] [**interval** *interval*]	Configure the switch to announce its candidacy as a bootstrap router (BSR). • *vlan-id* — A valid VLAN ID. • *hash-mask-length* — The length of a mask that is to be ANDed with the group address before the hash function is called. All groups with the same seed hash correspond to the same RP. For example, if this value is 24, only the first 24 bits of the group addresses matter. This allows you to get one RP for multiple groups. (Range 0–32 bits). • *priority* — The priority of the candidate BSR. The BSR with the higher priority is preferred. If the priority values are the same, the router with the higher IP address is the BSR. (Range 0–255). • *interval* — (Optional) Indicates the BSR candidate advertisement interval. The range is from 1 to 16383 seconds. The default value is 60 seconds.

Command	Purpose
ip pim rp-candidate vlan *vlan-id group-address group-mask* [**interval** *interval*]	Configure the router to advertise itself to the BSR router as a PIM candidate Rendezvous Point (RP) for a specific multicast group range. • *vlan-id* — A valid VLAN ID. • *group-address* — Group IP address supported by RP. • *group-mask* — Group subnet mask for group address. • *interval* — (Optional) Indicates the RP candidate advertisement interval. The range is from 1 to 16383 seconds. The default value is 60 seconds.
ip pim rp-address *rp-address group-address group-mask* [**override**]	(Optional) Statically configure the RP address for one or more multicast groups. Only one RP address can be used at a time within a PIM domain The optional keyword **override** indicates that if there is a conflict, the RP configured with this command prevails over the RP learned by BSR.
ip pim ssm {**default** | *group-address group-mask* }	Define the Source Specific Multicast (SSM) range of IP multicast addresses. • **default** — Defines the SSM range access list to 232.0.0.0/8. • *group-address group-mask* — defines the SSM range.
interface vlan *vlan-id*	Enter Interface Configuration mode for the specified VLAN.
ip pim hello-interval *seconds*	Specify the number of seconds (range: 0–65535) to wait between sending PIM hello messages on the interface.
ip pim bsr-border	Prevent bootstrap router (BSR) messages from being sent or received through the interface.
ip pim dr-priority *priority*	Set the priority value for which a router is elected as the designated router (DR). The election priority range is 0–2147483647.
ip pim join-prune-interval *interval*	Configure the interface join/prune interval for the PIM-SM router. The interval range is 0–18000 seconds.

Command	Purpose
exit	Exit to Global Config mode.
exit	Exit to Privileged EXEC mode.
show ip pim	View system-wide PIM information.
show ip pim interface vlan *vlan-id*	View the PIM information for the specified interface.
show ip pim neighbor [**interface vlan** *vlan-id* | **all**]	View a summary or all the details of the multicast table.
show ip pim rp-hash *groupaddr*	View the RP router being selected for the specified multicast group address from the set of active RP routers. The RP router for the group is selected by using a hash algorithm.
show ip pim bsr-router [**candidate** | **elected**]	View the bootstrap router (BSR) information.
show ip pim rp mapping	View group-to-RP mappings of which the router is aware (either configured or learned from the BSR)

Configuring and Viewing PIM-SM for IPv6 Multicast Routing

Beginning in Privileged EXEC mode, use the following commands to configure PIM-SM for IPv6 multicast routing on the switch and on VLAN routing interfaces and to view PIM-SM information.

Command	Purpose
configure	Enter global configuration mode.
ipv6 pim sparse	Enable PIM-SM as the multicast routing protocol on the switch.
ipv6 pim spt-threshold *threshold*	Set the Data Threshold rate, in Kbps (range: 0–2000) for the last-hop (or leaf) router to switch to the shortest path.
ipv6 pim register-threshold *threshold*	Set the rate, in Kbps (range: 0–2000), above which the Rendezvous Point router will switch to a source-specific shortest path tree.

Command	Purpose
ipv6 pim bsr-candidate vlan *vlan-id hash-mask-length* [*priority*] [**interval** *interval*]	Configure the switch to announce its candidacy as a bootstrap router (BSR) • *vlan-id* — A valid VLAN ID. • *hash-mask-length* — The length of a mask that is to be ANDed with the group address before the hash function is called. All groups with the same seed hash correspond to the same RP. For example, if this value is 24, only the first 24 bits of the group addresses matter. This allows you to get one RP for multiple groups. (Range 0–32 bits). • *priority* — The priority of the candidate BSR. The BSR with the higher priority is preferred. If the priority values are the same, the router with the higher IPv6 address is the BSR. (Range 0–255). • *interval* — (Optional) Indicates the BSR candidate advertisement interval. The range is from 1 to 16383 seconds. The default value is 60 seconds.
ipv6 pim rp-candidate vlan *vlan-id group-address/prefix-length* [**interval** *interval*]	Configure the router to advertise itself to the BSR router as a PIM candidate Rendezvous Point (RP) for a specific multicast group range. • *vlan-id* — A valid VLAN ID. • *group-address/prefix-length*— Group IPv6 address and prefix length supported by RP. • *interval* — (Optional) Indicates the RP candidate advertisement interval. The range is from 1 to 16383 seconds. The default value is 60 seconds.
ipv6 pim rp-address *rp-address group-address/prefix-length* [**override**]	(Optional) Statically configure the RP address for one or more multicast groups. Only one RP address can be used at a time within a PIM domain The optional keyword **override** indicates that if there is a conflict, the RP configured with this command prevails over the RP learned by BSR.

