SonicWall® SonicOS 6.5 System Setup

Administration

SONICWALL®

Contents

Part 1. System Setup | About System Setup

About Setting Up Your SonicOS System 16
About the SonicOS Management Interface 16

Part 2. System Setup | Appliance

Configuring Base Settings 18
About Appliance | Base Settings 19
One-Click Wireless and Non-Wireless Controller Modes 19
Configuring the Appliance 21
Accessing the Appliance | Base Settings Page 21
Configuring the Firewall Name 21
Changing the Administrator Name & Password 22
Enabling Wireless LAN and IPv6 23
Configuring Login Security 24
Configuring Multiple Administrator Access 28
Enabling Enhanced Audit Logging Support 32
Configuring the Wireless LAN Controller 32
Configuring the Management Interface 34
Configuring the Front-Panel Administrative Interface (SuperMassive Firewalls Only) 40
Configuring Client Certificate Verification 40
Checking Certificate Expiration 44
Configuring SSH Management 44
Enabling SonicOS API 44
Configuring Advanced Management Options 45
Downloading SonicPoint Images Manually 48
Selecting a Language 49

Administering SNMP 50
About **Appliance > SNMP** **50**
About SNMP 50
Setting Up SNMP Access 51
Configuring SNMP as a Service and Adding Rules 59
About SNMP Logs 60

Managing Certificates 61
About Certificates 61
About Digital Certificates 61
About the Certificates and Certificate Requests Table 63
Importing Certificates 65
Deleting a Certificate 67
Generating a Certificate Signing Request 67
Configuring Simple Certificate Enrollment Protocol 71

Configuring Time Settings 72
About **Appliance > Time 72**
Setting System Time 73
Configuring NTP Settings 74

Setting Schedules 76
About Schedules 76
About **Appliance > Schedules 77**
Adding a Custom Schedule 78
Modifying Schedules 79
Deleting Custom Schedules 80

Part 3. System Setup | User Management

About Managing Users 83
About User Management 83
Using Local Users and Groups for Authentication 84
Using RADIUS for Authentication 87
Using LDAP/Active Directory/eDirectory Authentication 87
Using TACACS+ 91
About Single Sign-On 92
Installing the Single Sign-On Agent and/or Terminal Services Agent 102
About Multiple Administrator Support 112
Configuring Multiple Administrator Support 114

Configuring Settings for Managing Users 116
Users > Settings 116
Configuring User Authentication and Login Settings 117
Configuring User Sessions 128
Customization 132
Accounting 139
Configuring RADIUS Authentication 142
Configuring the SonicWall for LDAP 149
About LDAP Test 165
Configuring TACACS + for Authentication 166
Configuring SonicOS to Use the SonicWall SSO Agent 167

Managing Authentication Partitions 190
About Authentication Partitioning 190
About User Authentication Partitioning 191
About Subpartitions 192
About Inter-Partition User Roaming 194
About Authentication Partition Selection 195
About Extended Support for Multiple LDAP Servers 198
Per-Partition DNS Servers and Split DNS 198
About RADIUS Authentication 198
Upgrading from a Non-Partitioned Configuration 198
Configuring Authentication Partitions and Policies 199
Displaying and Filtering Users/Partitions 199

Configuring and Managing Partitions 201
Configuring Partition Selection Policies 215
Configuring Servers, Agents, and Clients for Authentication Partitioning 219

Configuring Local Users and Groups 221
About Authentication and Passwords 221
Using Two-Factor Authentication 221
Enforcing First Login Password Change 222
Configuring Local Users 222
Quota Control for all Users 223
Viewing Local Users 223
Adding Local Users 224
Editing Local Users 230
Configuring Global Settings 231
Importing Local Users from LDAP 232
Configuring a Guest Administrator 232
Configuring Local Groups 234
Creating or Editing a Local Group 235
Importing Local Groups from LDAP 241
Setting User Membership by LDAP Location 241

Managing Guest Services 242
Users > Guest Services 242
Global Guest Settings 243
Guest Profiles 243

Managing Guest Accounts 247
Users > Guest Accounts 247
Viewing Guest Account Statistics 248
Adding Guest Accounts 250
Enabling Guest Accounts 255
Enabling Auto-Prune for Guest Accounts 255
Editing Guest Accounts 256
Deleting Guest Accounts 256
Printing Account Details 257

Part 4. System Setup | Network

Configuring Interfaces 259
About Interfaces 260
Physical Interfaces 260
Virtual Interfaces 262
SonicOS Secure Objects 264
Transparent Mode 264
IPS Sniffer Mode 264
Firewall Sandwich 266
HTTP/HTTPS Redirection 267
Enabling DNS Proxy on an Interface 267
LTE Modem Support 267

LAN Bypass 268
Network > Interfaces 268
Show/Hide PortShield Interfaces (IPv4 Only) 269
Interface Settings 270
Interface Traffic Statistics 271
Configuring Interfaces 272
Configuring a Static Interface 272
Configuring Routed Mode 279
Enabling Bandwidth Management on an Interface 281
Configuring Interfaces in Transparent IP Mode (Splice L3 Subnet) 282
Configuring Wireless Interfaces 286
Configuring a WAN Interface 292
Configuring Tunnel Interfaces 297
Configuring Link Aggregation and Port Redundancy 301
Configuring Virtual Interfaces (VLAN Subinterfaces) 305
Configuring IPS Sniffer Mode 306
Configuring Security Services (Unified Threat Management) 309
Configuring Wire and Tap Mode 310
Wire Mode with Link Aggregation 313
Layer 2 Bridged Mode 315
Configuring Layer 2 Bridged Mode 334
Asymmetric Routing 341
Configuring Interfaces for IPv6 342
31-Bit Network 342
PPPoE Unnumbered Interface Support 344

Configuring PortShield Interfaces 347
Network > PortShield Groups 347
About PortShield 347
SonicOS Support of X-Series/N-Series Switches 348
SonicOS Support of N-Series Switches 356
Managing Ports 361
Configuring PortShield Groups 370

Configuring PoE 375
Network > PoE Settings 375
Enabling PoE 376
Configuring Generic PoE Settings 377
Configuring Power Settings for a Port 378

Setting Up Failover and Load Balancing 379
Network > Failover & Load Balancing 379
About Failover and Load Balancing 379
How Failover and Load Balancing Work 380
Multiple WAN (MWAN) 381
Network > Failover & Load Balancing 382
Configuring Failover and LB Groups 384
Configuring Probe Settings for Group Members 388

Configuring Network Zones 390
About Zones 390
How Zones Work 391
Predefined Zones 392
Security Types 392
Allow Interface Trust 393
Enabling SonicWall Security Services on Zones 393
Effect of Wireless and Non-Wireless Controller Modes 394
Network > Zones 396
The Zone Settings Table 396
Adding a New Zone 397
Configuring a Zone for Guest Access 400
Configuring a Zone for Open Authentication and Social Login 402
Configuring a Zone for Captive Portal Authentication with Radius 403
Configuring a Zone for a Customized Policy Message 405
Configuring a Zone for a Customized Login Page 406
Configuring the WLAN Zone 407
Configuring the RADIUS Server 409
Configuring DPI-SSL Granular Control per Zone 410
Enabling Automatic Redirection to the User-Policy Page 411
Deleting a Zone 412

Configuring Wire Mode VLAN Translation 413
Network > VLAN Translation 413
About VLAN Translation 413
Creating and Managing VLAN Maps 415

Configuring DNS Settings 421
Network > DNS 421
About Split DNS 421
Managing DNS Servers 423
DNS and IPv4 432

Configuring DNS Proxy Settings 435
Network > DNS Proxy 436
About DNS Proxy 437
Enabling DNS Proxy 439
Configuring DNS Proxy Settings 440
Monitoring DNS Server Status 441
Monitoring Split DNS Server Status 442
Viewing and Managing Static DNS Cache Entries 442
Viewing DNS Proxy Cache Entries 444

Configuring DNS Security 445
About Sinkholes 445
Network > DNS Security 445
Configuring DNS Security Settings 446
Deleting Entries in the Lists 447
About DNS Tunneling Detection 447

Configuring Route Advertisements and Route Policies **450**
About Routing 451
About Metrics and Administrative Distance 451
Route Advertisement 452
ECMP Routing 453
Policy-based Routing 453
Policy-based TOS Routing 453
PBR Metric-based Priority 454
Policy-based Routing and IPv6 455
OSPF and RIP Advanced Routing Services 455
Drop Tunnel Interface 463
App-Based Routing 464
Network > Routing **464**
Network > Routing > Settings 464
Network > Routing > Route Policies 465
Network > Routing > Route Advertisement 466
Network > Routing > OSPFv2 467
Network > Routing > RIP 469
Network > Routing > OSPFv3 470
Network > Routing > RIPng 472
Configuring Routing 472
Prioritizing Routes by Metric 473
Configuring Metrics for Default Routes Learned through Router Advertisement 473
Configuring Route Advertisement 474
Configuring Static and Policy-based Routes 475
Configuring a Static Route for a Drop Tunnel Interface 479
Configuring OSPF and RIP Advanced Routing Services 481
Configuring BGP Advanced Routing 490
Configuring App-based Routes 491

Managing ARP Traffic **492**
Network > ARP **492**
Static ARP Entries 493
ARP Settings 496
ARP Cache 496

Configuring Neighbor Discovery Protocol **498**
Network > Neighbor Discovery (IPv6 Only) 498
Static NDP Entries 499
NDP Settings 500
NDP Cache 500
Configuring a Static NDP Entry 501
Editing a Static NDP Entry 501
Flushing the NDP Cache 502

Configuring MAC-IP Anti-spoof **503**
About MAC-IP Anti-spoof Protection 503
Extension to IP Helper 504

Network > MAC-IP Anti-spoof 504
Settings for Interface(s) 505
Anti-Spoof Cache 507
Spoof Detected List 508
Configuring MAC-IP Anti-spoof Protection 508
Displaying Traffic Statistics 509
Editing MAC-IP Anti-spoof Settings for an IPv6 Interface 509
Editing MAC-IP Anti-spoof Settings for an IPv4 Interface 510
Adding Devices to Anti-Spoof Cache 512
Deleting Anti-Spoof Cache Entries 512
Filtering What Is Displayed 513
Adding Static Entries from Spoof Detected List 513

Setting Up the DHCP Server 514
Network > DHCP Server 514
DHCP Server Options Feature 516
Multiple DHCP Scopes per Interface 517
About DHCP Server Persistence 519
Configuring the DHCP Server 519
DHCP Server Lease Scopes 521
Current DHCP Leases 521
DHCPv6 Relay 523
Configuring Advanced Options 523
Configuring DHCP Server for Dynamic Ranges 528
Configuring Static DHCP Entries 534
Configuring DHCP Generic Options for DHCP Lease Scopes 536
RFC-Defined DHCP Option Numbers 536
DHCP and IPv6 540

Using IP Helper 541
About IP Helper 541
VPN Tunnel Interface Support for IP Helper 542
DHCPv6 Relay 543
Network > IP Helper 545
Relay Protocols 546
Policies 547
DHCP/DHCPv6 Relay Leases 547
Configuring IP Helper 548
Enabling IP Helper 548
Managing Relay Protocols 548
Managing IP Helper Policies 550
Filtering What DHCP Relay Leases are Displayed 552
Displaying IP Helper Cache from TSR 552

Setting Up Web Proxy Forwarding 554
Network > Web Proxy 554
Configuring Automatic Proxy Forwarding (Web Only) 555
Configuring User Proxy Servers 556

Configuring Dynamic DNS 558
Network > Dynamic DNS 558
About Dynamic DNS 558
Supported DDNS Providers 559
Dynamic DNS Profiles Table 560
Configuring a Dynamic DNS Profile 561
Editing a DDNS Profile 564
Deleting DDNS Profiles 564

Configuring AWS Credentials 565
Network > AWS Configuration 565
About AWS 565
Configuring AWS 566
Troubleshooting the Connection 567

Part 5. System Setup | SD-WAN

About SD-WAN 570

Configuring SD-WAN Groups 571
SD-WAN Groups 571
Creating an SD-WAN Group 572
Editing an SD-WAN Group 572
Deleting SD-WAN Groups 573

Configuring Performance Probes 575
About Performance Probes 575
Configuring Performance Probes 578
Editing an SD-WAN Performance Probe 579
Deleting SD-WAN Performance Probes 579
Deleting a Performance Probe 579
Deleting Multiple Performance Probes 579
Deleting All Performance Probes 580

Configuring Performance Class Objects 581
About Performance Class Objects 581
Configuring Performance Class Objects 583
Editing a Performance Class Object 584
Deleting SD-WAN Performance Class Objects 584
Deleting a Performance Class Object 584
Deleting Multiple Performance Class Objects 584
Deleting All Performance Class Objects 585

Configuring Path Selection Profiles 586
About Path Selection Profiles 586
Configuring Path Selection Profiles 588
Editing a Path Selection Profile 589
Deleting SD-WAN Path Selection Profiles 589
Deleting a Path Selection Profile 589

Deleting Multiple Path Selection Profiles . . . 589
Deleting All Path Selection Profiles . . . 590

Configuring SD-WAN Route Policies . . . 591
About SD-WAN Route Policies . . . 591
Configuring SD-WAN Route Policies . . . 592
Editing SD-WAN Route Policies . . . 594
Deleting SD-WAN Route Policies . . . 594
Deleting SD-WAN Route Policies . . . 595
Deleting Multiple SD-WAN Route Policies . . . 595
Deleting All SD-WAN Route Policies . . . 595

Monitoring SD-WAN . . . 597
SD-WAN > SD-WAN Monitor . . . 597

Viewing SD-WAN Route Policy Connections . . . 599
SD-WAN > SD-WAN Connection Logs . . . 599

Part 6. System Setup | Switching

About Switching . . . 602
About Switching . . . 602
SonicOS Switching . . . 602
Benefits of Switching . . . 603
How Switching Works . . . 604
Glossary . . . 604

Configuring VLAN Trunking . . . 605
Switching > VLAN Trunking . . . 606
About Trunking . . . 607
Viewing VLANs . . . 607
Editing VLANs . . . 609
Adding a VLAN Trunk Port . . . 609
Enabling a VLAN on a Trunk Port . . . 610
Deleting VLAN Trunk Ports . . . 610

Managing Layer 2 Discovery and LLDP/LLTD . . . 612
Switching > Port Mirroring . . . 612
About L2 Discovery and LLDP . . . 613
Viewing L2 Discovery and LLDP/LLTD Interfaces . . . 616
Associating an LLDP Profile with a L2 Discovery Interface . . . 619
Refreshing the Page . . . 619
Globally Enabling/Disabling LLDP . . . 619
Discovering Neighbors . . . 620
Switching > Port Mirroring > LLDP Profile . . . 621
Viewing LLDP Profiles . . . 621
Adding a LLDP Custom Profile . . . 623
Editing a Custom LLDP Profile . . . 625
Deleting Custom Profiles . . . 625

Configuring Link Aggregation **627**
Switching > Link Aggregation **627**
About Link Aggregation 627
Viewing Link Aggregation 629
Creating a Logical Link (LAG) 631
Deleting a LAG 632

Configuring Port Mirroring **633**
Switching > Port Mirroring **633**
About Port Mirroring 633
Viewing Mirrored Ports 634
Configuring a Port Mirroring Group 634
Enabling a Mirrored Group 635
Editing a Port Mirroring Group 636
Deleting Port Mirroring Groups 636

Configuring Switch Shield **639**
Switching > Switch Shield 639

Configuring from the Firewall **640**
Before Adding a Switch 641
Checking Switch Details 641
Adding a Switch to a Firewall with Zero-Touch 642
Adding a Switch to a Firewall Manually 645
Changing the Switch Configuration 647
Check the Status and Link Details 647
Enabling the Switch 648
Shutting Down the Switch 648
Restarting the Switch 649
Adding a VLAN 650
Adding Static Routes 653
Editing DNS 654
Setting Up the Ports 655
Setting Up QoS 656
Setting Up PoE 658
Setting Up Users 659
Setting Up 802.1x Authentication 660
Daisy-Chaining Switches 661
Connecting Access Points 663
Modifying the MAC Address Table 669
Checking Port Statistics 670
Setting Spanning Tree Protocol 671

Part 7. System Setup | High Availability

About High Availability and Active/Active Clustering **674**
High Availability 674

About High Availability . . . 675
About Active/Standby HA . . . 679
About Stateful Synchronization . . . 681
About Active/Active DPI HA . . . 683
Active/Standby and Active/Active DPI Prerequisites . . . 683
Maintenance . . . 686
Active/Active Clustering . . . 688
About Active/Active Clustering . . . 688

Configuring High Availability . . . 703
High Availability > Base Setup . . . 703
Configuring Active/Standby High Availability Settings . . . 704
Configuring HA with Dynamic WAN Interfaces . . . 706
Configuring Active/Active DPI High Availability Settings . . . 707
Configuring Active/Active Clustering . . . 709
Verifying Active/Active Clustering Configuration . . . 714
Configuring Network DHCP and Interface Settings . . . 716
Active/Active Clustering Full-Mesh . . . 718

Fine Tuning High Availability . . . 725
High Availability > Advanced Settings . . . 725
Configuring Advanced High Availability . . . 726

Monitoring High Availability . . . 728
High Availability > Monitoring Settings . . . 728
Configuring Active/Standby High Availability Monitoring . . . 729
IPv6 High Availability Monitoring . . . 730

Part 8. System Setup | WAN Acceleration

Using WAN Acceleration . . . 733
About WAN Acceleration . . . 733
Supported Platforms . . . 734
Transmission Control Protocol Acceleration . . . 734
Windows File Sharing Acceleration . . . 734
Web Cache . . . 735
Deployment Prerequisites for the WAN Acceleration Service . . . 735
About WXA Clustering . . . 736
How Does WXA Clustering Work? . . . 738
Permitting Acceleration on a Route Policy . . . 739
System Setup > WAN Acceleration . . . 739
Enabling WAN Acceleration . . . 740
Managing Groups . . . 740
Managing WXAs through the WXAs Table . . . 744
Configuring WXA for VPN Polices . . . 760
Configuring Acceleration of SSL VPN Traffic . . . 760
Displaying and Editing Route Polices for WXA . . . 761
Monitoring Group Connections . . . 762

Part 9. System Setup | VoIP

About VoIP 765
About VoIP 765
What is VoIP? 765
VoIP Security 765
VoIP Protocols 766
SonicWall's VoIP Capabilities 768

Configuring SonicWall VoIP Features 776
Configuration Tasks 776
Configuring VoIP 776
Configuring VoIP Logging 782

Part 10. System Setup | Virtual Assist

Configuring Virtual Assist 784
About Virtual Assist 784
Maximizing Virtual Assist Flexibility 785
Configuring Virtual Assist 786

Part 11. System Setup | Appendices

Configuring Open Authentication, Social Login, and LHM 792
About OAuth and Social Login 792
What are OAuth and Social Login? 793
Benefits of OAuth and Social Login 793
How Do OAuth and Social Login Work? 794
Supported Platforms 795
Requirements for Development and Production 795
About Lightweight Hotspot Messaging (LHM) 796
Configuring Facebook for Social Login 798
Facebook Settings 799
Client OAuth Settings 800
Guest Status (demo) 800
Configuring Open Authentication and Social Login 800
About Configuring Guest Services 800
About Configuring Social Login 801
Configuring Social Login in SonicOS 801
Verifying the Social Login Configuration 807
Using Social Login, LHM, and ABE 807
About ABE 808
Session Life Cycle 809
Session Update 816
Message Format 816
LHM RESTful API 822
Frequently Asked Questions (FAQs) 822
LHM Script Library 829

IPv6 932
IPv6 932
About IPv6 932
Configuring IPv6 938
IPv6 Visualization 961
IPv6 High Availability Monitoring 962
IPv6 Diagnostics and Monitoring 963

BGP Advanced Routing 964
BGP Advanced Routing 964
About BGP 964
Caveats 971
Configuring BGP 971
Verifying BGP Configuration 981
IPv6 BGP 984

SonicWall Support 1005
About This Document 1006

Part 1

System Setup | About System Setup

- About Setting Up Your SonicOS System

1

About Setting Up Your SonicOS System

- About the SonicOS Management Interface on page 16

About the SonicOS Management Interface

The web-based SonicOS Management Interface allows you to configure SonicWall Security Appliances (firewalls) running SonicOS 6.5 and above. For a complete list of appliances supported by SonicOS, see *SonicOS 6.5 About SonicOS.*

SonicOS provides an easy-to-use, graphical Management Interface for configuring your SonicWall Security Appliance. For information about the dynamic Management Interface and its features, such as tooltips and dynamic tables, see *SonicOS 6.5 About* SonicOS.

This guide provides instructions about configuring:

- Passwords, login security, web management, certificates, and schedules.
- User authentication, groups, guest services and accounts, and partitioning.
- Network settings, such as interfaces, zones, and routing.
- SD-WAN.
- Switching settings for VLAN trunking, L2 discovery, link aggregation, and port mirroring.
- High availability.
- WAN acceleration.
- VOIP.
- Virtual Assist.