Command	Purpose
ipv6 pim ssm {**default** | *group-address/prefix-length* }	Define the Source Specific Multicast (SSM) range of IPv6 multicast addresses. • **default** — Defines the SSM range access list to FF3x::/32. • *group-address/prefix-length* — defines the SSM range.
interface vlan *vlan-id*	Enter Interface Configuration mode for the specified VLAN.
ipv6 pim hello-interval *seconds*	Specify the number of seconds (range: 0–65535) to wait between sending PIM hello messages on the interface.
ipv6 pim bsr-border	Prevent bootstrap router (BSR) messages from being sent or received through the interface.
ipv6 pim dr-priority *priority*	Set the priority value for which a router is elected as the designated router (DR). The election priority range is 0–2147483647.
ipv6 pim join-prune-interval *interval*	Configure the interface join/prune interval for the PIM-SM router. The interval range is 0–18000 seconds.
exit	Exit to Global Config mode.
exit	Exit to Privileged EXEC mode.
show ipv6 pim	View system-wide PIM information.
show ipv6 pim interface vlan *vlan-id*	View the PIM information for the specified interface.
show ipv6 pim neighbor [**interface vlan** *vlan-id* | **all**]	View a summary or all the details of the multicast table.
show ipv6 pim rp-hash *groupaddr*	View the RP router being selected for the specified multicast group address from the set of active RP routers. The RP router for the group is selected by using a hash algorithm.
show ipv6 pim bsr-router	View the bootstrap router (BSR) information.
show ipv6 pim rp mapping	View group-to-RP mappings of which the router is aware (either configured or learned from the BSR)

Configuring and Viewing DVMRP Information

Beginning in Privileged EXEC mode, use the following commands to configure DVMRP on the switch and on VLAN routing interfaces and to view DVMRP information.

Command	Purpose
configure	Enter global configuration mode.
ip dvmrp	Enable DVMRP on the switch.
interface vlan *vlan-id*	Enter Interface Configuration mode for the specified VLAN routing interface.
ip dvmrp	Enable DVMRP on the interface.
ip dvmrp metric *metric*	Configure the metric (range: 1–31) for an interface. This value is used in the DVMRP messages as the cost to reach this network.
exit	Exit to Privileged EXEC mode.
show ip dvmrp interface vlan *vlan-id*]	View the multicast information for the specified interface.
show ip dvmrp neighbor	View neighbor information for DVMRP.
show ip dvmrp nexthop	View the next hop information on outgoing interfaces for routing multicast datagrams.
show ip dvmrp prune	View the table that lists the router's upstream prune information
show ip dvmrp route	View the multicast routing information for DVMRP.

L3 Multicast Configuration Examples

This section contains the following configuration examples:

- Configuring Multicast VLAN Routing With IGMP and PIM-SM
- Configuring DVMRP

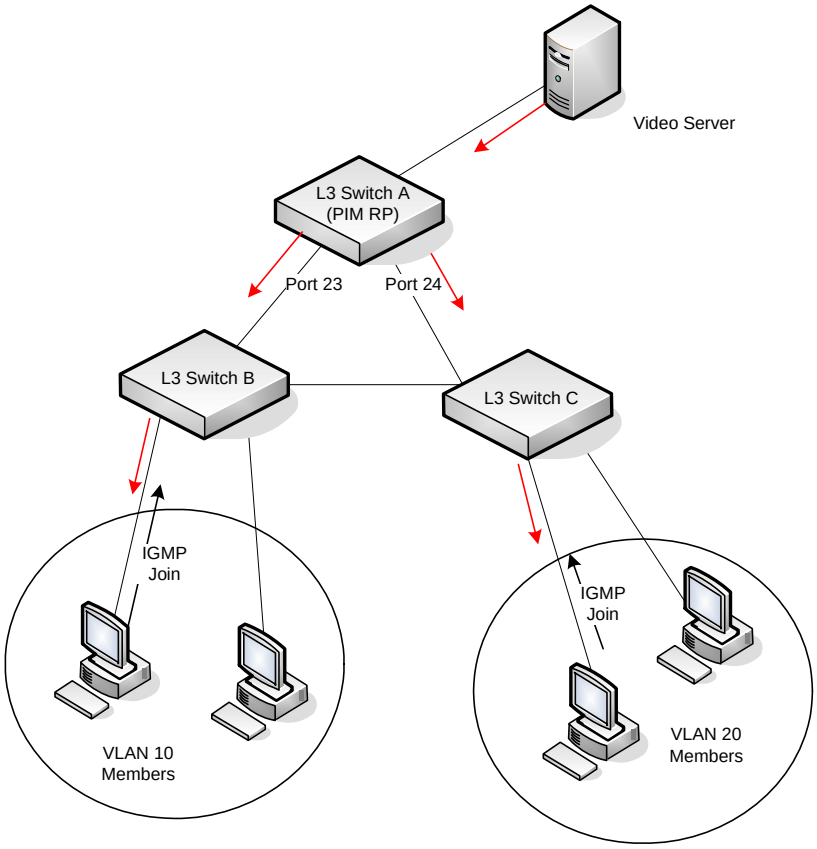
Configuring Multicast VLAN Routing With IGMP and PIM-SM

This example describes how to configure a PowerConnect switch with two VLAN routing interfaces that route IP multicast traffic between the VLANs. PIM and IGMP are enabled on the switch and interfaces to manage the multicast routing. IGMP snooping is enabled on the VLAN interfaces to control the multicast subscriptions within each VLAN. VLAN 10 is statically configured as the RP for the multicast group.

NOTE: PIM does not require OSPF specifically; static routing or RIP could also be configured for unicast routing.

The configuration in this example takes place on L3 switch A shown in Figure 43-47. The red arrows indicate the path that multicast traffic takes. L3 Switch A is configured as the RP for the PIM domain, so it is in charge of sending the multicast stream to L3 Switch B and L3 Switch C, and these switches forward the multicast data to the hosts that have requested to receive the data.