For Information about Configuring	See the
Connectivity: VPN, SSL VPN, SonicPoint/SonicWave, wireless	*SonicOS 6.5 Connectivity*
Polices: access rules, NAT polices, and all the objects, such as address, action, match, service, and bandwidth	*SonicOS 6.5 Policies*
Licenses, updating firmware, and backing/restarting your system	*SonicOS 6.5 Updates*
Monitoring: dashboard, threat prevention, traffic, capture ATP	*SonicOS 6.5 Monitoring*
Security: Security Appliance settings, security services, anti-SPAM, Deep Packet Inspection (DPI)	*SonicOS 6.5 Security Configuration*
Logs and reporting: AppFlow settings, logs, legal	*SonicOS 6.5 Logs and Reporting*
Quick configuration	*SonicOS Quick Configuration*
SonicOS API	*SonicOS 6.5 About SonicOS*

Part 2

System Setup | Appliance

- Configuring Base Settings
- Administering SNMP
- Managing Certificates
- Configuring Time Settings
- Setting Schedules

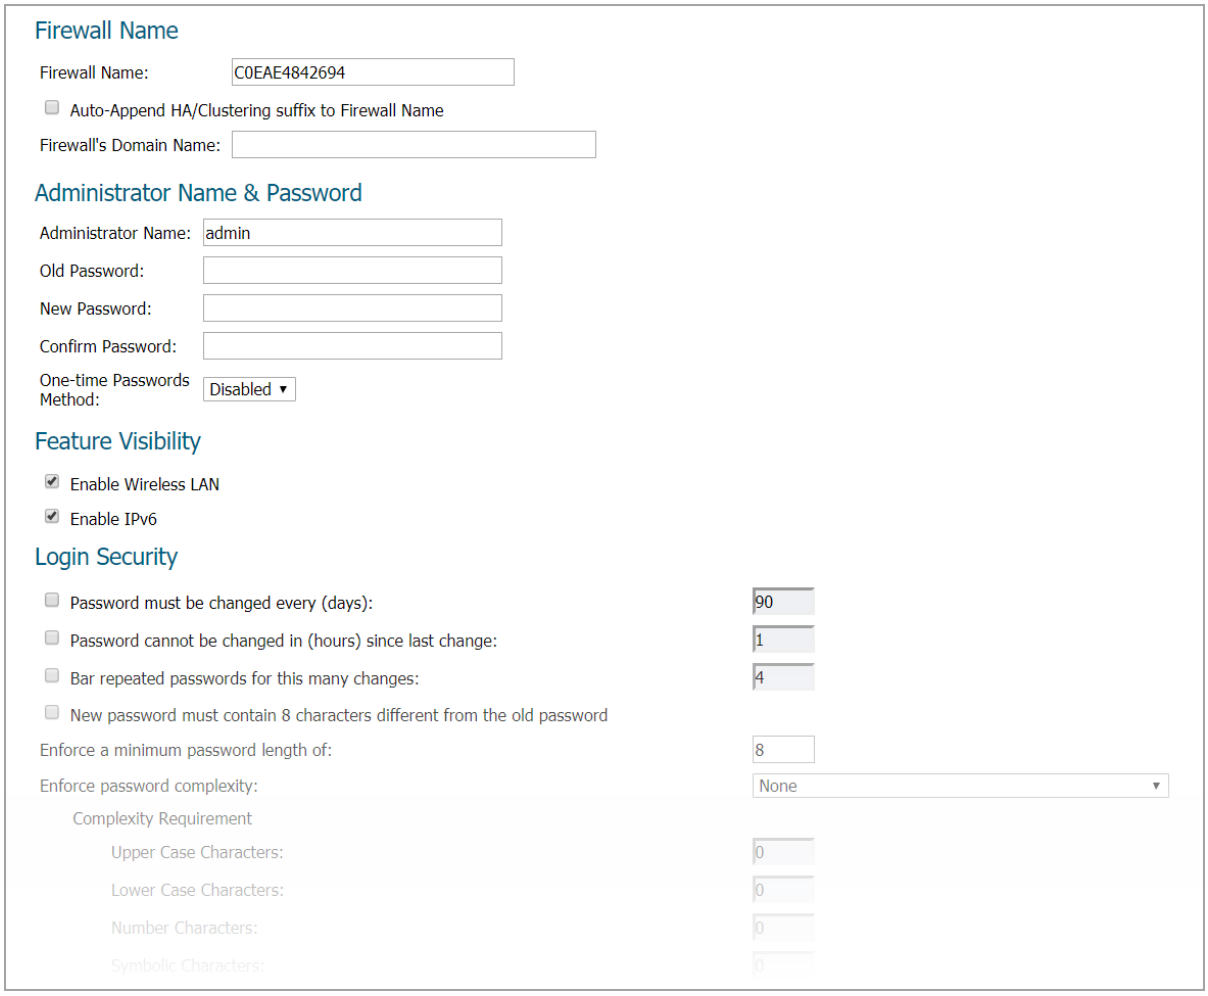
2

Configuring Base Settings

- About Appliance | Base Settings on page 19
  - One-Click Wireless and Non-Wireless Controller Modes on page 19
- Configuring the Appliance on page 21
  - Configuring the Firewall Name on page 21
  - Changing the Administrator Name & Password on page 22
  - Enabling Wireless LAN and IPv6 on page 23
  - Configuring Login Security on page 24
  - Configuring Multiple Administrator Access on page 28
  - Enabling Enhanced Audit Logging Support on page 32
  - Configuring the Wireless LAN Controller on page 32
  - Configuring the Management Interface on page 34
  - Configuring the Front-Panel Administrative Interface (SuperMassive Firewalls Only) on page 40
  - Configuring Client Certificate Verification on page 40
  - Checking Certificate Expiration on page 44
  - Configuring SSH Management on page 44
  - Enabling SonicOS API on page 44
  - Configuring Advanced Management Options on page 45
  - Downloading SonicPoint Images Manually on page 48
  - Selecting a Language on page 49

About Appliance | Base Settings

MANAGE | System Setup > **Appliance | Base Settings** provides settings for configuring the SonicWall Security Appliance for secure and remote management.



You can manage the firewall using a variety of methods, including HTTPS, SNMP or SonicWall Global Management System (SonicWall GMS).

NOTE: To apply all changes to the SonicWall appliance, click **ACCEPT**; a message confirming the update displays at the bottom of the browser window.

One-Click Wireless and Non-Wireless Controller Modes

IMPORTANT: When you change the **Wireless Controller Mode**, you must restart the firewall after clicking **ACCEPT** on the **MANAGE | System Setup > Appliance > Base Settings** page.

SonicOS 6.5 introduced Wireless Controller Mode for deployments in which the firewall is being solely used for providing secure wireless access. Alternatively, you can select Non-Wireless Controller Mode for deployments in which the firewall should not provide any wireless access. Normal, Full-Feature-Gateway mode, which allows all wireless and non-wireless functions, is the default.

This feature allows you to either:

- Enable wireless controller mode, which disables and renders uneditable:
  - SSL VPN and VPN zones.
  - Group VPN and SSL VPN policies as well as the updating of all zones using these policies.
  - VPN.
  - WAN Acceleration (WXA).
  - SIP and H.323 transformations.
- Enable non-wireless controller mode, which disables and renders uneditable:
  - Wireless zones, including the default WLAN zone, as well as disabling the creation of wireless zones.
  - Internal wireless functions.
  - Access points, including L2 and L3.

A new section, **Wireless Controller**, has been added to the **MANAGE | System Setup > Appliance | Base Settings** page.

Topics:

- Effects of Enabling Non-Wireless Controller Mode
- Effects of Enabling Wireless Controller Mode
- Enabling Wireless Controller Mode
- Enabling Non-Wireless Controller Mode
- Enabling Normal Firewall Mode

Effects of Enabling Non-Wireless Controller Mode

Enabling Non Wireless Mode affects several management interface pages.

- The **Edit** and **Delete** icons for wireless zones become dimmed on the **MANAGE | System Setup > Network > Zones** page.
- The status of access point objects becomes **Disabled** on the **MANAGE I Connectivity > Access Points > Base Settings** page.
- Internal wireless zones are disabled.
- Any attempt to enabled and access point or internal wireless is rejected:

Effects of Enabling Wireless Controller Mode

Enabling Wireless Controller Mode affects several management interface pages. Attempts to enable and/or configure features on these pages are denied.

- The **Edit** and **Delete** icons for VPN and SSL VPN zones become dimmed on the **MANAGE | System Setup > Network > Zones** page.
- VPN is disabled, as shown on the **MANAGE I Connectivity > VPN > Base Settings** page.
- All interfaces to SSL VPN are disabled on the **MANAGE I Connectivity > SSL VPN > Server Settings** page.
- Any attempt to enable SIP and/or H.323 options on **MANAGE | System Setup > VOIP** displays an error message in the lower right corner of the browser window:

  Clicking the **View List** link displays an error log:

- WXA is shown as disabled.
- Any attempt to enable a zone with VPN and/or SSL VPN results in an error.

Configuring the Appliance

Topics:

- Accessing the Appliance | Base Settings Page on page 21
- Configuring the Firewall Name on page 21
- Changing the Administrator Name & Password on page 22
- Enabling Wireless LAN and IPv6 on page 23
- Configuring Login Security on page 24
- Configuring Multiple Administrator Access on page 28
- Enabling Enhanced Audit Logging Support on page 32
- Configuring the Wireless LAN Controller on page 32
- Configuring the Management Interface on page 34
- Configuring the Front-Panel Administrative Interface (SuperMassive Firewalls Only) on page 40
- Configuring Client Certificate Verification on page 40
- Checking Certificate Expiration on page 44
- Configuring SSH Management on page 44
- Enabling SonicOS API on page 44
- Configuring Advanced Management Options on page 45
- Downloading SonicPoint Images Manually on page 48
- Selecting a Language on page 49

Accessing the Appliance | Base Settings Page

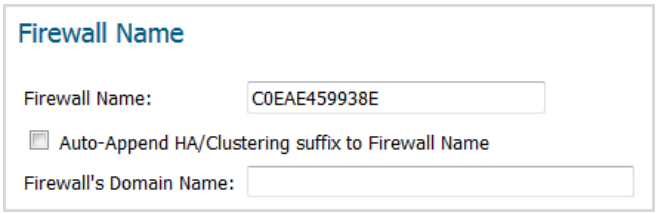
Accessing the Appliance | Base Settings page:

1. Navigate to **MANAGE | System Setup > Appliance** to expand the navigation pane.
2. Click **Base Settings**.

Configuring the Firewall Name

To configure the firewall name:

1. Navigate to **MANAGE | System Setup > Appliance > Base Settings**.
2. Scroll to the **Firewall Name** section.



3 Enter the hexadecimal serial number of the firewall in the **Firewall Name** field. This number uniquely identifies the SonicWall Security Appliance and defaults to the serial number of the firewall. The serial number is also the MAC address of the unit. To change the **Firewall Name**, enter a unique alphanumeric name in the **Firewall Name** field. It must be at least 8 characters in length and can be up to 63 characters long.

4 To facilitate recognition of the primary/secondary firewalls in the Event Logs, check **Auto-Append HA/Clustering suffix to Firewall Name**. When this option is enabled, an appropriate suffix is appended automatically to the firewall name in **MANAGE | INVESTIGATE > Log > Event Logs**:

- **Primary**
- **Secondary**
- **Primary Node** *<nodeNumber>*
- **Secondary Node** *<nodeNumber>*

This option is not selected by default. For more information about Event Logs, see *SonicOS 6.5 Investigate*.

5 Enter a friendly name in the **Firewall's Domain Name**. The name can be private, for internal users, or an externally registered domain name. This domain name is used in conjunction with **User Web Login Settings** on the **System Setup > Users > Settings** view for user-authentication redirects. For more information about user web login settings, see User Web Login Settings on page 121.

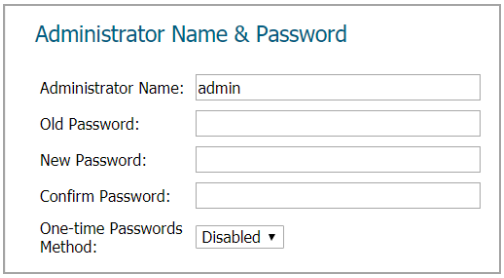
Changing the Administrator Name & Password

Each SonicWall Security Appliance has a default administrator name of `admin` and a password of `password`. If you did not change the password with the Initial **Setup Guide** or **Startup Guide**, or with the **Setup Quick Configuration Guide**, it is highly recommended that you do so now.

To change the administrator name and/or password:

1 Navigate to **MANAGE | System Setup > Appliance > Base Settings**.

2 Scroll to **Administrator Name & Password**.



3 Type the new name in the **Administrator Name** field. The **Administrator Name** can be changed from the default setting of **admin** to any word using alphanumeric characters up to 32 characters in length. It must be at least one character long.

4 Click **ACCEPT**.

To set a new password for SonicWall management interface access:

1. Type the old password in the **Old Password** field. The old password is shown encrypted in the **Old Password** field.
2. Type the new password in the **New Password** field. The new password can be up to 32 alphanumeric and special characters.

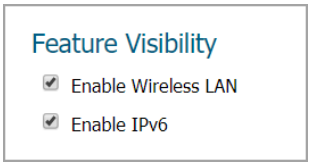
   IMPORTANT: It is recommended you change the default password, **password**, to your own custom password. Enter a strong password that cannot be easily guessed by others. A strong password should have at least one uppercase letter, one lowercase letter, one number, and one special character. For example, `MyP@ssw0rd`.

3. Type the new password again in the **Confirm Password** field.
4. To enforce the use of a One-time Passwords Method, select TOTP from the drop-down. Leave this feature disabled to allow for the reuse of old passwords.
5. Click **ACCEPT**.

Enabling Wireless LAN and IPv6

To enable the visibility of a wireless LAN and/or IPv6:

1. Navigate to **MANAGE | System Setup > Appliance > Base Settings**.
2. Scroll to **Feature Visibility**.

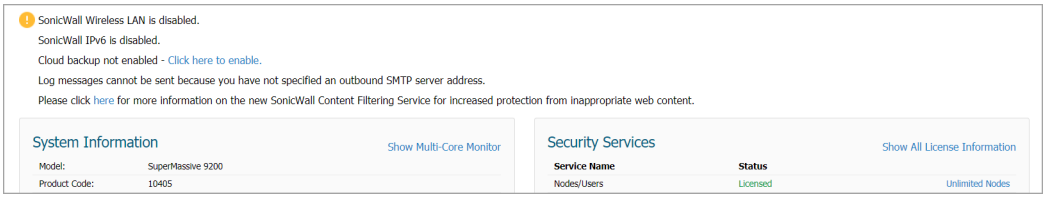


3. Click **Enable Wireless LAN** and/or **Enable IPv6**. These options are selected by default. A confirmation message displays.



IMPORTANT: Enabling or disabling the Wireless LAN feature requires a restart of the firewall.

By selecting either or both of these options, you allow these features to be used and they are visible on the **MONITOR | Current Status > System Status** page. If the options are not selected, the **System Status** page displays that they are disabled.

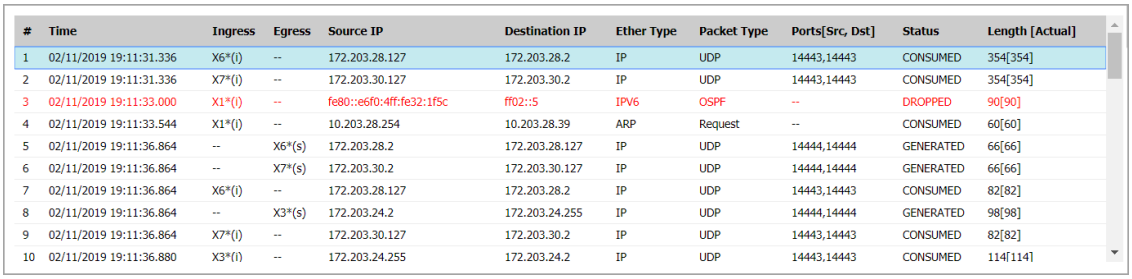


When WLAN is disabled:

- All access point and wireless-related management interface pages do not display.
- WLAN is not displayed as a zone type.

- Any existing WLAN zones or objects become uneditable.

When IPv6 is disabled, all IPv6 packets are dropped by the firewall and the **INVESTIGATE | Tools > Packet Monitor** page displays the log messages.



#	Time	Ingress	Egress	Source IP	Destination IP	Ether Type	Packet Type	Ports[Src, Dst]	Status	Length [Actual]
1	02/11/2019 19:11:31.336	X6*(i)	--	172.203.28.127	172.203.28.2	IP	UDP	14443,14443	CONSUMED	354[354]
2	02/11/2019 19:11:31.336	X7*(i)	--	172.203.30.127	172.203.30.2	IP	UDP	14443,14443	CONSUMED	354[354]
3	02/11/2019 19:11:33.000	X1*(i)	--	fe80::e6f0:4ff:fe32:1f5c	ff02::5	IPV6	OSPF	--	DROPPED	90[90]
4	02/11/2019 19:11:33.544	X1*(i)	--	10.203.28.254	10.203.28.39	ARP	Request	--	CONSUMED	60[60]
5	02/11/2019 19:11:36.864	--	X6*(s)	172.203.28.2	172.203.28.127	IP	UDP	14444,14444	GENERATED	66[66]
6	02/11/2019 19:11:36.864	--	X7*(s)	172.203.30.2	172.203.30.127	IP	UDP	14444,14444	GENERATED	66[66]
7	02/11/2019 19:11:36.864	X6*(i)	--	172.203.28.127	172.203.28.2	IP	UDP	14443,14443	CONSUMED	82[82]
8	02/11/2019 19:11:36.864	--	X3*(s)	172.203.24.2	172.203.24.255	IP	UDP	14444,14444	GENERATED	98[98]
9	02/11/2019 19:11:36.864	X7*(i)	--	172.203.30.127	172.203.30.2	IP	UDP	14443,14443	CONSUMED	82[82]
10	02/11/2019 19:11:36.880	X3*(i)	--	172.203.24.255	172.203.24.2	IP	UDP	14443,14443	CONSUMED	114[114]

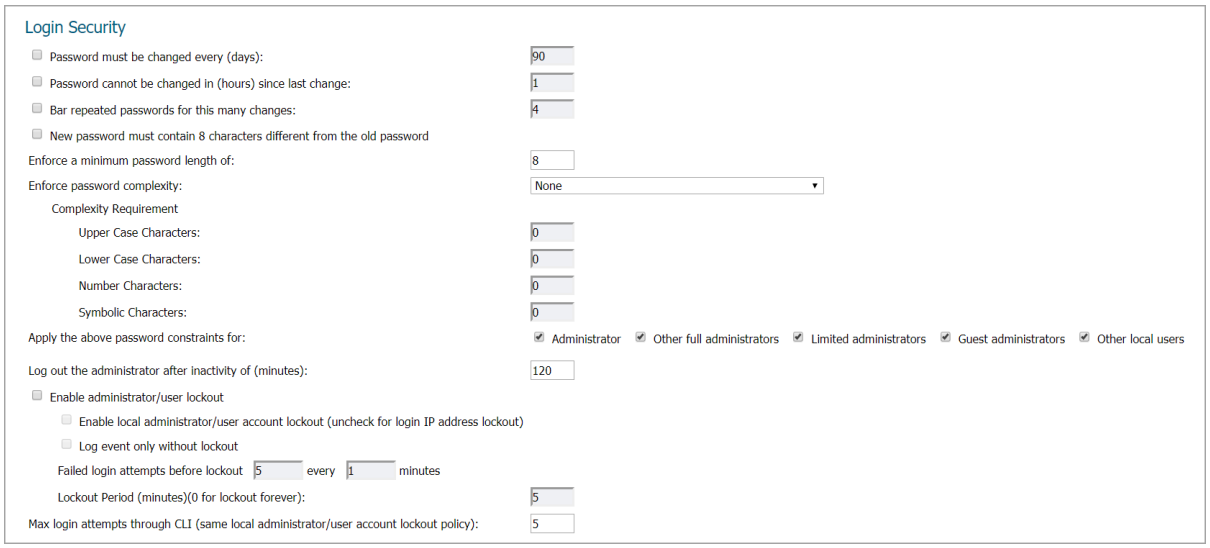
4 Click **OK**.

Configuring Login Security

The internal SonicOS Web-server supports TLS 1.1 and above with strong ciphers (128 bits or greater) when negotiating HTTPS management sessions. SSL implementations are not supported. This heightened level of HTTPS security protects against potential SSLv2 rollback vulnerabilities and ensures compliance with the Payment Card Industry (PCI) and other security and risk-management standards.

TIP: SonicOS uses advanced browser technologies, such as HTML5, which are supported in most recent browsers. SonicWall recommends using the latest Chrome, Firefox, Internet Explorer, or Safari (does not operate on Windows platforms) browsers for administration of SonicOS. Mobile device browsers are not recommended for SonicWall system administration.

Configuring SonicOS password constraint enforcement ensures that administrators and users are using secure passwords. This password constraint enforcement can satisfy the confidentiality requirements as defined by current information security management systems or compliance requirements, such as Common Criteria and the Payment Card Industry (PCI) standard.



Topics:

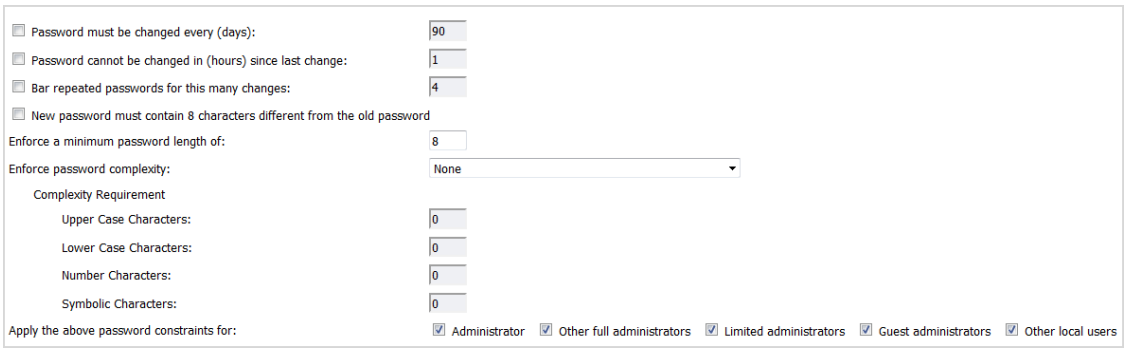
- Configuring Password Compliance on page 25
- Configuring Login Constraints on page 26

Configuring Password Compliance

To configure password compliance:

1 Navigate to **MANAGE | System Setup > Appliance > Base Settings**.

2 Scroll to **Login Security**.



3 To require users to change their passwords after a designated number of days has elapsed:

   a Select **Password must be changed every (days)**. The field becomes active. This option is not selected by default.

   b Enter the elapsed time in the field. The default number of days is **90**, the minimum is 1 day, and the maximum is 9999.

   When a user attempts to login with an expired password, a popup window prompts the user to enter a new password. The **User Login Status** window now includes **Change Password** so users can change their passwords at any time.

4 To specify the minimum length of time, in hours, allowed between password changes:

   a Select **Password cannot be changed in (hours) since the last change**. The field becomes active. This option is not selected by default.

   b Enter the number of hours. The minimum – and default – time is **1** hour; the maximum is 9999 hours.

5 To require users to use unique passwords for the specified number of password changes:

   a Select **Bar repeated passwords for this many changes**. The field becomes active. This option is not selected by default.

   b Enter the number of changes. The default number is **4**, the minimum number is 1, and the maximum number is 32.

6 To require users to change at least 8 alphanumeric/symbolic characters of their old password when creating a new one, select **New password must contain 8 characters different from the old password**. For how to specify what characters are allowed, see Step 8.

7 Specify the shortest allowed password, enter the minimum number of characters in the **Enforce a minimum password length of** field. The default number is **8**, the minimum is 1, and the maximum is 99.

8 Choose how complex a user's password must be to be accepted from the **Enforce password complexity** drop-down menu:

- **None** (default)
- **Require both alphabetic and numeric characters**
- **Require alphabetic, numeric, and symbolic characters** – for symbolic characters, only **!**, **@**, **#**, **$**, **%**, **^**, **&**, *****, **(**, and **)** are allowed; all others are denied

9 When a password complexity option other than **None** is selected, the options under **Complexity Requirement** become active. Enter the minimum number of alphanumeric and symbolic characters required in a user's password. The default number for each is **0**, but the total number of characters for all options cannot exceed 99.

- **Upper Case Characters**
- **Lower Case Characters**
- **Number Characters**
- **Symbolic Characters**

> **NOTE:** The **Symbolic Characters** field becomes active only if **Require alphabetic, numeric and symbolic characters** is selected.

10 Select to which classes of users the password constraints are applied under **Apply the above password constraints for**. By default, all options are selected:

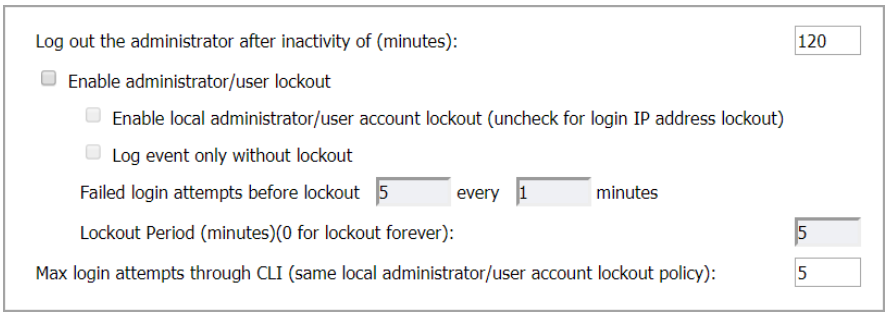
- **Administrator** – Refers to the default administrator with the username **admin**.
- **Other full administrators**
- **Limited administrators**
- **Guest administrators**
- **Other local users**

Configuring Login Constraints

To configure login constraints:

1 Navigate to **MANAGE | System Setup > Appliance > Base Settings**.

2 Scroll to **Login Security**.



3 To specify the length of inactivity time that elapses before you are automatically logged out of the Management Interface, enter the time, in minutes, in the **Log out the Administrator after inactivity of (minutes)** field. By default, the SonicWall Security Appliance logs out the administrator after **5** minutes of inactivity. The inactivity timeout can range from 1 to 9999 minutes.

> **TIP:** If the Administrator Inactivity Timeout is extended beyond five minutes, you should end every management session by clicking **Logout** in the upper right corner of the view to prevent unauthorized access to the firewall's Management Interface.

4 To configure the SonicWall Security Appliance to lockout an administrator or a user if the login credentials are incorrect, select **Enable administrator/user lockout**. Both administrators and users are

locked out of accessing the firewall after the specified number of incorrect login attempts. This option is disabled by default. When this option is enabled, the following fields become active.

CAUTION: If the administrator and a user are logging into the firewall using the same source IP address, the administrator is also locked out of the firewall. The lockout is based on the source IP address of the user or administrator.

a Select **Enable local administrator/user account lockout (uncheck for login IP address lockout)**. This option locks out user accounts and IP addresses when they have surpassed a specified number of incorrect login attempts. This option is only available when **Enable administrator/user lockout** is selected.

b Select **Log event only without lockout** for SonicOS to log failed user login attempts that have reached the established threshold, but does not lock out the user or IP address. This option is only available when **Enable administrator/user lockout** is selected.

After a user or IP address is locked out, a "User login denied - User is locked out" message displays on the login screen and the login is rejected.



NOTE: You can review and edit all locked out user accounts on the **MONITOR | Current Status | User Sessions > Active Users** page when the **Enable local administrator/user account lockout (uncheck for login IP address lockout)** checkbox is selected.

c Enter the number of failed attempts within a specified time frame before the user is locked out in the first **Failed login attempts per minute before lockout** field. The default number is **5**, the minimum is 1, and the maximum is 99.

d Enter the maximum time in which failed attempts can be made. The default is **5** minutes, the minimum is 1 minute, and the maximum is 240 minutes (4 hours)

e Enter the length of time that must elapse before the user is allowed to attempt to log into the firewall again in the **Lockout Period (minutes)** field. The default is **5** minutes, the minimum is 0 (permanent lockout), and the maximum is 60 minutes.

5 Enter the number of incorrect login attempts from the command line interface (CLI) that triggers a lockout in the **Max login attempts through CLI** field. The default is **5**, the minimum is 3, and the maximum is 15.

6 Click **ACCEPT**.

Locked Out User Accounts

SonicOS maintains an Unauthenticated Users record indicating how many unauthenticated user accounts or IP addresses have been locked out. The **Enable local administrator/user account lockout (uncheck for login IP address lockout)** feature is meant to lockout unauthenticated users and IP addresses when the user or

administrator authentication/login fails. See that record on the **MONITOR | Current Status | User Sessions > Active Users** page.

Unlocking Unauthenticated User Accounts

There are three ways that you can unlock a locked user account:

- **Lockout timer** - This method unlocks the unauthenticated account automatically after the established waiting period has elapsed.
- **GUI unlock** - This method requires that an administrator unlocks unauthenticated users manually by clicking an **Unlock** icon on the **MONITOR | Current Status | User Sessions > Active Users** page.
- **Disable checkbox** - Deselecting the **Enable local administrator/user account lockout (uncheck for login IP address lockout)** checkbox unlocks all unauthenticated users unless a specific user has been marked locked "forever." The lockout period for that particular unauthenticated user would show as finite minutes, and would not be unlocked by deselecting the option.

Configuring Multiple Administrator Access

SonicOS supports multiple concurrent administrators with full administrator privileges, read-only privileges, and limited privileges.

Topics:

- About Multiple Administrator Support on page 28
- Configuring Multiple Administrator Access on page 31

About Multiple Administrator Support

Topics:

- What is Multiple Administrators Support? on page 28
- Benefits on page 28
- How Does Multiple Administrators Support Work? on page 29

What is Multiple Administrators Support?

The original version of SonicOS supported only a single administrator to log on to a firewall with full administrative privileges. Additional users can be granted "limited administrator" access, but only one administrator can have full access to modify all areas of the SonicOS GUI at one time.

SonicOS provides support for multiple concurrent administrators. This feature allows for multiple users to log-in with full administrator privileges. In addition to using the default **admin** user name, additional administrator user names can be created.

Because of the potential for conflicts caused by multiple administrators making configuration changes at the same time, only one administrator is allowed to make configuration changes. The additional administrators are given full access to the GUI, but they cannot make configuration changes.

Benefits

Multiple Administrators Support provides the following benefits:

Improved productivity	Allowing multiple administrators to access a firewall simultaneously eliminates auto logout, a situation that occurs when two administrators require access to the appliance at the same time and one is automatically forced out of the system.
Reduced configuration risk	The new read-only mode allows users to view the current configuration and status of a firewall without the risk of making unintentional changes to the configuration.

How Does Multiple Administrators Support Work?

Topics:

- Configuration Modes on page 29
- User Groups on page 30
- Priority for Preempting Administrators on page 31
- GMS and Multiple Administrator Support on page 31

Configuration Modes

To allow multiple concurrent administrators, while also preventing potential conflicts caused by multiple administrators making configuration changes at the same time, these configuration modes have been defined:

Configuration mode	Administrator has full privileges to edit the configuration. If no administrator is already logged into the appliance, this is the default behavior for administrators with full and limited administrator privileges (but not read-only administrators). **NOTE:** Administrators with full configuration privilege can also log in using the Command Line Interface.
Read-only mode	Administrator cannot make any changes to the configuration, but can view the entire management UI and perform monitoring actions. Only administrators who are members of the **SonicWall Read-Only Admins** user group are given read-only access, and it is the only configuration mode they can access.
Non-configuration mode	Administrator can view the same information as members of the read-only group and they can also initiate management actions that do not have the potential to cause configuration conflicts. Only administrators who are members of the **SonicWall Administrators** user group can access non-configuration mode. This mode can be entered when another administrator is already in configuration mode and the new administrator chooses not to preempt the existing administrator. By default, when an administrator is preempted out of configuration mode, he or she is converted to non-configuration mode. On the **System Setup > Administration** page, this behavior can be modified so that the original administrator is logged out.

Access rights available to configuration modes provides a summary of the access rights available to the configuration modes. Access rights for limited administrators are included also, but note that this table does not include all functions available to limited administrators.

Access rights available to configuration modes

Function	Full admin in config mode	Full admin in non-config mode	Read-only administrator	Limited administrator
Import certificates	X			
Generate certificate signing requests	X			
Export certificates	X			
Export appliance settings	X	X	X	
Download TSR	X	X	X	
Use other diagnostics	X	X		X
Configure network	X			X
Flush ARP cache	X	X		X
Setup DHCP Server	X			
Renegotiate VPN tunnels	X	X		
Log users off	X	X		X guest users only
Unlock locked-out users	X	X		
Clear log	X	X		X
Filter logs	X	X	X	X
Export log	X	X	X	X
Email log	X	X		X
Configure log categories	X	X		X
Configure log settings	X			X
Generate log reports	X	X		X
Browse the full UI	X	X	X	
Generate log reports	X	X		X

User Groups

The Multiple Administrators Support feature supports two new default user groups:

SonicWall Administrators	Members of this group have full administrator access to edit the configuration.
SonicWall Read-Only Admins	Members of this group have read-only access to view the full management interface, but they cannot edit the configuration and they cannot switch to full configuration mode.

It is not recommended to include users in more than one of these user groups. If you do so, however, the following behavior applies:

If members of this user group	Are
SonicWall Administrators	Also included in the **Limited Administrators** or **SonicWall Read-Only Admins** user groups, the members have full administrator rights.

If members of this user group	Are
Limited Administrators	Included in the **SonicWall Read-Only Admins** user group, the members have limited administrator rights.
Read-Only Admins	Later included in another administrative group, the **If this read-only admin group is used with other administrative groups** option in the **SonicWall Read-Only Admins** group configuration determines whether the members are still restricted to read-only access or have the full administration capabilities set by their other group.

Priority for Preempting Administrators

These rules govern the priority levels that the various classes of administrators have for preempting administrators that are already logged into the appliance:

1. The **admin** user and SonicWall Global Management System (GMS) have the highest priority and can preempt any users.
2. A user who is a member of the **SonicWall Administrators** user group can preempt any users except for the **admin** and SonicWall GMS.
3. A user who is a member of the **Limited Administrators** user group can only preempt other members of the **Limited Administrators** group.

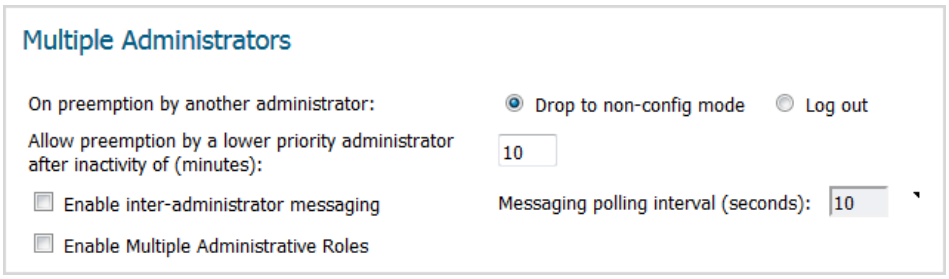
GMS and Multiple Administrator Support

When using SonicWall GMS to manage a firewall, GMS frequently logs in to the appliance (for such activities as ensuring that GMS management IPSec tunnels have been created correctly). These frequent GMS log-ins can make local administration of the appliance difficult because the local administrator can be preempted by GMS.

Configuring Multiple Administrator Access

To configure multiple administrator access:

1. Navigate to **MANAGE | System Setup > Appliance > Base Settings**.
2. Scroll to **Multiple Administrators**.



3. To configure what happens when one administrator preempts another administrator, from the **On preemption by another administrator** options, select whether the preempted administrator can be converted to non-config mode or logged out:

To Allow	Select
More than one administrator to access the appliance in non-config mode without disrupting other administrators. This option is not selected by default.	**Drop to non-config mode**
The new administrator to preempt other sessions. **NOTE:** Selecting **Log Out** disables Non-Config mode and prevents entering Non-Config mode manually.	**Log Out**

4 To allow a lower-priority administrator to preempt the current administrator after a specified time, enter the time, in minutes, in the **Allow preemption by a lower priority administrator after inactivity of (minutes)** field. The default is **10** minutes, the minimum is 1 minute, and the maximum is 9999 minutes.

5 The SonicOS Management Interface allows administrators to send text messages through the Management Interface to other administrators logged into the appliance. The message appears in the browser's status bar. This option is not selected by default. To enable this option:

   a Select **Enable inter-administrator messaging**. The **Messaging polling interval (seconds)** field becomes active.

   b Specify how often an administrator's browser checks for inter-administrator messages in the **Messaging polling interval (seconds)** field. Specify a reasonably short interval to ensure timely delivery of messages, especially if there are likely to be multiple administrators who need to access the appliance. The default is **10** seconds, the minimum is 1 second, and the maximum is 99 seconds.

6 To enable access by System Administrators, Cryptographic (Crypto) Administrators, and Audit Administrators, select **Enable Multiple Administrator Roles**. When this option is disabled, the these administrators cannot access the system, and all related user groups and information about them are hidden. This option is not selected by default.

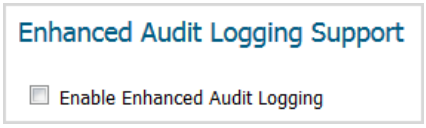
Enabling Enhanced Audit Logging Support

An enhanced log entry contains the parameter changed and user name in the **Investigate | Logs > Event Logs** page. For further information about logs, see *SonicOS 6.5 Investigate*.

To enable logging of all configuration changes in the INVESTIGATE | Logs > Event Logs page:

1 Navigate to **MANAGE | System Setup > Appliance | Base Settings**.

2 Scroll to **Enhanced Audit Logging Support**.



3 Select **Enable Enhanced Audit Logging**. This option is not selected by default.

4 Click **ACCEPT**.

Configuring the Wireless LAN Controller

You can configure the Wireless LAN Controller in one of several ways:

- Full-Feature-Gateway (default)

- Non-Wireless
- Wireless-Controller-Only

For information about the modes, see One-Click Wireless and Non-Wireless Controller Modes on page 19.

Topics:

- Enabling Normal Firewall Mode on page 33
- Enabling Wireless Controller Mode on page 33
- Enabling Wireless Controller Mode on page 33

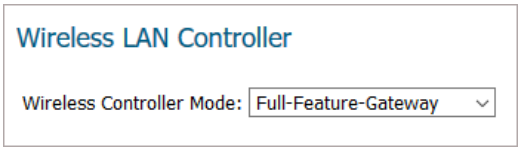
Enabling Normal Firewall Mode

To enable normal firewall mode:

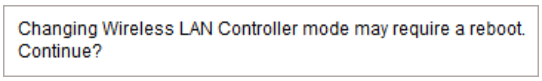
IMPORTANT: You must reboot the firewall after changing Wireless Controller modes.

1 Navigate to **MANAGE | System Setup > Appliance > Base Settings**.

2 Scroll to **Wireless Controller.**



3 From **Wireless Controller Mode**, select **Full-Feature-Gateway**. A warning message displays:



4 Click **OK**.

5 Click **ACCEPT**.

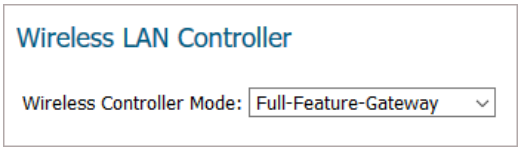
Enabling Wireless Controller Mode

To enable wireless controller mode:

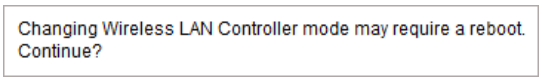
IMPORTANT: You must reboot the firewall after changing Wireless Controller modes.

1 Navigate to **MANAGE | System Setup > Appliance > Base Settings**.

2 Scroll to **Wireless LAN Controller.**



3 From **Wireless Controller Mode**, select **Wireless-Controller-Only**. A warning message displays:



4 Click **OK**.

5 Click **ACCEPT**.

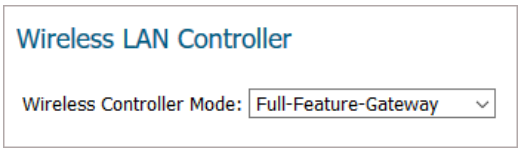
Enabling Non-Wireless Controller Mode

To enable non-wireless controller mode:

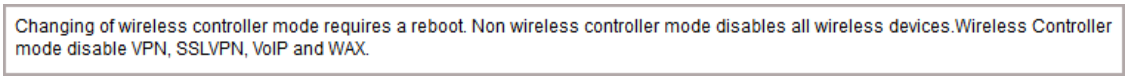
IMPORTANT: You must reboot the firewall after changing Wireless Controller modes.