Figure 43-47. IPv4 Multicast VLAN Routing



In addition to multicast configuration, this example includes commands to configure STP and OSPF on L3 Switch A. STP is configured on the ports that connects the switch to other switches. OSPF is configured to route unicast traffic between the VLANs.

To configure the switch:

1 Create the two VLANs.

```
console#configure
console(config)#vlan database
```

```
console(config-vlan)#vlan 10,20
```

2 While in VLAN Database mode, enable IGMP snooping on the VLANs.

```
console(config-vlan)#ip igmp snooping 10
console(config-vlan)#ip igmp snooping 20
console(config-vlan)#exit
```

3 Configure port 23 and 24 as trunk ports.

```
console(config)#interface te1/0/23
console(config-if-Te1/0/23)#switchport mode trunk
console(config-if-Te1/0/23)#switchport trunk
allowed vlan add 20
console(config-if-Te1/0/23)#exit

console(config)#interface te1/0/24
console(config-if-Te1/0/24)#switchport mode trunk
console(config-if-Te1/0/24)#switchport trunk
allowed vlan add 10
console(config-if-Te1/0/24)#exit
```

4 Enable routing on the switch and configure the OSPF router ID.

```
console(config)#ip routing
console(config)#router ospf
console(config-router)#router-id 3.3.1.1
console(config-router)#exit
```

5 Configure VLAN 10 as a VLAN routing interface and specify the OSPF area. When you assign an IP address to the VLAN, routing is automatically enabled.

```
console(config)#interface vlan 10
console(config-if-vlan10)#ip address 192.168.10.4
255.255.255.0
console(config-if-vlan10)#ip ospf area 0
```

6 Enable IGMPv2 and PIM-SM on the VLAN routing interface.

```
console(config-if-vlan10)#ip igmp
console(config-if-vlan10)#ip igmp version 2
console(config-if-vlan10)#ip pim
console(config-if-vlan10)#exit
```

7 Configure VLAN 20 as a VLAN routing interface and specify the OSPF area.

```
console(config)#interface vlan 20
console(config-if-vlan20)#ip address 192.168.20.4
255.255.255.0
console(config-if-vlan20)#ip ospf area 0
```

8 Enable IGMPv2 and PIM-SM on the VLAN routing interface.

```
console(config-if-vlan20)#ip igmp
console(config-if-vlan10)#ip igmp version 2
console(config-if-vlan20)#ip pim
console(config-if-vlan20)#exit
```

9 Globally enable IGMP snooping, IP multicast, IGMP, and PIM-SM on the switch.

```
console(config)#ip igmp snooping
console(config)#ip multicast
console(config)#ip igmp
console(config)#ip pim sparse
```

10 Configure VLAN 10 as the RP and specify the range of multicast groups for PIM-SM to control.

```
console(config)#ip pim rp-address 192.168.10.4
225.0.0.0 240.0.0.0
```

Configuring DVMRP

The following example configures two DVMRP interfaces on the switch to enable inter-VLAN multicast routing.

To configure the switch:

1. Globally enable IP routing and IP multicast.

```
console#configure
console(config)#ip routing
console(config)#ip multicast
```

2. Globally enable IGMP so that this L3 switch can manage group membership information for its directly-connected hosts. Enabling IGMP is not required if there are no directly-connected hosts.

```
console(config)#ip igmp
```

3. Globally enable DVMRP.

```
console(config)#ip dvmrp
```

4. Enable DVMRP and IGMP on VLAN routing interfaces 10 and 20.

```
console(config)#interface vlan 10
console(config-if-vlan10)#ip address 192.168.10.1
255.255.255.0
console(config-if-vlan10)#ip dvmrp
console(config-if-vlan10)#ip igmp
console(config-if-vlan10)#exit

console(config)#interface vlan 20
console(config-if-vlan20)#ip address 192.168.20.1
255.255.255.0
console(config-if-vlan20)#ip dvmrp
console(config-if-vlan20)#ip igmp
console(config-if-vlan20)#exit
```

A

Feature Limitations and Platform Constants

- Table A-1 lists feature limitations.
- Table A-2 lists platform constants.

Table A-1. Feature Limitations

Feature	PC8000-series	PC8100-series
Base software features		
Link Dependency Max Groups	16	72
Switching features		
RMON 1, 2, 3, 9		
Max History entries	270	270
Max buckets per History entry	50	50
Max Alarm entries	32	32
Max Event entries	32	32
Max Log entries per Event entry	100	100
Management ACL (MACAL) Max Rules	64	64
Routing features		
IP Helper Max entries	512	512
User management features		
User ID configuration		
Max number of configured users	8	8
Max user name length	60	64
Max number of IAS users (internal user database)	50	50

Table A-1. Feature Limitations (Continued)

Feature	PC8000-series	PC8100-series
Authentication login list	Y	Y
Max Count	5	5
Max methods per list	6	6
Max name length	15	15
Authentication Enable lists	Y	Y
Max Count	5	5
Max methods per list	6	6
Max name length	15	15
Authentication HTTP lists	Y	Y
Max Count	1	1
Max methods per list	6	6
Max name length	15	15
Authentication Dot1x lists	Y	Y
Max Count	1	1
Max methods per list	6	6
Max name length	15	15
Accounting Exec lists	Y	Y
Max Count	5	5
Max methods per list	2	2
Max name length	15	15
Accounting Commands lists	Y	Y
Max Count	5	5
Max methods per list	1	1
Max name length	15	15
Login History	50	50

Table A-1. Feature Limitations (Continued)

Feature	PC8000-series	PC8100-series
QoS features		
iSCSI	Y	Y
Dot1p Marking	N	N
IP DSCP Marking	Y	Y
Max Monitored TCP Ports/IP Addresses	16	16
Max Sessions	192	192
Max Connections	192	192
Stacking features		
Max physical units per stack	6	12
Max physical slots per unit	1	1
Max physical ports per slot	24	56
Max physical ports per stack	144	672
Nonstop forwarding	Y	Y

Table A-2. Platform Constants

Feature	PC8000-series	PC8100-series
MAC addresses assigned per system	3	3
Reference CPU	PPC8533	NetLogic XLP308L
Reference CPU speed	1 GHz	1.2 GHz
Reference RAM	I Gbyte	2 Gbyte DDR3
Reference Flash	32 Mbyte	256 Mbyte
Maximum number of remote Telnet connections	4	4
Maximum number of remote SSH connections	4	4
Number of MAC addresses supported (stacking/nonstacking)	32768	131072
Number of VLANs	1024	1024
Maximum VLAN IDs	4093	4093
Number of 802.1p traffic classes	7	7
Number of IEEE 802.1x clients per port	16	16
Number of LAGs (max. lags/ports/max dynamic LAG ports per system	128/8/144	128/8/144
Maximum multiple spanning tree instances	15	15
Number of MAC-based VLANs supported	256	256
Number of records in log	400	400
Number of subnet-based VLANs supported	128	128
Protocol-based VLANs		
Max number of groups	128	128
Max protocols	16	16
Maximum MFDB entries	2048	2048
Jumbo frame support	Y	Y
Max size supported	9216	9216
Number of DHCP snooping bindings	8192	32768

Table A-2. Platform Constants (Continued)

Feature	PC8000-series	PC8100-series
Number of DHCP snooping static entries	1024	1024
LLDP-MED		
Number of remote nodes	288	816
LLDP Remote Management address buffers	388	916
LLDP Unknown TLV address buffers	100	100
LLDP Organizationally Defined TLV buffers	1872	5304
Port MAC locking		
Dynamic addresses per port	600	600
Static addresses per port	100	100
sFlow		
Number of samplers	32	52
Number of pollers	416	416
Number of receivers	8	8
RADIUS		
Max Authentication servers	32	32
Max Accounting servers	32	32
Number of routes (IPv4/IPv6)		
IPv4/IPv6 build		
IPv4 routes (–32 due to prefix len. sharing)	6112	8160
IPv6 routes	3072	4096
Number of static routes (IPv4/IPv6)	64/64	64/64

Table A-2. Platform Constants (Continued)

Feature	PC8000-series	PC8100-series
OSPF		
Max OSPFv2 LSAs		
IPv4-only build	N/A	36968
IPv4/IPv6 build	18536	24680
Max OSPFv3 LSAs	9416	12488
OSPFv3 max neighbors	400	400
OSPFv3 max neighbors per interface	100	100
Tunnels		
Number of configured v6-over-v3 tunnels	8	8
Number of automatic (6to4) tunnels	1	1
Number 6to4 next hops	16	16
DHCP server		
Max number of pools	16	16
Total max leases	256	256
DNS client		
Concurrent requests	16	16
Name server entries	8	8
Search list entries	6	6
Static host entries	64	64
Cache entries	128	128
Domain search list entries	32	32
DHCPv6		
Max number of pools	16	16
DNS domain names within a pool	5	5
DNS server addresses within a pool	8	8
Delegated prefix definitions within a pool	10	10
Number of VLAN routing interfaces	128	128

Table A-2. Platform Constants (Continued)

Feature	PC8000-series	PC8100-series
Number of ARP entries (Hosts)		
Device hardware capacity (v4/v6)	8192/4086	8192/4086
IPv4-only build	N/A	6144
IPv4/IPv6 build (v4/v6)	4096/2048	4096/1024
Static v4 ARP entries	64	128
Number of ECMP next hops per route	4	4
IP Multicast		
Number of IPv4/IPv6 Multicast Forwarding Entries	1024 (512 IPv4 and 256 IPv6)	1024 (512 IPv4 and 256 IPv6)
Hardware limit	2048	1024
IGMP Group Memberships per system	1024	2048
DVMRP Neighbors	256	256
PIM-DM Neighbors	256	256
PIM-SM Neighbors	256	256
PIM-SM Static RP entries	5	5
PIM-SM Candidate RP Group Range entries	20	20
PIM-SM SSM range entries	5	5
IGMP Sources processed per group per message	73	73

Table A-2. Platform Constants (Continued)

Feature	PC8000-series	PC8100-series
ACL limits		
Maximum number of ACLs (any type)	100	100
Maximum number of configurable rules per list.	127	1023
Maximum ACL Rules per Interface and Direction (IPv4/L2)	127	1023
Maximum ACL Rules per Interface and DIrection (IPv6)	125	1023
Maximum ACL Rules (system-wide)	4096	16384
Maximum ACL Logging Rules (system-wide)	128	128
CoS Device Characteristics		
Configurable Queues per port	7	7
Configurable Drop Precedence levels	3	3

Table A-2. Platform Constants (Continued)

Feature	PC8000-series	PC8100-series
DiffServ Device Limits		
Number of queues (stacking/nonstacking)	7/8	8
Requires LTLV to contain all policy instances combined	Y	Y
Max Rules per Class	13	13
Max Instances per Policy	12	28
Max Attributes per Instance	3	3
Max Service INterfaces (nonstacking/stacking)	58/624	116
Max table entries		
Class Table	32	32
Class Rule Table	416	416
Policy Table	64	64
Policy Instance Table	768	1792
Policy Attribute Table	2304	5376
Max Nested Class Reference Chain Rule Count	26	26
AutoVoIP number of voice calls	16	16
Voice VLAN number of devices	144	408

Index

Numerics

10GBase-T copper uplink module, 91

802.1p
see CoS queuing

A

AAA, 177

access lines, 199

access profiles, 59

accounting, 182

ACL
for controlling management access, 534

ACLs
Auto-Voip usage, 1152
binding configuration, 518
CLI configuration, 521
configuration steps, 506
defined, 501
examples, 529
iSCSI usage, 389
limitations, 505
logging, 504
management control, example, 536
preventing false matches, 506
supported types, 63
time-based, 533
web-based configuration, 508

ACLs. See also IP ACL, IPv6 ACL, and MAC ACL.

active images, 315

address table. See MAC address table.

administrative profiles, 181
defaults, 200
RADIUS authorization, 191
TACACS+ authorization, 188

alternate store and forward, 65

Amber, 95

ARP, 76
dynamic ARP inspection, 64

ARP inspection. see DAI.