1 Navigate to **MANAGE | System Setup > Appliance > Base Settings**.

2 Scroll to **Wireless LAN Controller.**



3 From **Wireless Controller Mode**, select **Non Wireless**. A warning message displays:



4 Click **OK**.

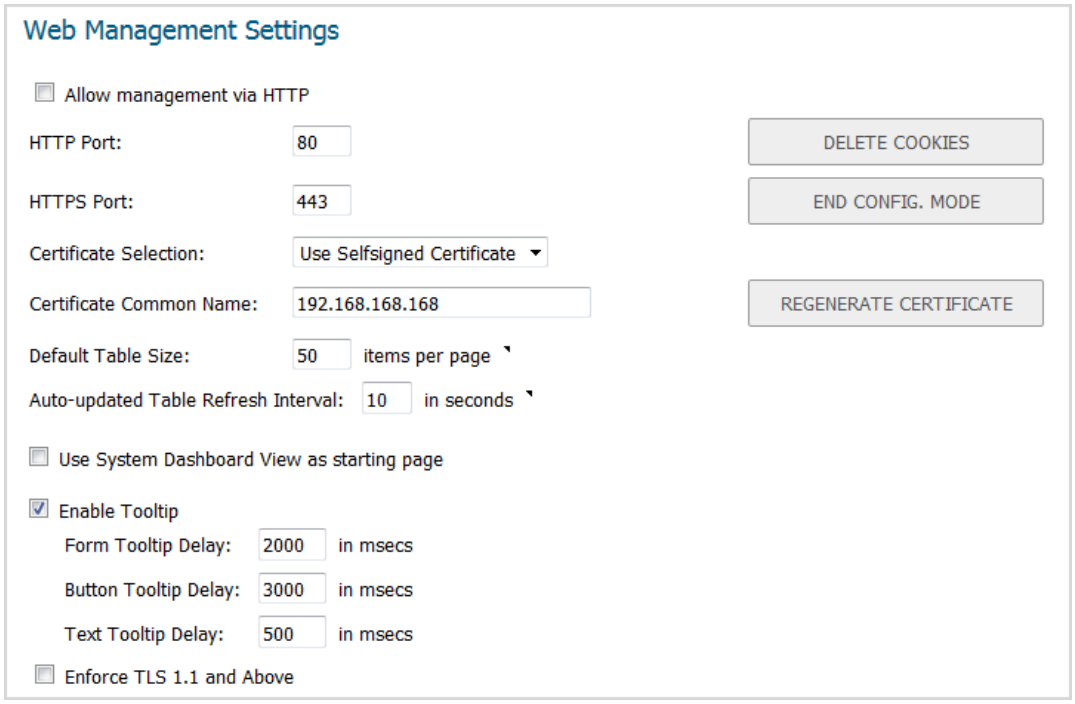
5 Click **ACCEPT**.

NOTE: The **Edit** and **Delete** icons for wireless zones become dimmed on the **MANAGE | System Setup > Network > Zones** and **MANAGE I Connectivity > Access Points > Base Settings** pages.

Configuring the Management Interface

In this section, you configure:

- How the Management Interface tables display.
- Certificate usage.
- Which page displays as a starting page.
- Whether you are operating in Configuration or Non-Config mode.
- Tooltip behavior.
- Other management options.



Topics:

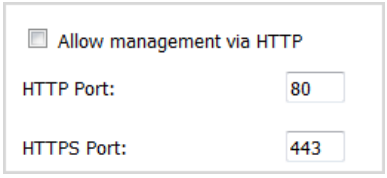
- Managing through HTTP/HTTPS on page 35
- Deleting Browser Cookies on page 36
- Switching Configuration Modes on page 36
- Switching Configuration Modes on page 36
- Controlling the Management Interface Tables on page 38
- Specifying the Starting Page on page 38
- Managing Tooltips on page 39
- Enforcing TLS Version on page 39

Managing through HTTP/HTTPS

You can manage the SonicWall Security Appliance using HTTP or HTTPS and a Web browser. HTTP web-based management is disabled by default. Use HTTPS to log into the SonicOS Management Interface with factory default settings.

To manage through HTTP or HTTPS:

1. Navigate to **MANAGE | System Setup > Appliance > Base Settings**.
2. Scroll to **Web Management Settings**.



3 To enable HTTP management globally, select **Allow management via HTTP**. This option is not selected by default.

4 The default port for HTTP is port **80**, but you can configure access through another port. Enter the number of the desired port in the **HTTP Port** field.

IMPORTANT: If you configure another port for HTTP management, you must include the port number when you use the IP address to log into the SonicWall Security Appliance. For example, if you configure the port to be `76`, then you must type *`LAN IP Address`*`:76` into the Web browser, for example, `http://192.18.16.1:76`.

5 The default port for HTTPS management is **443**. To add another layer of security for logging into the SonicWall Security Appliance by changing the default port, enter the preferred port number into the **HTTPS Port** field.

IMPORTANT: If you configure another port for HTTPS management Port, you must include the port number when you use the IP address to log into the SonicWall Security Appliance. For example, if you use `700` for the port, then you must log into the SonicWall using the port number as well as the IP address; for example, `https://192.18.16.1:700`.

Deleting Browser Cookies

IMPORTANT: Deleting cookies causes you to lose any unsaved changes made in the Management Interface.

To delete all browser cookies saved by the Security Appliance:

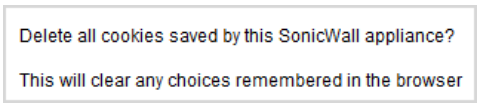
1 Navigate to **MANAGE | System Setup > Appliance > Base Settings**.

2 Scroll to **Web Management Settings**.

3 Click **Delete Cookies**.



A confirmation message displays.



4 Click **OK**. All cookies saved since the last time you deleted cookies are deleted.

Switching Configuration Modes

Each appliance includes a **Mode** option that toggles the configuration mode of the Management Interface. If you are in Configuration Mode, you can switch to Non-Config Mode at any time, or if you are in Non-Config Mode. you can switch to Configuration Mode.

TIP: This method is in addition to switching modes from the **Mode** setting on each view. For more information about modes, see the *SonicOS 6.5 About SonicOS* administration documentation.

To switch modes:

1 Navigate to **MANAGE | System Setup > Appliance > Base Settings**.

2 Scroll to **Web Management Settings**.

3 If you are in:

- Configuration Mode, click **END CONFIG. MODE**. The button changes to:

  The **Mode** indicator in the top right of the page displays **Non-Config**:

  If you attempt to save any changes on any view, an error message displays:

- Non-Config Mode, click **CONFIGURATION MODE**. The button changes to:

  The **Mode** indicator in the top right of the page displays **Configuration**:

There is no need to click **ACCEPT**.

4 To return to:

- Configuration Mode, click **CONFIGURATION MODE**.
- Non-Config Mode, click **END CONFIG. MODE**.

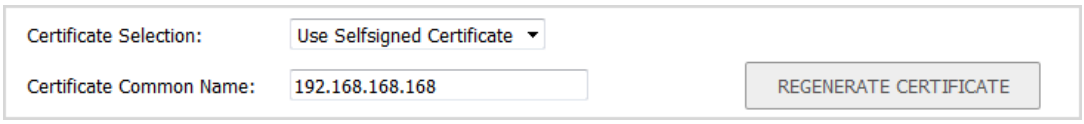
Selecting a Security Certificate

Security certificates provide data encryption and a secure web site.

To specify the type of security certificate:

1 Navigate to **MANAGE | System Setup | Appliance > Base Settings**.

2 Scroll to **Web Management Settings**.



3 From **Certificate Selection**, select the type of certificate for your web site:

- **Use Self-signed Certificate**, which allows you to continue using a certificate without downloading a new one each time you log into the SonicWall Security Appliance. This option is selected by default. Go to Step 4.
- **Import Certificate** to select an imported certificate from the **Appliance > Certificates** page to use for authentication to the management interface. A confirmation message displays:

  a) Click **OK**. The **Appliance > Certificates** page displays.

  b) Go to Managing Certificates on page 61.

4 In the **Certificate Common Name** field, enter the IP address or common name for the firewall. If you choose **Use Selfsigned Certificate**, SonicOS populates the field with the firewall's IP address.

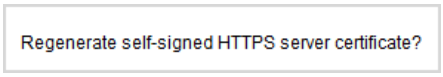
5 Click **ACCEPT**.

To regenerate a Self-Signed Certificate:

1 Navigate to **MANAGE | System Setup > Appliance > Base Settings**.

2 Scroll to **Web Management Settings**.

3 Click **REGENERATE CERTIFICATE**. A confirmation message displays:



4 Click **OK**.

Controlling the Management Interface Tables

The SonicWall Management Interface allows you to control the display of large tables of information across all tables in the Management Interface by changing the:

- Number of table entries displayed on a page.
- Frequency of background automatic refresh of tables.

Some tables have individual settings for items per page that are initialized at login to the value configured here. After these pages are viewed, their individual settings are maintained. Subsequent changes made here affect these pages only following a new login.

To change the display and refresh of tables:

1 Navigate to **MANAGE | System Setup > Appliance > Base Settings**.

2 Scroll to **Web Management Settings**.



3 Enter the desired number of **items per page** in the **Default Table Size** field. The minimum is 1, the maximum is 5000, and the default is **50**.

4 Enter the desired refresh interval, in seconds, in the **Auto-updated Table Refresh Interval** field. The minimum is 1 second, the maximum is 300 seconds, and the default is **10** seconds.

5 Click **ACCEPT**.

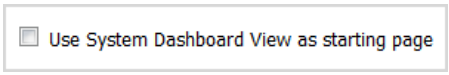
Specifying the Starting Page

When you log in to the Management Interface, the view where you logged out of the Management Interface is displayed. You can have the System Dashboard View displayed instead.

To see the MONITOR | Dashboard page first when you log in:

1 Navigate to **MANAGE | System Setup > Appliance > Base Settings**

2 Scroll to **Web Management Settings**.



3 Select **Use System Dashboard View as starting page**.

4 Click **ACCEPT**. The next time you log in, the Monitor Dashboard page displays regardless of which view was displayed when you logged out.

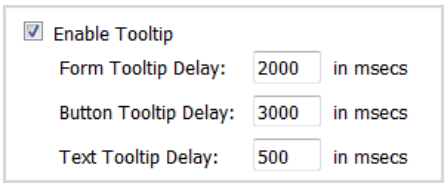
Managing Tooltips

The SonicOS Management Interface has embedded tooltips for many elements. For more information about tooltips, see *SonicOS 6.5 About* SonicOS.

To configure tooltip behavior:

1 Navigate to **MANAGE | System Setup > Appliance > Base Settings**.

2 Scroll to **Web Management Settings**.



3 To enable tooltips, select **Enable Tooltip**.

TIP: Tooltips are enabled by default. To disable tooltips, clear **Enable Tooltip**.

4 To configure the delay, in milliseconds, before tooltips display, enter the appropriate time(s):

In this field	Enter the delay for
Form Tooltip Delay	Fields. The default is **2000** ms., the minimum is 500 ms., and the maximum is 5000 ms.
Button Tooltip Delay	Radio buttons and checkboxes. The default is **3000** ms., the minimum is 500 ms., and the maximum is 5000 ms.
Text Tooltip Delay	Management Interface text. The default and minimum is **500** ms. and the maximum is 5000 ms.

5 Click **ACCEPT**.

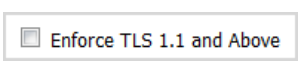
Enforcing TLS Version

SonicOS supports versions 1.0, 1.1, and 1.2 of the Transport Layer Security (TLS) protocol. You can ensure that the more secure version 1.1 and above are used.

To enforce use of TLS versions 1.1 and above:

1 Navigate to **MANAGE | System Setup > Appliance > Base Settings**.

2 Scroll to **Web Management Settings**.



3 Select **Enforce TLS 1.1 and Above**. This option is not selected by default.

4 Click **ACCEPT**.

Configuring the Front-Panel Administrative Interface (SuperMassive Firewalls Only)

NOTE: This section appears only for SuperMassive Security Appliances, which have an LCD panel in the front.

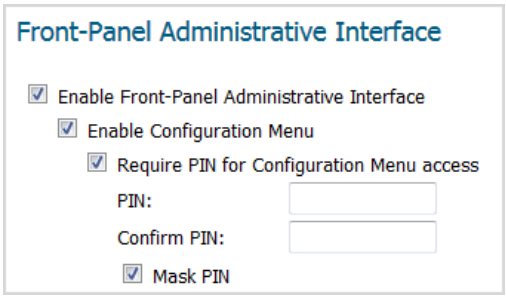
You can enable or disable access to the Configuration Menu in the front-panel administrative interface.

TIP: This feature is enabled automatically when a SuperMassive Security Appliance is first installed.

To allow access to the Configuration Menu in the front-panel administrative interface:

1 Navigate to **MANAGE | System Setup > Appliance > Base Settings**.

2 Scroll to **Front-Panel Administrative Interface**.



3 Select **Enable Front-Panel Administrative Interface**. This option is selected by default.

4 Select whether a PIN must be used to access the Configuration Menu access by checking **Require PIN for Configuration Menu access**. This option is selected by default.

   a Enter a PIN number in the **PIN** field. The default is **76642** (number dial equivalent to SONIC).

   TIP: To configure the front panel Configure menu, you have to enter the PIN number.

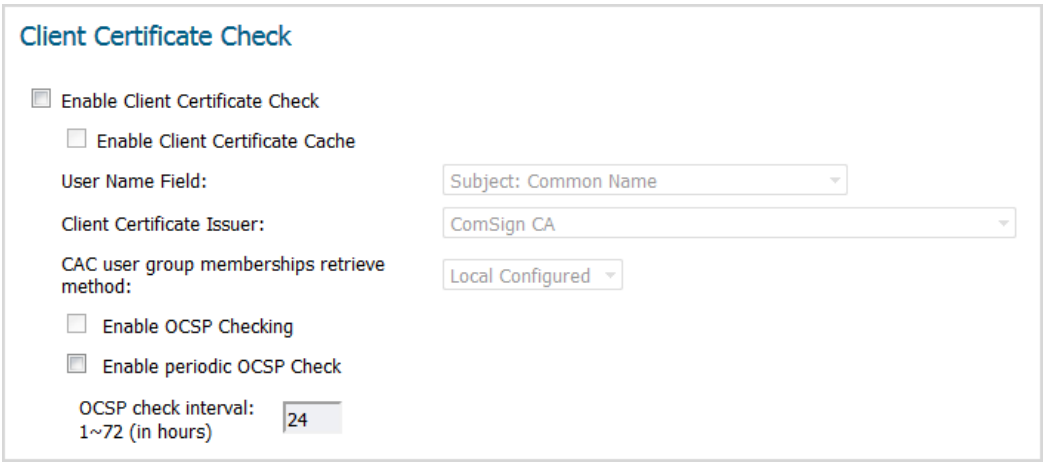
   b Enter the same PIN number in the **Confirm PIN** field.

5 Select whether the PIN is masked in the **PIN** and **Confirm PIN** fields by checking **Mask PIN**. If you mask the pin, it is displayed as a series of bullets. If this option is unchecked (not selected), the PIN is visible. This option is selected by default.

6 Click **ACCEPT**.

Configuring Client Certificate Verification

You can configure certificate verification with or without a Common Access Card (CAC).



NOTE: None of the options is selected by default.

Topics:

- About Common Access Card on page 41
- Configuring Client Certificate Verification on page 41
- Using the Client Certificate Check on page 43
- Troubleshooting User Lock Out on page 43

About Common Access Card

A Common Access Card (CAC) is a United States Department of Defense (DoD) smart card used by military personnel and other government and non-government personnel who require highly secure access over the Internet. A CAC uses PKI authentication and encryption.

NOTE: Using a CAC requires an external card reader connected on a USB port.

The Client Certificate Check was developed for use with a CAC; however, it is useful in any scenario that requires a client certificate on an HTTPS/SSL connection. CAC support is available for client certification only on HTTPS connections.

NOTE: CACs might not work with browsers other than Microsoft Internet Explorer.

Configuring Client Certificate Verification

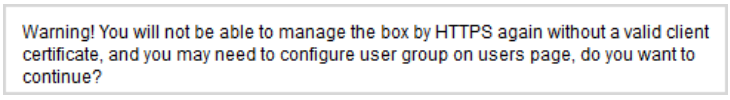
NOTE: By default, all options are not selected.

To configure Client Certificate Check:

1. Navigate to **MANAGE | System Setup > Appliance > Base Settings**.
2. Scroll to **Client Certificate Check**.



3 To enable client certificate checking and CAC support on the SonicWall Security Appliance, select **Enable Client Certificate Check**. If you enable this option, the other options become available. A warning confirmation message displays:



4 Click **OK**.

5 To activate the client certification cache, select **Enable Client Certificate Cache**.

NOTE: The cache expires 24 hours after being enabled.

6 To specify from which certificate field the user name is obtained, choose an option from **User Name Field**:

- **Subject: Common Name** (default)
- **Sub Alt: Email**
- **Sub Alt: Microsoft Universal Principal Name**

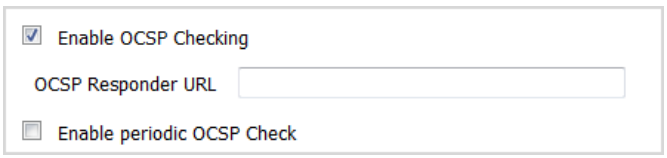
7 To select a Certification Authority (CA) certificate issuer, choose one from the **Client Certificate Issuer** drop-down menu. The default is **ComSign CA**.

NOTE: If the appropriate CA is not listed, you need to import that CA into the SonicWall Security Appliance. See Managing Certificates on page 61.

8 To select how to obtain the CAC user group membership and, thus, determine the correct user privilege, choose from the **CAC user group memberships retrieve method** drop-down menu:

- **Local Configured** (default) – If selected, you should create local user groups with proper memberships.
- **From LDAP** – If selected, you need to configure the LDAP server on the **Manage | Users > Settings** (see Configuring the SonicWall for LDAP on page 149).

9 To enable the Online Certificate Status Protocol (OCSP) check to verify the client certificate is still valid and has not been revoked, select **Enable OCSP Checking**. When this option is enabled, the **OCSP Responder URL** field displays and the **Enable periodic OCSP Check** option displays.



a Enter the URL of the OSCP server that verifies the status of the client certificate in the **OCSP Responder URL** field.

The **OCSP Responder URL** is usually embedded inside the client certificate and does not need to be entered. If the client certificate does not have an OCSP link, you can enter the URL link. The link should point to the Common Gateway Interface (CGI) on the server side, which processes the OCSP checking. For example: `http://10.103.63.251/ocsp`.

10 To enables a periodic OCSP check for the client certificate for verifying that the certificate is still valid and has not been revoked:

   a Select **Enable periodic OCSP Check**. the OCSP check interval field becomes available.

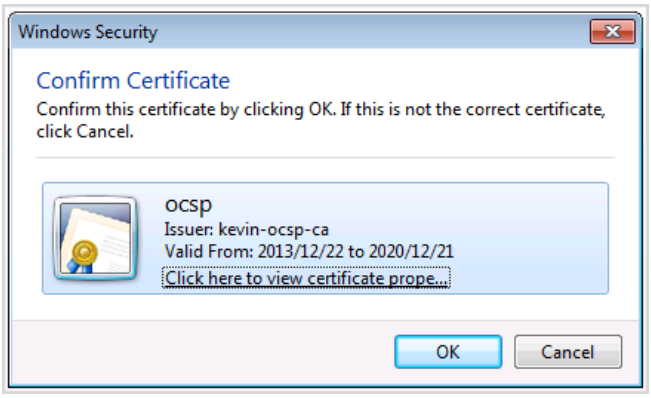
   b Enter the interval between OCSP checks, in hours, in the **OCSP check interval 1~72 (in hours)** field. The minimum interval is 1 hour, the maximum is 72 hours, and the default is **24** hours.

11 Click **ACCEPT**.

Using the Client Certificate Check

If you use the client certificate check without a CAC, you must manually import the client certificate into the browser.

If you use the **Client Certificate Check** with a CAC, the client certificate is automatically installed on the browser by middleware. When you begin a management session through HTTPS, a certificate selection window asks you to confirm the certificate.



After you select the client certificate from the drop-down menu, the HTTPS/SSL connection is resumed, and the SonicWall Security Appliance checks the **Client Certificate Issuer** to verify that the client certificate is signed by the CA. If a match is found, the administrator login page displays. If no match is found, the browser displays a standard browser connection fail message, such as:

```
.....cannot display web page!
```

If OCSP is enabled, before the administrator login page is displayed, the browser performs an OCSP check and displays the following message while it is checking.

```
Client Certificate OCSP Checking.....
```

If a match is found, the administrator login page is displayed, and you can use your administrator credentials to continue managing the SonicWall Security Appliance.

If no match is found, the browser displays:

```
OCSP Checking fail! Please contact system administrator!
```

Troubleshooting User Lock Out

When using the client certificate feature, these situations can lock the user out of the SonicWall Security Appliance:

- **Enable Client Certificate Check** is checked, but no client certificate is installed on the browser.
- **Enable Client Certificate Check** is checked and a client certificate is installed on the browser, but either no **Client Certificate Issuer** is selected or the wrong **Client Certificate Issuer** is selected.

- **Enable OSCP Checking** is enabled, but either the OSCP server is not available or a network problem is preventing the SonicWall Security Appliance from accessing the OSCP server.

To restore access to a user who is locked out, the following CLI commands are provided:

- `web-management client-cert disable`
- `web-management ocsp disable`

Checking Certificate Expiration

To activate periodic checks of certificate's expiration:

1 Navigate to **MANAGE | System Setup > Appliance > Base Settings**.

2 Scroll to **Check Certificate Expiration Settings**.



3 Select **Enable periodic certificate expiration check**. This option is selected by default. When enabled, the **Certificate expiration alert interval** field becomes available.

4 To set the interval between certificate checks, in hours, enter the interval in the **Certificate expiration alert interval: 1 - 168 (in hours)** field. The minimum time is 1 hour, the maximum is 168 hours, and the default is **168**.