ARP table
configuring (CLI), 900
configuring (web), 890

authentication, 179
examples, 183

authentication key, SNTP, 248

authentication profiles, 59

authentication server filter assignments, 491

authorization, 180
administrative profiles, 181

examples, 187
RADIUS, 191
auto configuration
auto save, 339
CLI configuration, 343
defaults, 341
defined, 333
DHCP, 344
configuration file, 337
image, 336
IP address, obtaining, 335
example, 344
files
USB, 334
files, managing, 339
stopping, 339
using DHCP, 335
web-based configuration, 342
auto image download
DHCP, 344
auto install, 56
auto install. See auto configuration.
auto negotiation, 67
auto save feature, 339
auto VoIP
CLI configuration, 1156
defaults, 1152
understanding, 1151
web-based configuration, 1153
auto-provisioning, iSCSI, 390
Auto-VoIP
and ACLs, 1152

B

back panel features, 92
back pressure, 66
banner, CLI, 264
baud rate, 93
BOOTP/DHCP relay agent, 77
BPDU
filtering, 74, 612
flooding, 612
guard, 74
protection, 614
bridge multicast address groups, configuring, 693
bridge multicast forwarding, 695
bridge multicast group table, 692
bridge table, 837
broadcast storm control. See storm control.
BSR, 1163

C

cable test, 203, 211
candidate BSR, 1164
captive portal, 62
CLI configuration, 429
client management, 434
configuring, 436
customizing pages, 405
defaults, 406

defined, 401
dependencies, 402
design considerations, 403
example, 435
localization, 405
understanding, 401, 404
user logout mode, 405
users, RADIUS server, 417
web-based configuration, 408

cards
supported, 257

CDP, interoperability through ISDP, 57

certificates, 310

checkpointing, 149

Cisco protocol filtering, 70

CLI
accessing the switch, 107
banner, 237
command completion, 112
command modes, 109
command prompt, 238
error messages, 113
negating commands, 112

CLI banner, configuring, 264

clock, system, 246

command modes, CLI, 109

commands
abbreviated, 112
entering, 111
history buffer, 113

Compellent storage arrays, 391

configuration file
defined, 309
DHCP auto configuration, 337
downloading, 312
editing, 312
SNMP, 313
USB device, 331

configuration scripts, 312, 329

configuration, saving the, 313

Configuring, 859

console port
connecting to, 107
description, 93

copy, files, 321

CoS
and iSCSI, 388
and PFC, 801
CLI configuration, 1144
configuration example, 1147
defaults, 1136
defined, 1133
queue management methods, 1135
traffic queues, 1135
traffic shaping, 1134
trusted mode ports, 1134
untrusted mode ports, 1134
web-based configuration, 1137

CoS queuing

D

DAI

defaults, 749
optional features, 748
purpose, 749
understanding, 748
data center
and DHCP snooping, 777
and NSF, 170
SDM template, 239
data center bridging exchange, 69
Data Center Bridging Exchange protocol, 807
date, setting, 262
daylight saving time, 238
DCBX
and iSCSI, 391
DCBx, 807
default gateway, configuring, 121, 127
default VLAN, 134
DHCP client, 132
IP address configuration, 125
denial of service, 61, 499
device discovery protocols, 638
device view, 106
DHCP
understanding, 859
DHCP auto configuration
dependencies, 340
enabling, 344
monitoring, 339
process, 335
DHCP client, 861
default VLAN, 132
OOB port, 132
DHCP relay, 70, 861
CLI configuration, 925
defaults, 913
example, 929
layer 2, 908
layer 3, 907
understanding, 907
VLAN, 909
web-based configuration, 914
DHCP server, 55
address pool configuration, 878
CLI configuration, 874
defaults, 861
examples, 878
leases, 133
options, 860
web-based configuration, 862
DHCP snooping, 64, 861
bindings database, 745
defaults, 749
example, 777
logging, 746
purpose, 749
understanding, 744
VLANs, 746
DHCPv6
client, 1068
defined, 79
examples, 1099
pool, 1084
prefix delegation, 1084

relay agent, configuring, 1101
relay agent, understanding, 1084
stateless server
configuring, 1099
stateless server,
understanding, 1084
understanding, 1083
dhcpv6, 1083
DHCPv6 pool
stateless server support, 1095
DHCPv6 relay
CLI configuration, 1095
defaults, 1085
web-based configuration, 1086
DHCPv6 server
CLI configuration, 1095
prefix delegation, 1100
web-based configuration, 1086
DHCPv6 server relay
defaults, 1085
DiffServ
and 802.1X, 465
and RADIUS, 465
and switch role, 1104
CLI configuration, 1119
defaults, 1105
elements, 1104
example, 1126
understanding, 1103
VoIP, 1130
web-based configuration, 1106
diffServ, 80
discovery, device, 637
document conventions, 50
domain name server, 128
domain name, default, 129
Dot1x, 62
dot1x
authentication, 180
double-VLAN tagging, 544
downloading files, 317
DSCP value and iSCSI, 388
dual images, 55
dual IPv4 and IPv6 template, 239
duplex mode, 86
DVMRP, 83
defaults, 1167
example, 1235
understanding, 1165
web-based configuration, 1209
when to use, 1166
dynamic ARP inspection, 64
dynamic LAGs, 797
dynamic VLAN creation, 490

E

EAP statistics, 360
email alert
statistics, 224
email alerting, 233
log messages, 230
enable authentication, 180

enhanced transmission selection, 69, 819
EqualLogic and iSCSI, 390
error messages, CLI, 113
EtherType numbers, common, 506
exec authorization, 181

F

failover, 59
failover, stacking, 148
false matches, ACL, 506
FC map value, 815
FCoE
  configuring CoS queues for, 1148
  frames, forwarding, 815
FCoE Initialization Protocol, 814
FCoE initialization protocol snooping, 69
FCoE, FC map value, 815
file management
  CLI, 322
  considerations, 311
  copying, 321
  purpose, 309
  supported protocols, 311
  web-based, 314
file system, 314
files
  downloading to the switch, 311
  types, 307
  uploading from the switch, 311
filter assignments, authentication server, 491
filter, DiffServ, 465
FIP snooping, 814
  enabling and disabling, 815
firmware
  managing, 311
  updating the stack, 147
  upgrade example, 326
firmware synchronization, stacking, 147
Flashing Green, 96-97
flow control
  configuring, 676
  default, 668, 800
  understanding, 666
flow-based mirroring, 1118
forwarding database, 837
  and port security, 747
front panel features, 85

G

GARP, 686
GARP and GVRP, 71
GMRP, 686
Green, 96
guest VLAN, 462

VLAN
guest, 489
GVRP, 544
statistics, 359

H

head of line blocking
prevention, 65
health, system, 209
help, accessing web-based, 111
host name, 237
host name mapping, 120

I

IAS
database, 474
understanding, 465
users, 481
icons, web-based interface, 104
identification
asset tag, 237
system contact, 237
system location, 237
system name, 237
IDSP
defaults, 639
IEEE 802.1d, 73
IEEE 802.1Q, 71
IEEE 802.1Qaz, 808
IEEE 802.1X, 62
and DiffServ, 465
authentication, 62
configuring, 481
defined, 458
monitor mode, 63, 463, 478
port authentication, 476
port states, 459
RADIUS-assigned VLANs, 479
reauthenticating ports, 469
VLAN assignment, 461
IEEE 802.1x
authentication, 180
IEEE 802.3x. See flow control.
IGMP, 83
defaults, 1167
understanding, 1161
web-based configuration, 1177
IGMP proxy, 83, 1161
IGMP snooping, 81
defaults, 689
querier, 82
querier, defined, 683
understanding, 683
image
activating, 322
auto configuration, 336
considerations, 311
defined, 307
downloading, 322
management, CLI, 322
management, web-based, 314
purpose, 309
in-band management, 121

interface, 843
configuration mode, 442
loopback, 844
OOB, 124
routing, 843
CLI configuration, 855
web configuration, 849
routing defaults, 848
supported types, 442
tunnel, 845
internal authentication server, see IAS
IP ACL
configuration, 508
defined, 503
example, 529
IP address
configuring, 121
default, 123
default VLAN, 125, 134
OOB port, 134
IP helper, 77, 909
IP multicast traffic
layer 2, 682
layer 3, 1158
IP protocol numbers, common, 507
IP routing
CLI configuration, 899
defaults, 885
example, 904
understanding, 883
web-based configuration, 887
IP source guard, 64
IPSG
and port security, 747
example, 779
purpose, 749
understanding, 747
IPv4 and IPv6 networks, interconnecting, 1005
IPv4 multicast
web-based configuration, 1169
IPv4 routing template, 239
IPv6
compared to IPv4, 1060
DHCP client, 1068
DHCPv6, 79
OSPFv3, 79
routes, 79
static reject and discard routes, 1080
tunnel, 78
IPv6 ACL configuration, 515
IPv6 interface
configuring, 1060
IPv6 management, 55
IPv6 multicast
web-based configuration, 1176
IPv6 routing
CLI configuration, 1073
defaults, 1061
features, 78
understanding, 1059
web-based configuration, 1063
IRDP, configuring, 901

iSCSI
ACL usage, 389
and Compellent storage arrays, 391
and CoS, 388
and DCBX, 391
and Dell EqualLogic arrays, 390
assigning flows, 388
CLI configuration, 397
defaults, 392
examples, 399
flow detection, 388
information tracking, 389
servers and a disk array, 399
understanding, 387
using, 388
web-based configuration, 393
ISDP
and CDP, 57
CLI configuration, 656
configuring, 657
enabling, 657
example, 661
understanding, 637
web-based configuration, 641

J

jumbo frames, 66

L

LACP, 74
adding a LAG port, 791
CLI configuration, 796
web-based configuration, 788
LAG
and STP, 784
CLI configuration, 793
defaults, 785
examples, 797
guidelines, configuration, 785
interaction with other features, 784
LACP, 74
purpose, 782
static and dynamic, 782
statistics, 374
threshold, minimum links, 793
understanding, 781
web-based configuration, 786
LAG hashing, 783
languages, captive portal, 405
LED
100/1000/10000Base-T port, 95
port, 92
SFP port, 95
system, 96
Link, 95
link aggregation group. See LAG.
link dependencies
CLI configuration, 452
creating, 448
example, 455
scenarios, 441
understanding, 440
web configuration, 448

link local protocol filtering, see LLPF
LLDP
CLI configuration, 656
defaults, 639
example, 662
understanding, 637
web-based configuration, 641
LLDP-MED
and voice VLANs, 547
configuring, 660
understanding, 638
viewing information, 661
LLPF
defaults, 668, 800
example, 679
understanding, 667
localization, captive portal, 405
log messages, 54
log server, remote, 217
logging
ACL, 504
CLI configuration, 225
considerations, 207
defaults, 207
destination for log messages, 204
example, 232
file, 216
log message format, 206
operation logs, 205
severity levels, 205
system startup logs, 205
trap log, 292
web-based configuration, 208
loopback, 78
loopback interface
configuring, 857
purpose, 847
understanding, 844
LSA, OSPF, 933