5 Click **ACCEPT**.

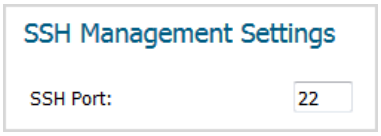
Configuring SSH Management

If you use SSH to manage the firewall, you can change the SSH port for additional security.

To change the SSH port:

1 Navigate to **MANAGE | System Setup > Appliance > Base Settings**.

2 Scroll to **SSH Management Settings**.



3 Enter the port in the **SSH Port** field. The default SSH port is **22**.

4 Click **ACCEPT**.

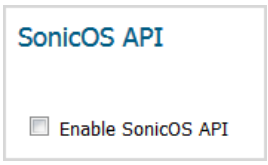
Enabling SonicOS API

You can use SonicOS API as an alternative to the SonicOS Command Line Interface (CLI) for configuring selected functions. To do so, you must first enable SonicOS API. For more information about SonicOS API, see the *SonicOS API Reference*.

To enable SonicOS API:

1 Navigate to **MANAGE | System Setup | Appliance > Base Settings**.

2 Scroll to **SonicOS API**.



3 Select **Enable SonicOS API**. This option is not selected by default.

4 Click **ACCEPT**.

Configuring Advanced Management Options

Advanced Management options allow you to specify:

- Whether the SonicWall Security Appliance is managed by SNMP (default) or the SonicWall Global Management System (GMS). For more information about GMS, see the *SonicWall GMS* and *SonicWall Management Services* administration documentation, available at https://www.sonicwall.com/support/technical-documentation.
- The creation of a Management Interface address object for the MGMT interface.

  This Management Interface provides a trusted interface to the management appliance. Network connections to this interface is very limited. If the NTP, DNS, and SYSLOG servers are configured in the MGMT subnet, the appliance uses the MGMT IP as the source IP and creates MGMT address object and route policies automatically. All traffic from the Management Interface is routed by this policy. Created routes display on the **System Setup > Network > Routing** page (for more information about routing, see Configuring Route Advertisements and Route Policies on page 450).

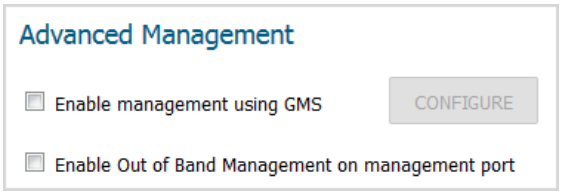
  The MGMT address object and route policies create/update IPv4 management IP. As the IPv6 management IP address object is created by default, this feature doesn't work on IPv6 management IP address object creation.

NOTE: By default, neither of these options is enabled.

To configure advanced management options:

1 Navigate to **MANAGE | System Setup > Appliance > Base Settings**.

2 Scroll to **Advanced Management**.



3 To allow SonicWall GMS to manage the firewall, select **Enable management using GMS**. **Configure** becomes available. For how to configure GMS management, see Enabling GMS Management on page 46.

4 To enables automatic creation of a Management Interface address object for the MGMT interface, which works as an out-of-band interface, and configures a route policy for the newly created address object, select **Out of Band Management on the management port**. Click **ACCEPT**.

IMPORTANT: To avoid confliction for delete/create route policies, updating this option to create a Management Interface address object and configure route policy causes system reboot.

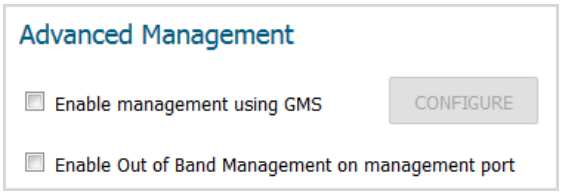
Enabling GMS Management

NOTE: For more information on SonicWall Global Management System, see the *SonicWall GMS* and *SonicWall Management Services* administration documentation, available at https://www.sonicwall.com/support/technical-documentation.

To configure the Security Appliance for GMS management:

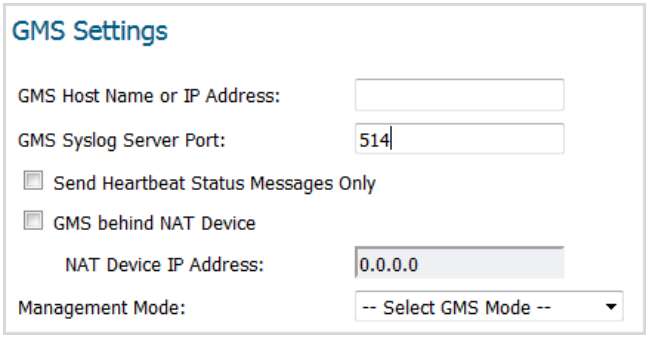
1 Navigate to **MANAGE | System Setup > Appliance > Base Settings**.

2 Scroll to **Advanced Management**.



3 Select **Enable Management using GMS**. **CONFIGURE** becomes available.

4 Click **CONFIGURE**. The **Configure GMS Settings** dialog displays.



5 Enter the host name or IP address of the GMS Console in the **GMS Host Name or IP Address** field.

6 Enter the port in the **GMS Syslog Server Port** field. The default value is **514**.

7 To send only heartbeat status instead of log messages, select **Send Heartbeat Status Messages Only**. This option is disabled by default.

8 If the GMS Console is placed behind a device using NAT on the network, select **GMS behind NAT Device**. This option is disabled by default.

When this option is selected, the **NAT Device IP Address** field becomes active.

a Enter the IP address of the NAT device in the **NAT Device IP Address** field.

9 Select one of the following GMS modes from **Management Mode**.

IPSEC Management Tunnel	Allows the firewall to be managed over an IPsec VPN tunnel to the GMS management console. Go to Step 10.
Existing Tunnel	Uses an existing VPN tunnel over the connection between the GMS server and the firewall. A message displays:  Go to Step 12.
HTTPS	Allows HTTPS management from two IP addresses: the GMS Primary Agent and the Standby Agent IP address. The SonicWall firewall also sends encrypted syslog packets and SNMP traps using 3DES and the firewall administrator's password. Options for configuring the GMS reporting server display. Go to Step 11.

10 The default IPsec VPN settings are displayed with values populated by SonicOS. Verify the settings.



a From **Encryption Algorithms**, select the appropriate algorithm.

b Optionally, enter a new encryption key in the **Encryption Key** field:

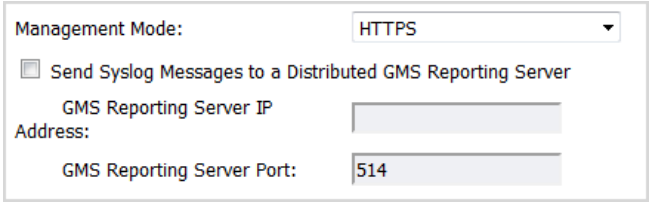
For	The key must be
DES	16 hexadecimal characters
3DES	48 hexadecimal characters

c Optionally, enter a new authentication key in the **Authentication Key** field:

For	The key must be
MD5	32 hexadecimal characters
SHA1	40 hexadecimal characters

d Go to Step 12.

11 SonicOS needs to know the GMS reporting server.



a Select **Send Syslog Messages to a Distributed GMS Reporting Server.** This option is not selected by default. The following options become available.

b In the **GMS Reporting Server IP Address** field, enter the IP address of the GMS server.

c In the **GMS Reporting Server Port** field, enter the port of the GMS server. The default port is **514**.

12 Click **OK**.

13 Click **ACCEPT**.

Downloading SonicPoint Images Manually

The **Download URL** section provides a field for specifying the URL address of a site for downloading SonicPoint images.

If your firewall:

- Has Internet connectivity, it automatically downloads the correct version of the SonicPoint image from the SonicWall server when you connect a SonicPoint device.
- Does not have Internet access, or has access only through a proxy server, you must manually specify a URL for the SonicPoint firmware. You do not need to include the `http://` prefix, but you do need to include the filename at the end of the URL. The filename should have a `.bin` extension. Here are examples using an IP address and a domain name:

  ```
  192.168.168.10/imagepath/sonicpoint.bin
  software.SonicWall.com/applications/sonicpoint/sonicpoint.bin
  ```

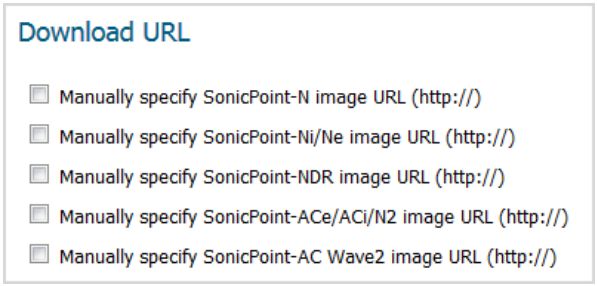
  For more information see *SonicOS 6.5 Updates*.

CAUTION: It is imperative that you download the corresponding SonicPoint image for the SonicOS firmware version that is running on your SonicWall Security Appliance. The MySonicWall website provides information about the corresponding versions. When upgrading your SonicOS firmware, be sure to upgrade to the correct SonicPoint image.

To select the type of SonicPoint image or images to download:

1 Navigate to **Manage | System Setup > Appliance > Base Settings**.

2 Scroll to **Download URL**.



- **Manually specify SonicPoint-N image URL (http://)**
- **Manually specify SonicPoint-Ni/Ne image URL (http://)**
- **Manually specify SonicPoint-NDR image URL (http://)**
- **Manually specify SonicPoint-ACe/ACi/N2 image URL (http://)**
- **Manually specify SonicPoint-AC Wave2 image URL (http://)**

3 Click the appropriate SonicPoint image URL. A field displays for that URL.



4 Enter the image download location in the associated field.

5 Click **ACCEPT**.

Selecting a Language

If your firmware contains other languages besides English, one can be selected from **Language Selection**.

NOTE: Changing the language of the SonicOS Management Interface requires that the SonicWall Security Appliance be rebooted.

To select a language for the Management Interface:

1 Navigate to **Manage | System Setup > Appliance > Base Settings**.

2 Scroll to **Language**.



3 Select the language from **Language Selection**.

4 Click **ACCEPT**.

3

Administering SNMP

- About Appliance > SNMP on page 50
  - About SNMP on page 50
  - Setting Up SNMP Access on page 51
  - Configuring SNMP as a Service and Adding Rules on page 59
  - About SNMP Logs on page 60

About Appliance > SNMP

You can manage the SonicWall Security Appliance using SNMP or SonicWall Global Management System (GMS). This section describes how to configure the SonicWall for management using SNMP. For information about managing the SonicWall appliance with GMS, see the *SonicWall GMS* and *SonicWall Management Services* administration documentation, available at https://www.sonicwall.com/support/technical-documentation.

Topics:

- About SNMP on page 50
- Setting Up SNMP Access on page 51
- Configuring SNMP as a Service and Adding Rules on page 59
- About SNMP Logs on page 60

About SNMP

SNMP (Simple Network Management Protocol) is a network protocol used over User Datagram Protocol (UDP) that allows network administrators to monitor the status of the SonicWall Security Appliance and receive notification of critical events as they occur on the network. The SonicWall Security Appliance supports SNMP v1/v2c/v3 and all relevant Management Information Base II (MIB-II) groups except **egp** and **at**.

SNMPv3 expands on earlier versions of SNMP and provides secure access to network devices by means of a combination of authenticating and encrypting packets.

Packet security is provided through:

- **Message Integrity**: ensures a packet has not been tampered with in transit
- **Authentication**: verifies a message comes from a valid source
- **Encryption**: encodes packet contents to prevent its being viewed by an unauthorized source.

SNMPv3 provides for both security models and security levels. A security model is an authentication strategy set up between a user and the group in which the user resides. The security level is the permitted level of security within a given security model. The security model and associated security level determine how an SNMP packet

is handled. SNMPv3 provides extra levels of authentication and privacy, as well as additional authorization and access control.

Security Level, Authentication, and Encryption Based on SNMP Version shows how security levels, authentication, and encryption are handled by the different versions of SNMP.

Security Level, Authentication, and Encryption Based on SNMP Version

Version	Level	Authentication Type	Encryption	Means of Authentication
v1	noAuthNoPriv	Community String	No	Community string match
v2c	noAuthNoPriv	Community String	No	Community string match
v3	noAuthNoPriv	Username	No	Username match
	authNoPriv	MD5 or SHA	No	Authentication is based on the HMAC-MD5 or HMSC-SRA algorithms.
	authPriv	MD5 or SHA	DES or AES	Provides authentication is based on the HMAC-MD5 or HMSC-SRA algorithms. Provides DES 56-bit encryption in addition to authentication based on the CBC-DES (DES-56) standard, or AES 128-bit encryption, as well.

The SonicWall Security Appliance replies to SNMP `Get` commands for MIB-II, using any interface, and supports a custom SonicWall MIB for generating trap messages. The custom SonicWall MIB is available for download from the SonicWall Web site and can be loaded into third-party SNMP management software such as HP Openview, Tivoli, or SNMPC.

You can view and configure SNMP settings. Settings cannot be viewed or modified by the user. SNMPv3 can be modified at the User or Group level. Access Views can be read, write, or both, and can be assigned to users or groups. A single View can have multiple Object IDs (OIDs) associated with it.

SNMPv3 settings for the SNMPv3 Engine ID are configurable under the **General Settings** menu of the **Configure SNMP** dialog. The Engine ID is used to authorize a received SNMP packet. Only matching packet EngineIDs are processed.

Setting Up SNMP Access

Setting up SNMP consists of:

- Enabling and Configuring SNMP Access on page 51
- Setting Up SNMPv3 Groups and Access on page 54

Enabling and Configuring SNMP Access

You can use either SNMPv1/v2 for basic functionality or configure the SonicWall Security Appliance to use the more extensive SNMPv3 options.

To use SNMP, you must first enable it.

Topics:

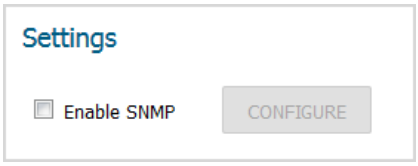
- Configuring Basic Functionality on page 52
- Configuring SNMPv3 Engine IDs on page 53
- Configuring Object IDs for SNMPv3 Views on page 55

- Creating Groups and Adding Users on page 56
- Adding Access on page 58

Configuring Basic Functionality

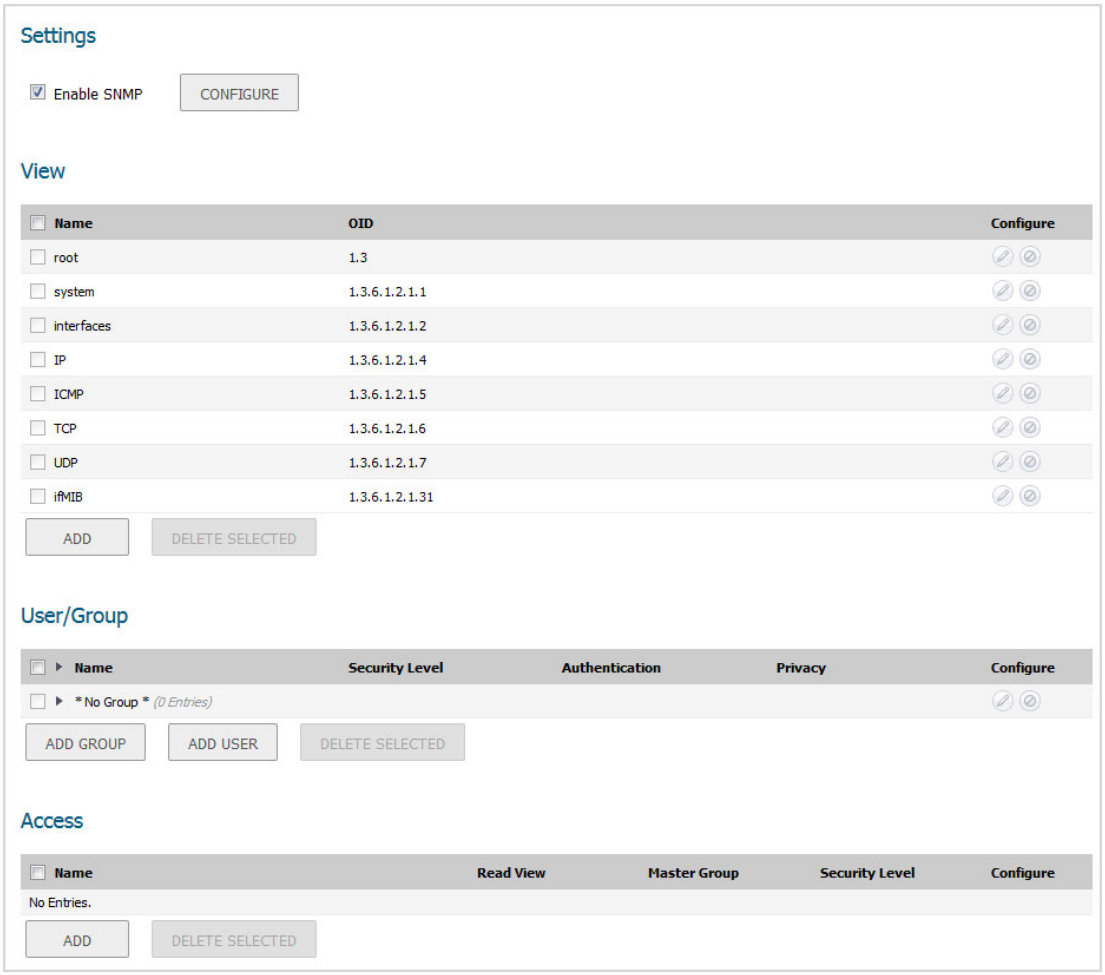
To enable SNMP:

1 Navigate to the **MANAGE | System Setup > Appliance > SNMP** page.

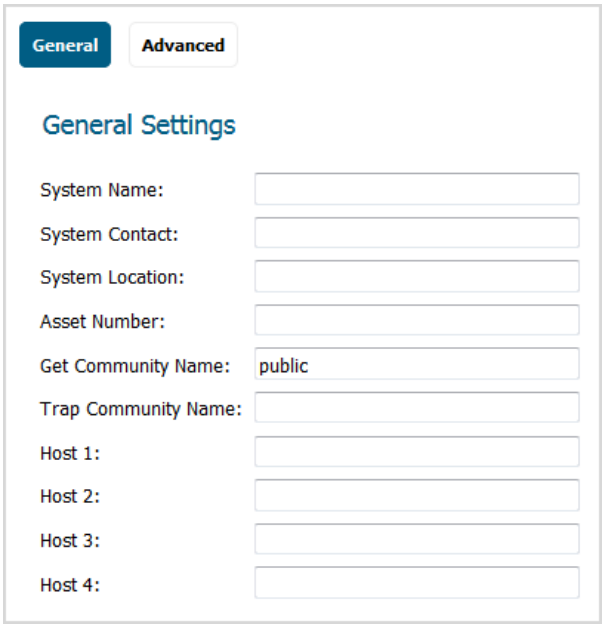


2 Select **Enable SNMP**. By default, SNMP is disabled.

3 Click **ACCEPT**. The SNMP information is populated on the SNMP page, and **CONFIGURE** becomes available.



4 To configure the SNMP interface, click **CONFIGURE**. The **Configure SNMP** dialog displays.



5 On the **General** page, enter the host name of the SonicWall Security Appliance in the **System Name** field.

6 Optionally, enter the network administrator's name in the **System Contact** field.

7 Optionally, enter an email address, telephone number, or pager number in the **System Location** field.

8 If the SNMPv3 configuration option is used, enter an asset number in the **Asset Number** field. Otherwise, this field is optional.

9 Enter a name for a group or community of administrators who can view SNMP data in the **Get Community Name** field. The default name is **public**.

10 Optionally, enter a name for a group or community of administrators who can view SNMP traps in the **Trap Community Name** field.

11 Enter the IP address(es) or host name(s) of the SNMP management system receiving SNMP traps in the **Host 1** through **Host *n*** fields. You must configure at least one IP address or host name, but up to the maximum number of addresses or host names for your system can be used.

12 If you:

- Want to set up SNMPV3, go to Configuring SNMPv3 Engine IDs on page 53.
- Finished setting up SNMP for now, click **OK**.

Configuring SNMPv3 Engine IDs

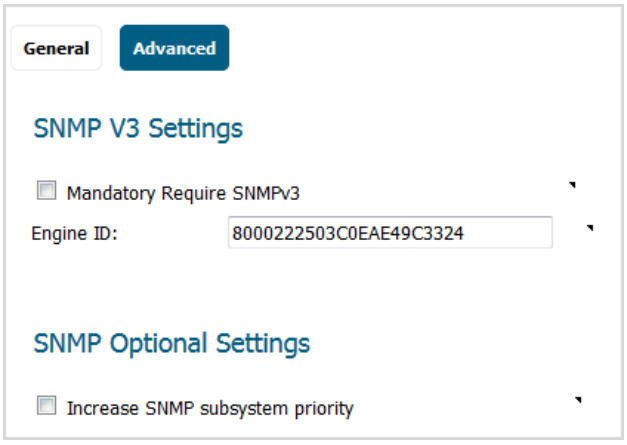
If SNMPv3 is used, you can configure the SNMPv3 Engine ID and SNMP priority. Configuring the SNMPv3 Engine ID provides maximum security for SNMP management.