M

M6348 and stacking, 144
MAC ACL
example, 531
understanding, 502
MAC address table
and port security, 747
contents, 838
defaults, 838
defined, 837
dynamic, 841
managing, CLI, 842
populating, 837
stacking, 838
web-based management, 839
MAC multicast support, 81
MAC port locking, 496
MAC-based 802.1X authentication
understanding, 460
MAC-based VLAN, 542
mail server
adding, 220
configuring, 229

email alert, 220
management
access control list, 534
access control using RADIUS, 192
access control using TACACS+, 197
management access list, example, 536
management, in-band and out-of-band, 121
MD5, 240
MDI/MDIX, auto, 66
MIB, SNMP, 271
Microsoft Network Load Balancing, 1166
mirror, ACL, 503
mirroring, flow-based, 1118
MLD, 84
defaults, 1167
understanding, 1162
web-based configuration, 1186
MLD snooping, 82
configuring, 730
defaults, 689, 749
understanding, 683
VLAN configuration, 731
mode
interface configuration, 442
monitor mode, IEEE 802.1X, 463
monitoring system information, 203
MSTP
example, 635
operation in the network, 607
support, 73
understanding, 605
MTU, configuring, 452
MTU, management interface, 122
Multicast
VLAN registration, 82
multicast
DVMRP, 83
IGMP, 83
IGMP proxy, 83
IGMP snooping, 81
IPv4, 1169
layer 2, 81
configuring (CLI), 725
configuring (web), 691
defaults, 689
understanding, 681
when to use, 685
layer 3, 83
CLI configuration, 1215
defaults, 1167
examples, 1231
understanding, 1157
when to use, 1160
MAC layer, 81
MLD snooping, 82
protocols
roles, 1159-1160

VLAN Routing with IGMP and PIM-SM, 1231
multicast bridging, 681, 725
multicast protocols, supported, 1159
multicast routing table, 1160
multicast snooping, 736
multicast tunneling, 1161
multicast VLAN registration, 684
MVR
adding an interface, 717

N

netinfo, 119
network information
CLI configuration, 132
default, 123
defined, 119
example, 136
purpose, 120
web-based configuration, 124
network pool, DHCP, 865
nonstop forwarding, see NSF
NSF
and DHCP snooping, 172
and routed access, 175
and the storage access network, 173
and VoIP, 171
in the data center, 170
network design considerations, 151
understanding, 147

O

OOB port, 92-93, 124
DHCP client, 132
OpenManage Switch Administrator, about, 101
optical transceiver diagnostics, 212
OSPF, 76
areas, 932
border router, 996
CLI configuration, 976
defaults, 940
difference from OSPFv3, 933
examples, 996
flood blocking, 938, 1013
LSA pacing, 937
NSSA, 999
static area range cost, 936, 1008
stub area, 999
stub routers, 934
topology, 932
trap flags, 290
understanding, 932
web-based configuration, 942
OSPFv3, 79
CLI configuration, 987
difference from OSPF, 933
global settings, 987
interface settings, 989

NSSA, 999
stub area, 999
trap flags, 291
web-based configuration, 959
out of band port, IP address, 134
out-of-band management, 121

P

password
protecting management access, 60
strong, 60
PFC, 801
PIM
defaults, 1167
IPv4 web-based configuration, 1196
IPv6 web-based configuration, 1196
PIM-DM, using, 1164
PIM-SM, using, 1163
SSM range, 1205
understanding, 1163
port
access control, 471
characteristics, 439
configuration examples, 454
configuring multiple, 446
defaults, 444
defined, 439
device view features, 106
example, 454
LEDs, 92
locking, 496
OOB, 92-93
protected, 65, 672, 677
statistics, 373
traffic control, 665, 799
USB, 92
port channel. See LAG.
port characteristics
CLI configuration, 451
web-based configuration, 445
port control, 470
port fast, STP, 612
Port LEDs, 95
port mirroring
configuring, 375
mode, enabling, 352
understanding, 351
port security
configuring, 498
MAC-based, 63
understanding, 495
port-based traffic control
CLI configuration, 676, 807
web-based configuration, 669
port-based VLAN, 542
port-MAC locking, 63
see port security
Ports, 49
power supplies, 93
priority flow control, 801

priority-based flow control, 68, 801
private VLAN edge, 65
private VLANs, 547, 602
protected port
  defined, 667
  example, 679
protocol filtering, Cisco, 70
protocol-based VLAN, 542

Q

QoS
  CoS queuing
  diffserv, 80
QSFP module, 91
queues, CoS, 1135

R

RADIUS, 60
  and DiffServ, 465
  authentication, 186
  authorization, 191
  for management access control, 192
  supported attributes, 194
  understanding, 192
RAM log, 215
real time clock, 238
Red, 96
redirect, ACL, 503
relay agent, DHCPv6, 1084
relay, DHCP, 907
remote logging, 228
rendezvous point, PIM, 1163
RIP, 77
  CLI configuration, 1027
  defaults, 1021
  determining route information, 1019
  example, 1031
  supported versions, 1020
  understanding, 1019
  web-based configuration, 1022
RMON, 57
  CLI management, 378
  defaults, 352
  example, 386
  understanding, 350
  web-based configuration, 353
router discovery, 901
router discovery protocol, 77
router, OSPF, 933
routes
  IPv4, 897
  IPv6, 1072
  selecting, 933
Routing
  table, 77
routing
  defaults (IPv4), 885
  defaults (IPv6), 1061

example, 904
IPv4, CLI configuration, 899
IPv4, web-based configuration, 887
IPv6, CLI configuration, 1073
IPv6, web-based configuration, 1063
understanding, 883
routing interfaces
CLI configuration, 855
defaults, 848
understanding, 843
using, 846
web-based configuration, 849
routing table
best routes, 894
configuring, 902
IPv6, 1077, 1079
RSTP
understanding, 605
running-config, saving, 313

S

save, system settings, 313
SDM template
configuration guidelines, 239
managing, 260
understanding, 239
SDM templates, 56
security
port, defined, 495
port-based
CLI configuration, 476
defaults, 466, 495
examples, 481
web-based configuration, 467
sFlow, 57
CLI management, 378
defaults, 352
example, 384
understanding, 347
web-based management, 353
SFP port LEDs, 95
SFP+ module, 91
SFTP, managing files, 325
SNMP
CLI configuration, 293
defaults, 273
examples, 302
MIB, 271
purpose, 273
traps, 272
understanding, 271
uploading files, 313
web-based configuration, 275
SNMPv1 example, 302
SNMPv2 example, 302
SNMPv3
engine ID, 293
example, 303
snooping,FIP, 814
SNTP
authentication, 260
authentication key, 248

example, 267
server, 260
server configuration, 251
understanding, 240
software image, 307
spanning tree. See STP.
Speed, 95
split horizon, 1020
SSH files, 310
SSH/SSL, 61
SSL files, 310
SSM range, 1205
stacking
adding a switch, 145
and NSF, 59
CLI configuration, 162
defaults, 152
defined, 141
design consideration, 150
failover, 59
failover, example, 166
failover, initiating, 148
features, 58
firmware synchronization, 147
firmware update, 147
MAC address table, 838
MAC addresses, 150
NSF usage scenario, 164
preconfiguration, 168
purpose, 151
removing a switch, 146
standby, 147
switch compatibility, 144
web-based configuration, 153
static reject route, 884
statistics
IPv6, 1066
statistics, Etherlike, 358
storage arrays and iSCSI, 390
storage arrays, Compellent, 391
storm control
configuring, 676
default, 668, 800
example, 679
understanding, 666
STP
and LAGs, 784
classic, 605
CLI configuration, 628
defaults, 615
defined, 605
examples, 633
loop guard, 613
MSTP, 73
optional features, 612
port fast, 612
port settings, 73
root guard, 613
RSTP, 73
understanding, 606
web-based configuration, 616
subnet mask, configuring, 121
subnet-based VLAN, 542
summer time, 238
switchport modes, VLAN, 542

switchport statistics, web view, 363
system description, 49
system health, monitoring, 208
system information
CLI configuration, 258
default, 240
defined, 237
example, 264
purpose, 238
web-based configuration, 241
system LEDs, 92
system time, 240

T

TACACS+, 60
authentication, 184
authorization, 187-188
management access control, 197
supported attributes, 198
understanding, 197
tagging, VLAN, 543
telnet
configuration options, 61
connecting to the switch, 108
TFTP, image download, 322
Thermal LEDs, 96
time domain reflectometry, 211
time management, 54
time range, 527
time zone, 255
time, setting the system, 269
time-based ACLs, 504, 533
traffic class queue, 388
traffic control
port based, 665, 799
traffic inspection, 743
traffic monitoring, 347
traffic snooping, 743
traps
OSPF, 290
trunk port
and 802.1X authentication, 489, 491
trunking, 579
tunnel, 78
tunnel interfaces, 845

U

UDP relay, 77, 909
upgrade, stack firmware, 58
uploading files, 319
USB auto configuration
files, 334
understanding, 334
USB flash drive, example, 331
USB port, 92
user security model, SNMP, 272

users
authenticated, 470
captive portal, 414
IAS database, 465
USM, 272

V

ventilation system, 94
virtual link, OSPF, 1003
VLAN, 784
authenticated and unauthenticated, 461
CLI configuration, 576
defaults, 554
double, 72
double-VLAN tagging, 544
dynamic, 462
example, 592, 597
guest, 72, 462, 490
IP subnet-based, 71
MAC-based, 71, 542
port-based, 70, 542
private, 547, 602
protocol-based, 71, 542
RADIUS-assigned, 490
routing, 76
routing interfaces, 843, 855
static, 542
support, 70
switchport modes, 542
trunk port, 579
understanding, 539
voice, 71, 546
voice traffic, 546
voice, example, 601
voice, understanding, 545
web-based configuration, 556
VLAN membership, defining, 556
VLAN priority tag and iSCSI, 388
VLAN routing, 843, 846
VLAN tagging, 543
VLANs
dynamically created, 490
RADIUS-assigned, 490
voice traffic, identifying, 546
voice VLAN, 546
and LLDP-MED, 547
example, 601
understanding, 545
VoIP, 80
VoIP and DiffServ, 1130
VoIP, auto, 1151
VRRP, 78
accept mode, 1037
CLI configuration, 1048
defaults, 1039
example, 1050
interface tracking, 1037
load sharing example, 1050
preemption, 1036
route and interface tracking example, 1054
route tracking, 1037
router priority, 1036

understanding, 1035
web-based configuration, 1040

W

web-based configuration, 102
web-based interface,
understanding, 103
writing to memory, 313