To configure SNMPv3 engine IDs:

1 Navigate to **MANAGE | System Setup > Appliance > SNMP**.

2 If you have not configured SNMP for your system, follow Step 1 through Step 11 in Configuring Basic Functionality on page 52.

3 Click **Advanced**. The **Advanced** page displays.



4 Select **Mandatory Require SNMPv3**. This disables SNMPv1/v2 and allows only SNMPv3 access, which provides maximum security for SNMP management.

IMPORTANT: If you select this option, you must specify an asset number on the **General** page before clicking **OK**.

5 Enter the hexadecimal Engine ID number in the **Engine ID** field. SonicOS automatically populates this field, but you can change it. This number is matched against received SNMP packets to authorize their processing; only packets whose Engine ID matches this number are processed.

6 Optionally, select the **Increase SNMP subsystem** priority checkbox.

For efficient system operation, certain operations might take priority over responses to SNMP queries. Enabling this option causes the SNMP subsystem to always respond and operate at a higher system priority.

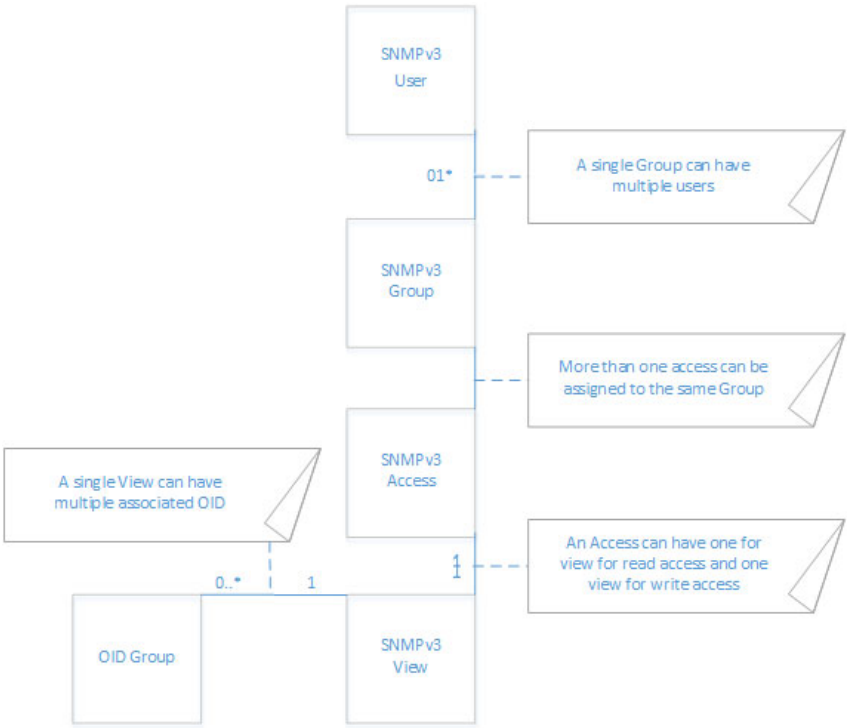
IMPORTANT: Enabling this option might affect the performance of the overall system.

7 Click **OK**. The SNMPv3 security options are now used in processing packets.

Setting Up SNMPv3 Groups and Access

SNMPv3 allows you to set up and assign groups and access with differing levels of security. Object IDs are associated with various levels of permissions, and a single view can be assigned to multiple objects. SNMPv3 group and user access shows how access for groups and users are associated with these different permission levels.

SNMPv3 group and user access



Configuring Object IDs for SNMPv3 Views

The SNMPv3 **View** shows access settings for Users and Groups. You create settings for users and groups, and these security settings are not user-modifiable. The SNMPv3 View defines the Object IDs (OID) and Object ID Groups, and is sometimes known as the SNMPv3 Access Object.

The SNMP **View** defines a collection of OIDs and OID groups. The initial set of default views cannot be changed or deleted. The default views reflect the most often used views, such as the root view, system view, IP, interfaces. The OIDs for these views are pre-assigned.

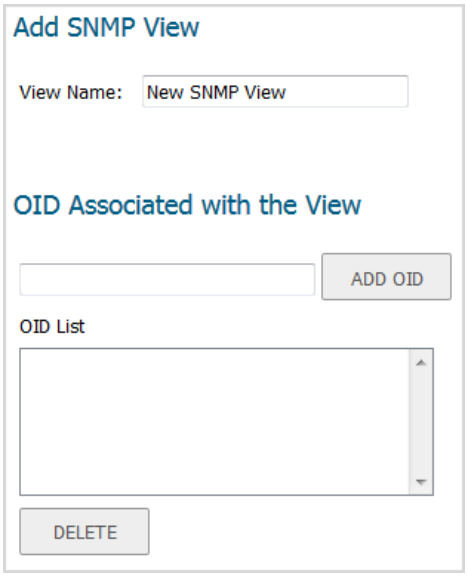
Additionally, you can create a custom view for specific users and groups.

You can modify views you create. You cannot modify the ones the system creates.

To configure OIDs for SNMPv3 views:

1 Navigate to **MANAGE | System Setup > Appliance > SNMP**.

2 To add a view, in the **View** section, click **ADD**. The **Add SNMP View** dialog displays.



3 Enter a meaningful name in the **View Name** field. The default name is **New SNMP View**.

NOTE: If editing an existing view, the name is not editable.

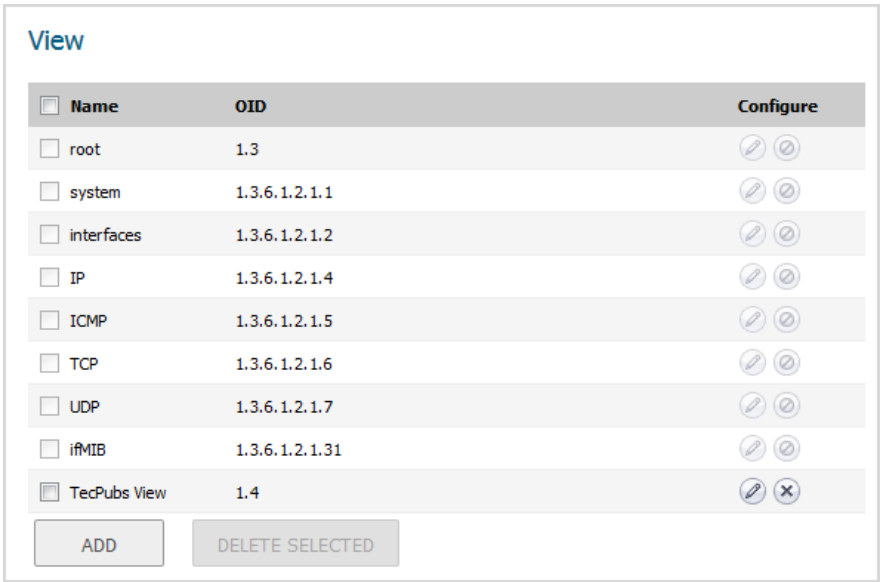
4 Enter an unassigned OID in the **OID Associated with the View** field.

5 Click **ADD OID**.

The new view appears in the **OID List**. To delete an OID from the OID List, select the OID and click **Delete**.

6 Add any more new views with associated OIDs.

7 Click **OK**. The new views are added to the **View** table.



Creating Groups and Adding Users

By default, there is one group, ***No Group***, that cannot be configured or deleted. You can, however, add users to this default group.

Topics:

- Creating a Group on page 57
- Adding Users on page 57

Creating a Group

To create a group:

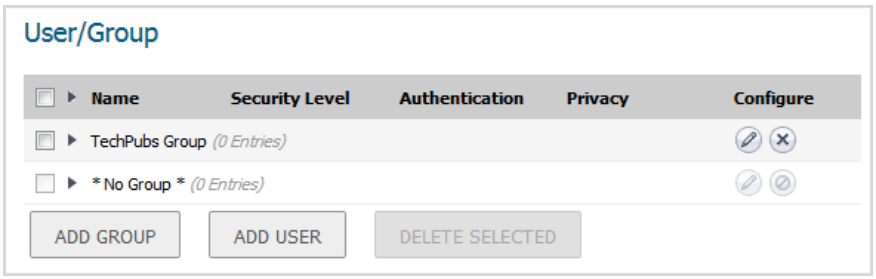
1 Navigate to **MANAGE | System Setup > Appliance > SNMP**.

2 Click **ADD GROUP** under the **User/Group** table. The **Add SNMP Group** dialog displays.



3 Enter a friendly name in the **Group Name** field. The group name can contain up to 32 alphanumeric characters.

4 Click **OK**. The **User/Group** table is updated, and the **Edit** and **Delete** icons in the **Configure** column are available.

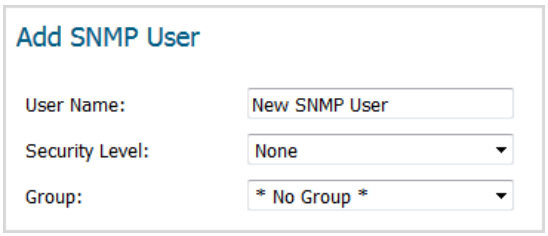


Adding Users

To add users:

1 Navigate to **MANAGE | System Setup > Appliance > SNMP**.

2 Click **ADD USER** under the **User/Group** table. The **Add SNMP User** dialog displays.

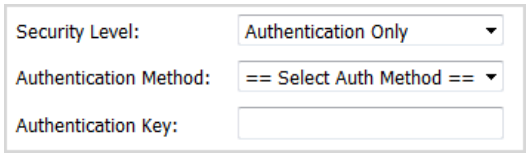


3 Enter the user name in the **User Name** field.

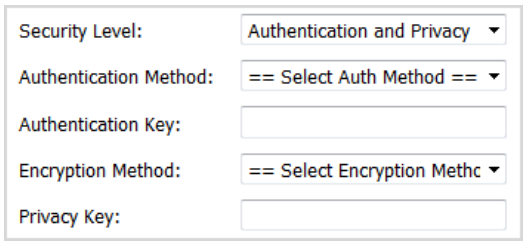
4 Select a security level from **Security Level**:

- **None** (default)

- **Authentication –** Two new options appear:



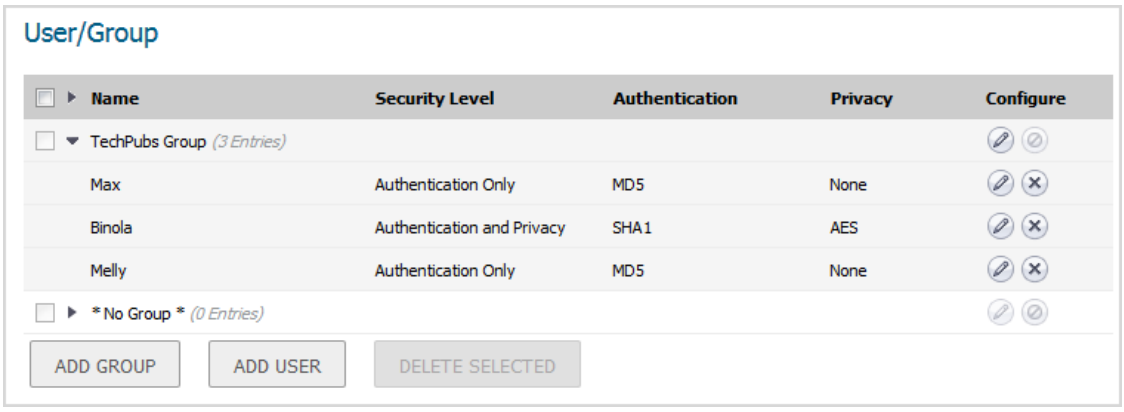
  - **Authentication Method** – Select one of these authentication methods: **MD5** or **SHA1**.
  - **Authentication Key** – Enter an authentication key in the field. The key can be any string of 8 to 32 printable characters.
- **Authentication and Privacy** – More options appear:



  - **Authentication Method** – See above.
  - **Authentication Key** – See above.
  - Select an encryption method from the **Encryption Method** drop-down menu: **AES** or **DES**.
  - Enter the encryption key in the **Privacy Key** field. The key can be any string of 8 to 32 printable characters.

5 Select a group from **Group**. The default is ***No Group***.

6 Click **OK** when finished. The user is added to the **User/Group** table and added to the appropriate group (including ***No Group***).



Adding Access

SNMPv3 Access is an object that:

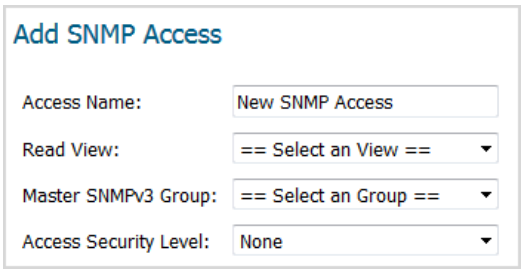
- Defines the read/write access rights of an SNMPv3 View.
- Can be assigned to an SNMPv3 Group.

Multiple groups can be assigned to the same Access object. An Access object can also have multiple views assigned to it.

To create an access object:

1 Navigate to **MANAGE | System Setup > Appliance > SNMP**.

2 Under the **Access** table, click **ADD**. The **Add SNMP Access** dialog displays.



3 Enter a friendly name in the **Access Name** field.

NOTE: Names of existing entries are non-editable.

4 From **Read View**, select a view from the list of available views.

5 From **Master SNMPv3 Group**, select a group from the list of available groups.

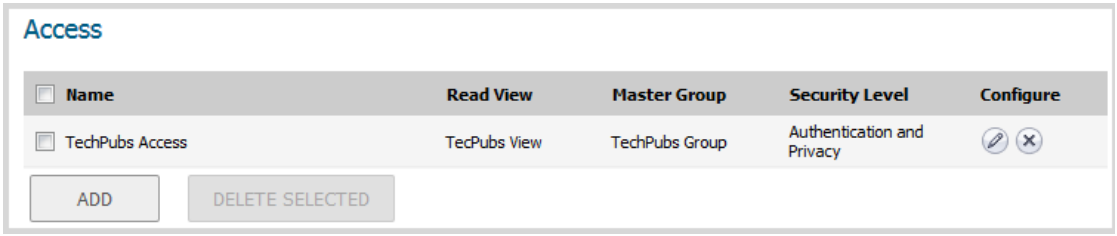
NOTE: Access can be assigned to only one SNMPv3 group, but a group can be associated with multiple Access objects.

Access cannot be given to ***No Group***.

6 From **Access Security Level**, select a security level:

- **None**
- **Authentication Only**
- **Authentication and Privacy**

7 Click **OK**. The Access object is added to the **Access** table.



Configuring SNMP as a Service and Adding Rules

By default, SNMP is disabled on the SonicWall Security Appliance. To enable SNMP, you must first enable SNMP on the **MANAGE | System Setup > Appliance > SNMP** page, and then enable it for individual interfaces. To do this, go to the **MANAGE | System Setup > Network > Interfaces** page and click on **Configure** for the interface you want to enable SNMP on. For more information about configuring SNMP as a service and adding rules, see Configuring Interfaces on page 259.

If your SNMP management system supports discovery, the SonicWall Security Appliance agent automatically discovers the SonicWall Security Appliance on the network. Otherwise, you must add the SonicWall Security Appliance to the list of SNMP-managed devices on the SNMP management system.

About SNMP Logs

SNMP logs can be viewed on the **INVESTIGATE | Log > Event Logs** page. For more information about Event Logs, see *SonicOS 6.5 Investigate*.

Trap messages are generated only for the alert message categories normally sent by the SonicWall Security Appliance. For example, attacks, system errors, or blocked Web sites generate trap messages. If none of the categories are selected on the **INVESTIGATE | Log > Event Logs** page, then no trap messages are generated.

4

Managing Certificates

Topics:

- About Certificates on page 61
  - About Digital Certificates on page 61
  - About the Certificates and Certificate Requests Table on page 63
  - Importing Certificates on page 65
  - Deleting a Certificate on page 67
  - Generating a Certificate Signing Request on page 67
  - Configuring Simple Certificate Enrollment Protocol on page 71

About Certificates

To implement the use of certificates for VPN policies, you must locate a source for a valid CA certificate from a third party CA service. When you have a valid CA certificate, you can import it into the firewall to validate your Local Certificates. You import the valid CA certificate into the firewall using the **MANAGE | System Setup > Appliance > Certificates** page. After you import the valid CA certificate, you can use it to validate your local certificates.

SonicOS provides a large number of certificates with the SonicWall Security Appliance; these are built-in certificates and cannot be deleted or configured.

SonicOS supports a local Certificate Revocation List (CRL), which is a list of digital certificates that have been revoked by the issuing Certificate Authority (CA) before their scheduled expiration date and should no longer be trusted. For further information about local CRL, contact Technical Support.

About Digital Certificates

A digital certificate is an electronic means to verify identity by a trusted third party known as a Certificate Authority (CA). The X.509 v3 certificate standard is a specification used with cryptographic certificates and allows you to define extensions that you can include with your certificate. SonicWall has implemented this standard in its third-party certificate support.

You can use a certificate signed and verified by a third-party CA to use with an IKE (Internet Key Exchange) VPN policy. IKE is an important part of IPsec VPN solutions, and it can use digital certificates to authenticate peer devices before setting up Security Associations (SAs). Without digital certificates, VPN users must authenticate by manually exchanging shared secrets or symmetric keys. Devices or clients using digital signatures do not require configuration changes every time a new device or client is added to the network.

A typical certificate consists of two sections: a data section and a signature section. The data section typically contains information such as the version of X.509 supported by the certificate, a certificate serial number, information about the user's public key, the Distinguished Name (DN), validation period for the certificate, and

optional information such as the target use of the certificate. The signature section includes the cryptographic algorithm used by the issuing CA, and the CA digital signature.

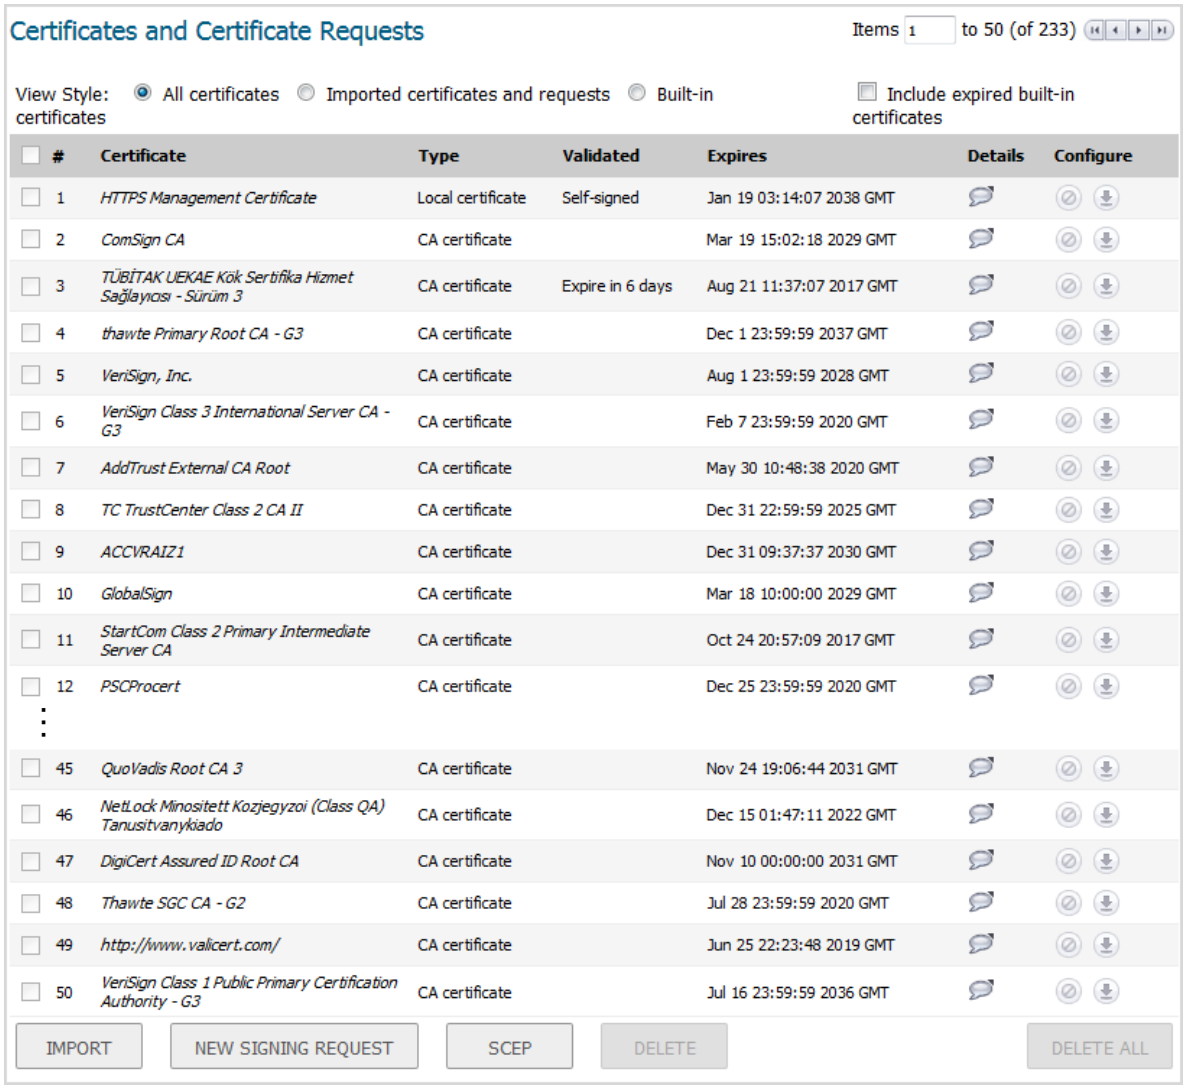
SonicWall Security Appliances interoperate with any X.509v3-compliant provider of Certificates. SonicWall Security Appliance have been tested with the following vendors of Certificate Authority Certificates:

- Entrust
- Microsoft
- OpenCA
- OpenSSL and TLS
- VeriSign

Topics:

- About the Certificates and Certificate Requests Table on page 63
- Importing Certificates on page 65
- Deleting a Certificate on page 67
- Generating a Certificate Signing Request on page 67
- Configuring Simple Certificate Enrollment Protocol on page 71

About the Certificates and Certificate Requests Table



The **Certificate and Certificate Requests** table provides all the settings for managing CA and Local Certificates.

The **View Style** menu allows you to display your certificates based on these criteria:

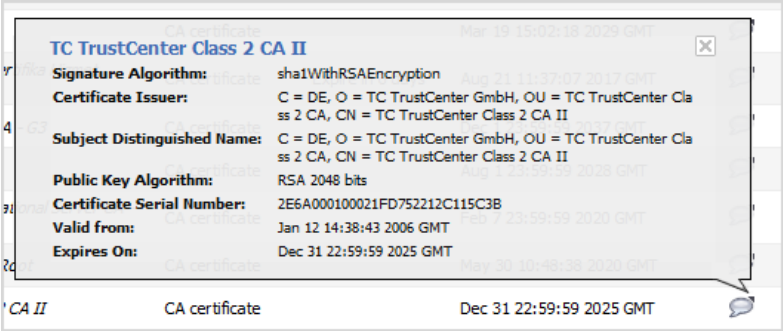
This criterion	Displays
All Certificates	All built-in and imported certificates and certificate requests. This is the default.
Imported certificates and requests	Only imported certificates and generated certificate requests. This option is not selected by default.
Built-in certificates	Only built-in certificates. This option is not selected by default.
Include expired and built-in certificates	All expired and built-in certificates. This option is not selected by default.

The **Certificates and Certificate Requests** table displays this information about certificates:

This column	Displays the
Certificate	Name of the certificate.
Type	Type of certificate: • **CA certificate** • **Local certificate** • **Pending request**
Validated	Validation information: • Blank • **Self-signed** • **Expire in *n* days** • **Expired**
Expires	Date and time the certificate expires.
Details	Details of the certificate. Moving the pointer over the **Comment** icon displays the details of the certificate. For information about certificate details, see About Certificate Details on page 64.
Configure	Contains the • **Delete** icon to delete a certificate entry • **Import** icon to import either certificate revocation lists (for CA certificates) or signed certificates (for Pending requests). **NOTE:** You cannot delete or import built-in certificates.

About Certificate Details

Clicking on the **Comment** icon in the **Details** column displays information about the certificate, which might include the following, depending on the type of certificate:



- Signature Algorithm
- Certificate Issuer
- Subject Distinguished Name
- Public Key Algorithm
- Certificate Serial Number
- Valid from
- Expires On
- Status (for Pending requests and local certificates)

The details depend on the type of certificate. **Certificate Issuer**, **Certificate Serial Number, Valid from**, and **Expires On** are not shown for Pending requests as this information is generated by the Certificate provider.

Importing Certificates

After your CA service has issued a Certificate for your Pending request, or has otherwise provided a Local Certificate, you can import it for use in VPN or Web Management authentication. CA Certificates might also be imported to verify local Certificates and peer Certificates used in IKE negotiation.

Topics:

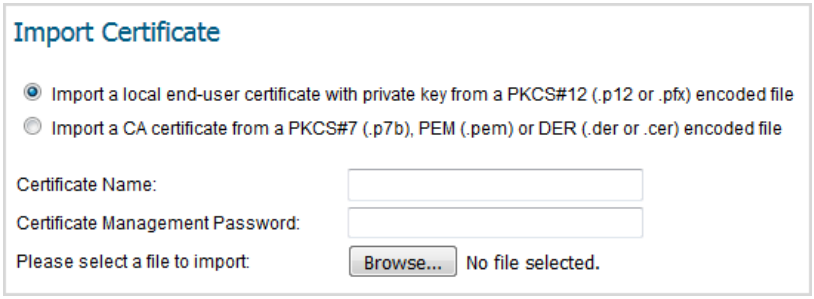
- Importing a Local Certificate on page 65
- Importing a Certificate Authority Certificate on page 65
- Creating a PKCS-12 Formatted Certificate File (Linux Systems Only) on page 66

Importing a Local Certificate

To import a local certificate:

1 Navigate to **MANAGE | System Setup > Appliance > Certificates**.

2 CLICK **IMPORT**. THE **IMPORT CERTIFICATE** DIALOG DISPLAYS.



3 Enter a certificate name in the **Certificate Name** field.

4 Enter the password used by your Certificate Authority to encrypt the PKCS#12 file in the **Certificate Management Password** field.

5 Click **Browse** to locate the certificate file.

6 Click **Open** to set the directory path to the certificate.

7 Click **Import** to import the certificate into the firewall. When it is imported, you can view the certificate entry in the **Certificates and Certificate Requests** table.

8 Moving your pointer to the **Comment** icon in the **Details** column displays the certificate details information.

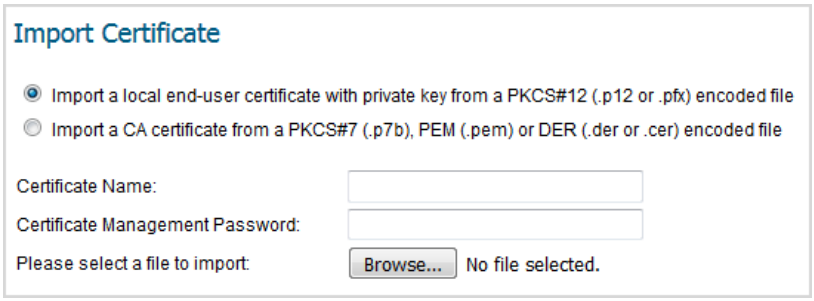
NOTE: If the certificate was uploaded successfully, the **Status** in the mouseover popup is **Verified**.

Importing a Certificate Authority Certificate

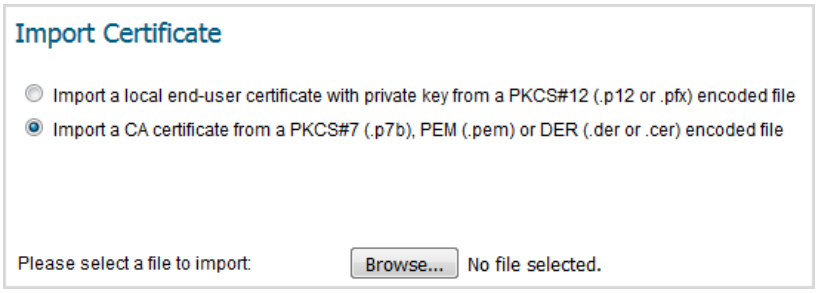
To import a certificate from a certificate authority:

1 Navigate to **MANAGE | System Setup > Appliance > Certificates**.

2 Click **IMPORT**. The **Import Certificate** dialog displays.



3 Choose **Import a CA certificate from a PKCS#7 (*.p7b) or DER (.der or .cer) encoded file**. The **Import Certificate** dialog settings change.



4 Click **Browse** to locate the certificate file.

5 Click **Open** to set the directory path to the certificate.

6 Click **Import** to import the certificate into the firewall. When it is imported, you can view the certificate entry in the **Certificates and Certificate Requests** table.

7 Moving your pointer to the **Comment** icon in the **Details** column displays the certificate details information.

Creating a PKCS-12 Formatted Certificate File (Linux Systems Only)

A PKCS12-formatted certificate file can be created using Linux system with OpenSSL. To create a PKCS-12 formatted certificate file, one needs to have two main components of the certificate:

- Private key (typically a file with `.key` extension or the word key in the filename)
- Certificate with a public key (typically a file with `.crt` extension or the word cert as part of filename).

For example, the Apache HTTP server on Linux has its private key and certificate in these locations:

- `/etc/httpd/conf/ssl.key/server.key`
- `/etc/httpd/conf/ssl.crt/server.crt`

With these two files available, run the following command:

```
openssl pkcs12 -export -out out.p12 -inkey server.key -in server.crt
```

In this example **`out.p12`** become the PKCS-12 formatted certificate file and **`server.key`** and **`server.crt`** are the PEM-formatted private key and the certificate file respectively.

After running the **`openssl`** command, you are prompted for the password to protect/encrypted the file. After choosing the password, the creation of the PKCS-12-formatted certificate file is complete, and it can be imported into the appliance.

Deleting a Certificate

NOTE: Built-in certificates cannot be deleted.

You can delete an imported certificate if it has expired or if you decide not to use third-party certificates for VPN authentication. You can always delete certificates you created.

To delete:

- A certificate, click its **Delete** icon.
- One or more certificates:
  a Click their checkbox(es). **DELETE** and **DELETE ALL** become available.
  b Click either **DELETE** or **DELETE ALL**.
- All non built-in certificates:
  a Click the checkbox in the table heading. **DELETE** and **DELETE ALL** become available.
  b Click either **DELETE** or **DELETE ALL**.

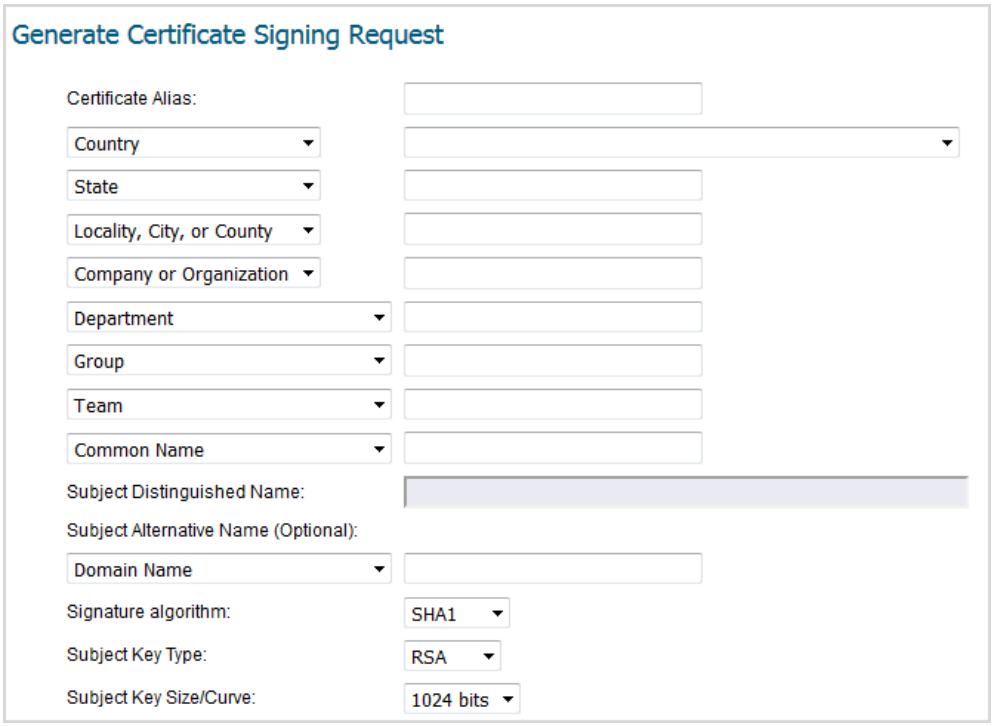
Generating a Certificate Signing Request

TIP: You should create a Certificate Policy to be used in conjunction with local certificates. A Certificate Policy determines the authentication requirements and the authority limits required for the validation of a certificate.

To generate a certificate signing request:

1 Navigate to **MANAGE | System Setup > Appliance > Certificates**.

2 Click **NEW SIGNING REQUEST**. The **Certificate Signing Request** dialog displays.



3 Enter an alias name for the certificate in the **Certificate Alias** field.

4 Create a Distinguished Name (DN) using the drop-down menus shown in Distinguished name components, then enter information for the certificate in the associated fields.

NOTE: For each DN, you can select your country from the associated drop-down menu; for all other components, enter the information in the associated field.

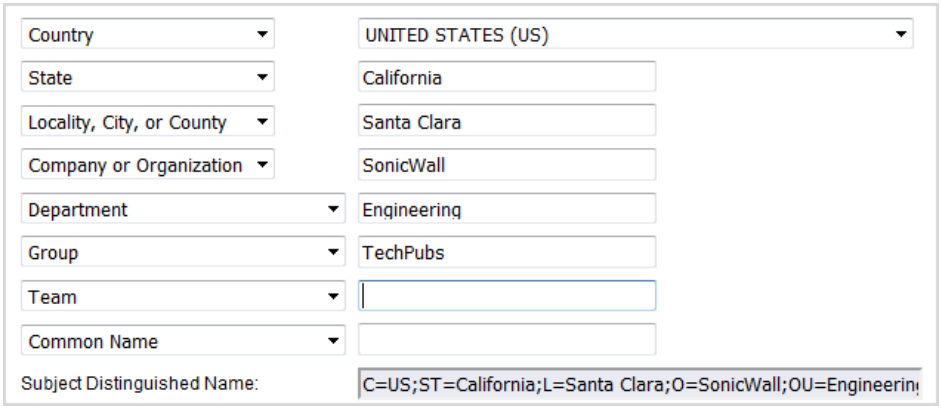
Distinguished name components

From this drop-down menu	Select/enter the appropriate information
Country	**Country** (default) **State** **Locality or County** **Company or Organization**
State	**Country** **State** (default) **Locality, City, or County** **Company or Organization** **Department**
Locality, City, or County	**Locality, City, or County** (default) **Company or Organization** **Department** **Group** **Team**
Company or Organization	**Company or Organization** (default) **Department** **Group** **Team** **Common Name** **Serial Number** **E-Mail Address**
Department	**Department** (default) **Group** **Team** **Common Name** **Serial Number** **E-Mail Address**
Group	**Group (default)** **Team** **Common Name** **Serial Number** **E-Mail Address**

Distinguished name components

From this drop-down menu	Select/enter the appropriate information
Team	**Team** (default)
	Common Name
	Serial Number
	E-Mail Address
Common Name	**Common Name** (default)
	Serial Number
	E-Mail Address

As you enter information for the components, the Distinguished Name (DN) is created in the **Subject Distinguished Name** field.



5 Optionally, you can also attach a **Subject Alternative Name** to the certificate after selecting the type from the drop-down menu:

- **Domain Name**
- **Email Address**
- **IPv4 Address**

6 Select a signature algorithm from the **Signature algorithm** drop-down menu:

- **MD5**
- **SHA1** (default)
- **SHA256**
- **SHA384**
- **SHA512**

7 Select a subject key type from the **Subject Key Type** drop-down menu:

RSA (default)	A public key cryptographic algorithm used for encrypting data,
ECDSA	Encrypts data using the Elliptic Curve Digital Signature Algorithm, which has a high strength-per-key-bit security.

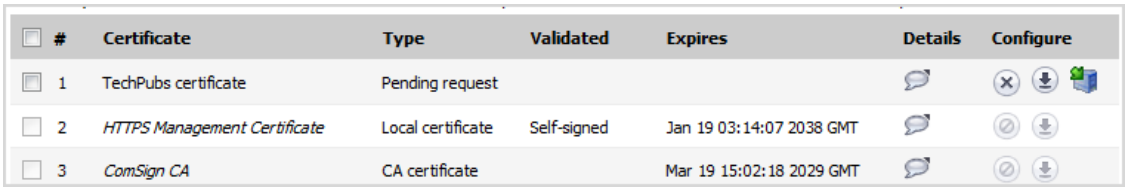
8 Select a subject key size or curve from the **Subject Key Size/Curve** drop-down menu.

NOTE: Not all key sizes or curves are supported by a Certificate Authority, therefore, you should check with your CA for supported key sizes.

If you selected a Key Type of	
RSA, select a key size	**ECDSA, select a curve**
1024 bits (default)	**prime256vi: X9.62.SECG curve over a 256 bit prime field** (default)
1536 bits	**secp384r1: NIST/SECG curve over a 384 bit prime field**
2048 bits	**secp521r1: NIST/SECG curve over a 521 bit prime field**
4096 bits	

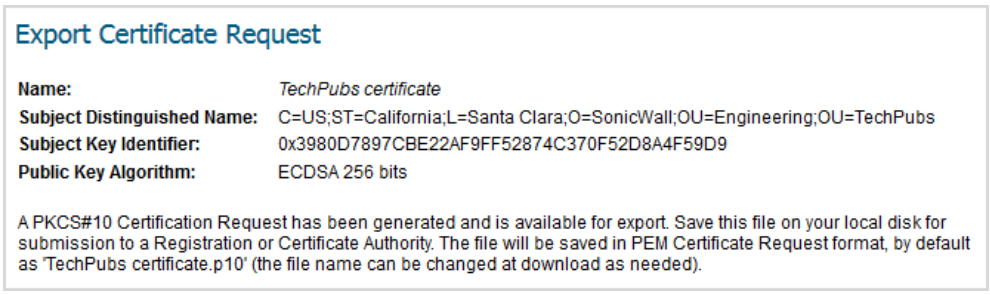
9 Click **GENERATE** to create a certificate signing request file.

When the **Certificate Signing Request** is generated, a message describing the result is displayed in the Status area at the bottom of the browser window and a new entry appears in the **Certificates and Certificate Requests** table with the type **Pending request**.



#	Certificate	Type	Validated	Expires	Details	Configure
1	TechPubs certificate	Pending request				
2	*HTTPS Management Certificate*	Local certificate	Self-signed	Jan 19 03:14:07 2038 GMT		
3	*ComSign CA*	CA certificate		Mar 19 15:02:18 2029 GMT		

10 Click the **Export** icon. The **Export Certificate Request** dialog displays.

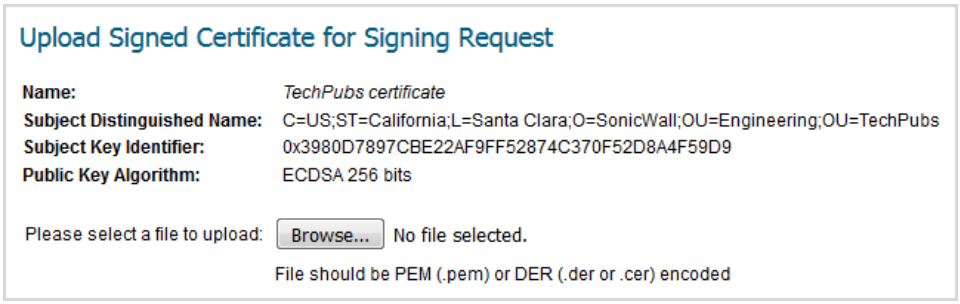


11 Click the **Export** icon to download the file to your computer. An **Opening *<certificate>*** dialog displays.

12 Click **OK** to save the file to a directory on your computer.

You have generated the **Certificate Request** that you can send to your Certificate Authority for validation.

13 Click the **Upload** icon to upload the signed certificate for a signing request. the **Upload Certificate** dialog displays.



14 Click **Browse** to select a file. The Open File dialog displays.

15 Select the file.

16 Click **Open**.

17 Click **UPLOAD**.

Configuring Simple Certificate Enrollment Protocol

The Simple Certificate Enrollment Protocol (SCEP) is designed to support the secure issuance of certificates to network devices in a scalable manner. There are two enrollment scenarios for SCEP:

- SCEP server CA automatically issues certificates.
- SCEP request is set to PENDING and the CA administrator manually issues the certificate.

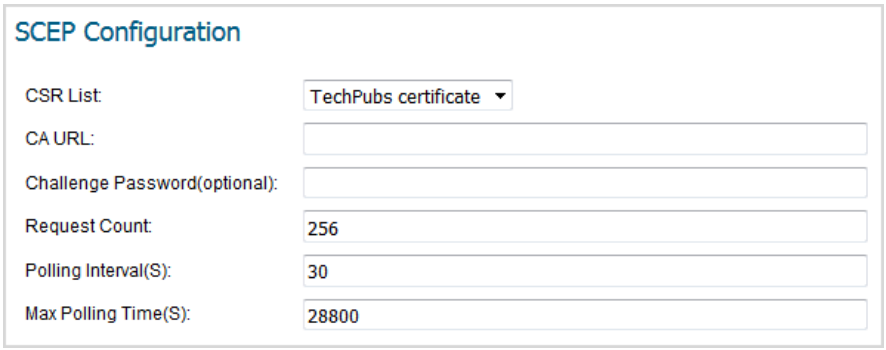
More information about SCEP can be found at: http://tools.ietf.org/html/draft-nourse-scep-18 (*Cisco Systems' Simple Certificate Enrollment Protocol draft-nourse-scep-18*).

To use SCEP to issue certificates:

1 Generate a signing request as described above in Generating a Certificate Signing Request on page 67.

2 Scroll to the bottom of the **Appliance > Certificates** page.

3 Click **SCEP**. The **SCEP Configuration** dialog displays.



4 From **CSR List**, SonicOS selects a default CSR list automatically. If you have multiple CSR lists configured, you can modify this.

5 In the **CA URL** field, enter the URL for the Certificate authority.

6 If the **Challenge Password(optional)** field, enter the password for the CA if one is required.

7 In the **Request Count** field, enter the number of requests. The default value is **256**.

8 In the **Polling Interval(S)** field, you can modify the default value for duration of time, in seconds, between the sending of polling messages. the default value is **30** seconds.

9 In the **Max Polling Time(S)** field, you can modify the default value for the duration of time, in seconds, the firewall waits for a response to a polling message before timing out. The default value is **28800** seconds (8 hours).

10 Click **SCEP** to submit the SCEP enrollment.

The firewall contacts the CA to request the certificate. The time this takes depends on whether the CA issues certificates automatically or manually. After the certificate is issued, it is displayed in the list of available certificates on the **Appliance > Certificates** page, under the **Imported certificates and requests** or **All certificates** category.

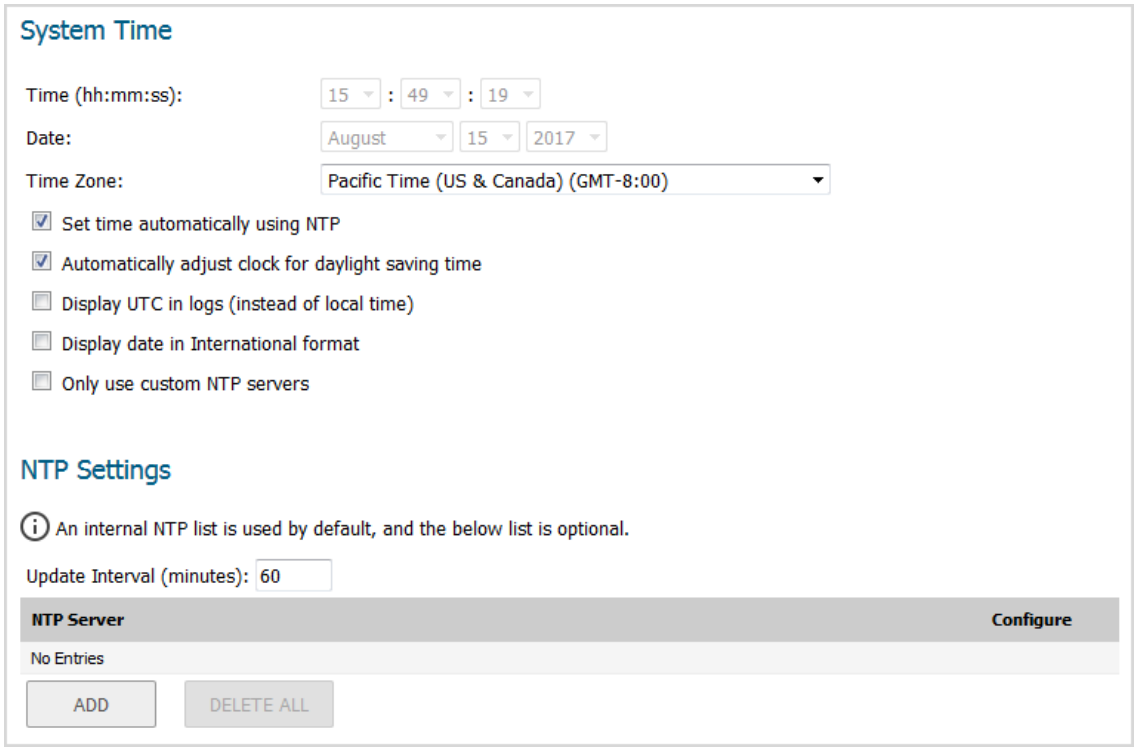
5

Configuring Time Settings

- About Appliance > Time on page 72
  - Setting System Time on page 73
  - Configuring NTP Settings on page 74

About Appliance > Time

MANAGE | System Setup > Appliance > Time defines the time and date settings to time stamp log events, to automatically update SonicWall Security Services, and for other internal purposes.



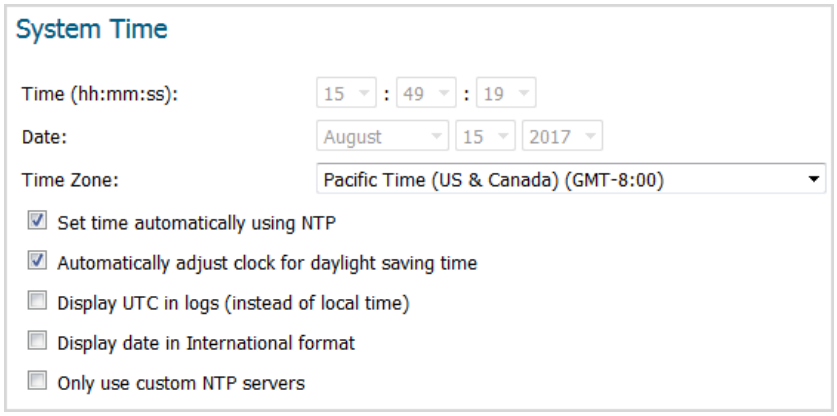
By default, the SonicWall Security Appliance uses an internal list of public NTP servers to update the time automatically. Network Time Protocol (NTP) is a protocol used to synchronize computer clock times in a network of computers. NTP uses Coordinated Universal Time (UTC) to synchronize computer clock times to a millisecond, and sometimes to a fraction of a millisecond.

Topics:

- Setting System Time on page 73
- Configuring NTP Settings on page 74

Setting System Time

You set the system time in the **System Time** section of **Appliance > Time**.



To set the system time:

1 Navigate to **MANAGE | System Setup > Appliance > Time**.

2 Select the time zone you are in from **Time Zone**.

3 To set the time:

- Automatically, select **Set time automatically using NTP** to use NTP (Network Time Protocol) servers from an internal list. This option is selected by default.
- Manually, clear **Set time automatically using NTP**. The **Time** and **Date** options become available.



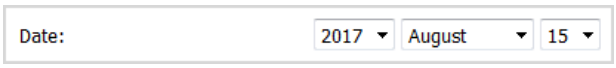
   1) Select the time in the 24-hour format using the **Time (hh:mm:ss)** drop-down menus

   2) Select the date from the **Date** drop-down menus.

4 To enable automatic adjustments for daylight savings time, select **Automatically adjust clock for daylight saving time**. For those areas that observe daylight savings time, this option is selected by default.

5 To use universal time (UTC) rather than local time for log events, select **Display UTC in logs (instead of local time)**. This option is not selected by default.

6 To display the date in International format, with the day preceding the month, select **Display date in International format**.



   This option is not selected by default.

7 To use the manually entered list of NTP servers to set the firewall clock rather than the internal list of NTP servers, select **Only use custom NTP servers**.

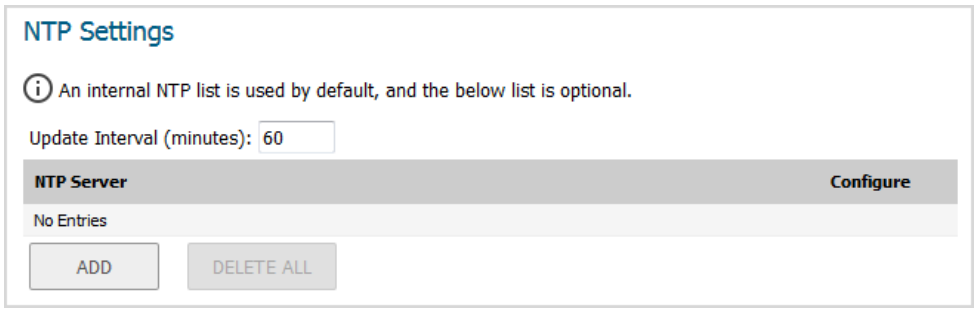
   IMPORTANT: Select this option only if you have configured one or more NTP servers. For more information about NTP servers, see Configuring NTP Settings on page 74.

8 Click **ACCEPT**.

Configuring NTP Settings

Network Time Protocol (NTP) is a protocol used to synchronize computer clock times in a network of computers. NTP uses Coordinated Universal Time (UTC) to synchronize computer clock times to a millisecond, and sometimes, to a fraction of a millisecond.

TIP: The SonicWall Security Appliance uses an internal list of NTP servers, so manually entering a NTP server is optional.



Topics:

- Using an NTP Server for Updating the Firewall Clock on page 74
- Adding an NTP Server on page 74
- Editing an NTP Server Entry on page 75
- Deleting NTP Server Entries on page 75

Using an NTP Server for Updating the Firewall Clock

To use a local server to set the firewall clock:

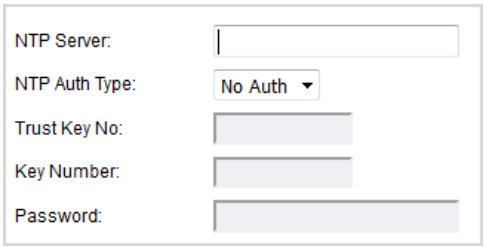
1. Navigate to **MANAGE | System Setup > Appliance > Time**.
2. Add one or more NTP servers as described in Configuring NTP Settings on page 74.
3. Select **Use NTP to set time automatically** (see Setting System Time on page 73). This option is not selected by default.
4. To configure the frequency for the NTP server to update the firewall, enter the interval in **Update Interval (minutes)**. The default value is **60** minutes.
5. Click **Accept**.

Adding an NTP Server

To add an NTP server to the firewall configuration

1. Navigate to **MANAGE | System Setup > Appliance > Time**.

2 In the **NTP Settings** section, click **ADD**. The **Add NTP Server** dialog displays.



3 Type the IP address of the remote NTP server in the **NTP Server** field.

4 Select the authentication type from the **NTP Auth Type** drop-down menu:

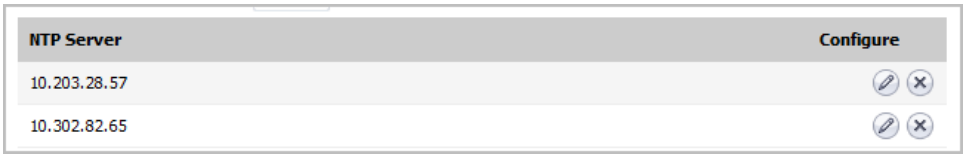
- **No Auth** – Authentication is not required and the following three options are dimmed. Go to Step 8.
- **MD5** – Authentication is required and the following three options are active.

5 Enter the Trust Key number in the **Trust Key No** field. The minimum is 1 and the maximum is 99999.

6 Enter the Key number in the **Key Number** field. The minimum is 1 and the maximum is 99999.

7 Enter the password in the **Password** field.

8 Click **OK**. The **NTP Server** section shows the server.



Editing an NTP Server Entry

To edit an NTP server entry:

1 Navigate to **MANAGE | System Setup > Appliance > Time**.

2 In the **NTP Server** table, click the entry's **Edit** icon. The **Edit NTP Server** dialog displays; it is the same as the **Add NTP Server** dialog; see Adding an NTP Server on page 74.

3 Make the changes.

4 Click **OK**.

Deleting NTP Server Entries

To delete an NTP server entry:

1 Navigate to **MANAGE | System Setup > Appliance > Time**.

2 In the **NTP Server** table, click the entry's **Delete** icon.

To delete all servers:

3 Navigate to **MANAGE | System Setup > Appliance > Time**.

4 Below the **NTP Server** table, click **DELETE ALL**.

6

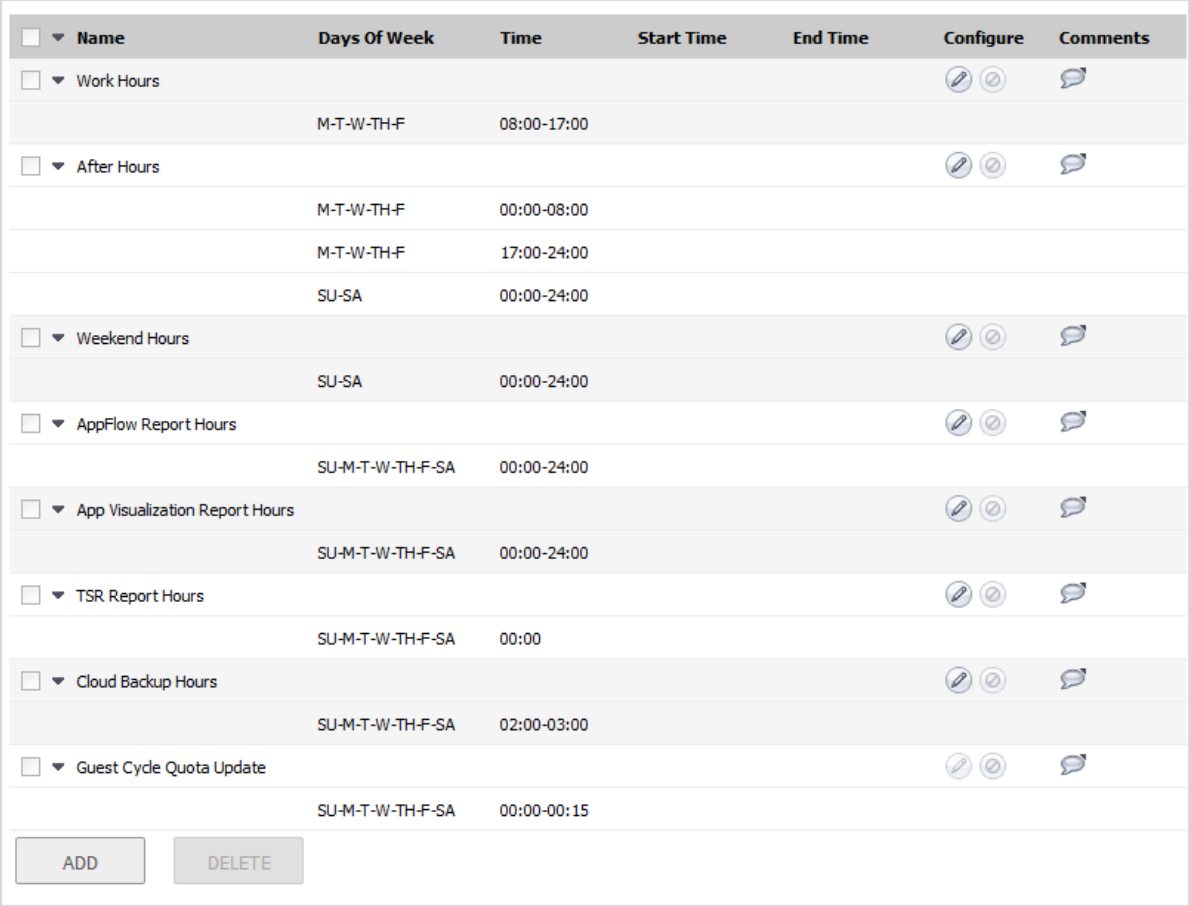
Setting Schedules

- About Schedules on page 76
- About Appliance > Schedules on page 77
  - Adding a Custom Schedule on page 78
  - Modifying Schedules on page 79
  - Deleting Custom Schedules on page 80

About Schedules

SonicOS uses schedule objects in conjunction with its security features and policies. You create schedule objects with **MANAGE | System Setup > Appliance > Schedules**. You apply schedule objects for a specific security feature or policy (rule). For example, if you add an access rule in the **MANAGE | Policies > Rules > Access Ruless** page, the **Add Rule** dialog lists all the available predefined schedule objects as well as the schedule objects you create with the **MANAGE | System Setup > Appliance > Schedules** page. A schedule can include multiple day and time increments for rule enforcement with a single schedule.

About Appliance > Schedules



MANAGE | System Setup > Appliance > Schedules allows you to create and manage default and custom schedule objects for enforcing schedule times for a variety of SonicWall Security Appliance features.

NOTE: You can modify default schedules, but you cannot delete them.

The **Schedules** table displays all predefined and custom schedules. The default schedules consist of:

Work Hours	**AppFlow Report Hours**	**Cloud Backup Hours**
After Hours	**App Visualization Report Hours**	**Guest Cycle Quota Update**
Weekend Hours	**TSR Report Hours**	

Topics:

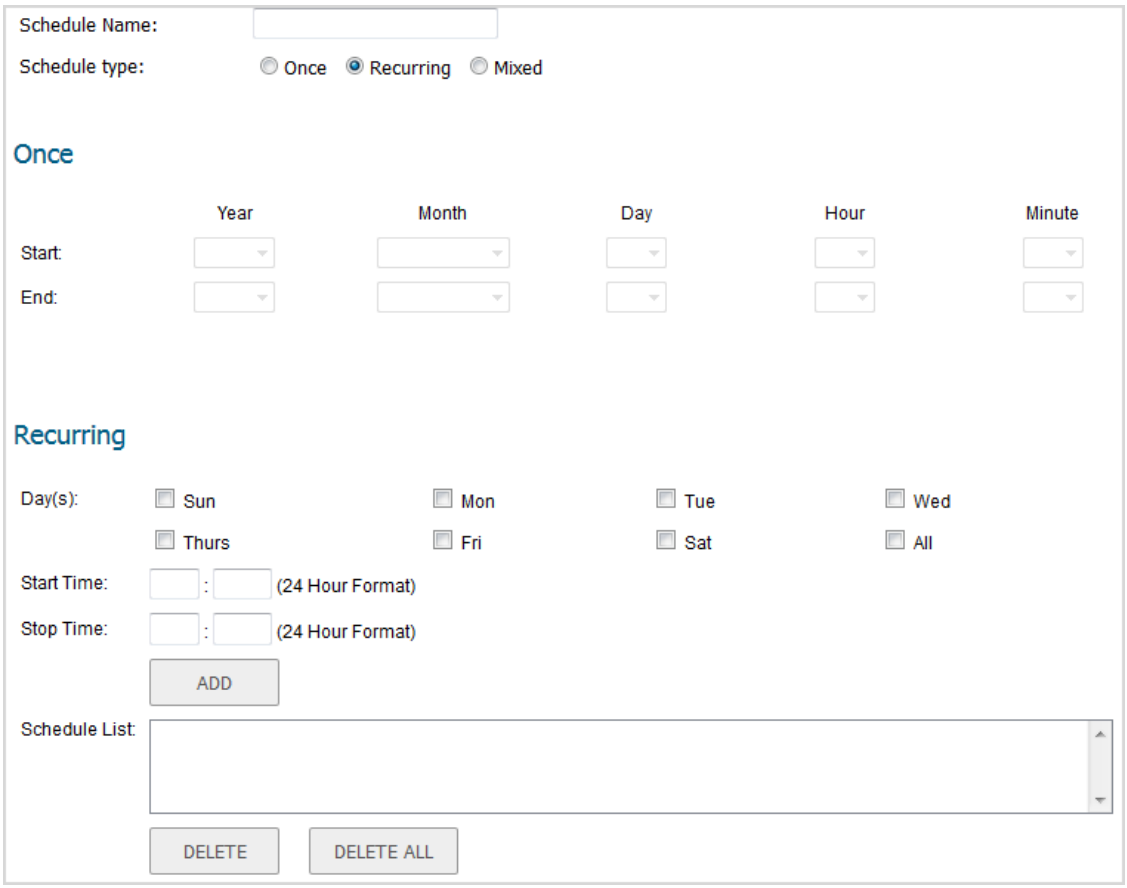
- Adding a Custom Schedule on page 78
- Modifying Schedules on page 79
- Deleting Custom Schedules on page 80

Adding a Custom Schedule

To create custom schedules:

1 Navigate to **MANAGE | System Setup > Appliance > Schedules**.

2 Click **ADD**. The **Add Schedule** dialog displays.



3 Enter a descriptive name for the schedule in the **Schedule Name** field.

4 Choose one of the following radio buttons for **Schedule type**:

Once	For a one-time schedule between the configured **Start** and **End** times and dates. When selected, the fields under **Once** become available, and the fields under **Recurring** become dimmed.
Recurring (default)	For a schedule that occurs repeatedly during the same configured hours and days of the week, with no start or end date. When selected, the fields under **Recurring** become available, and the fields under **Once** become dimmed.
Mixed	For a schedule that occurs repeatedly during the same configured hours and days of the week, between the configured start and end dates. When selected, all fields on the page become active.

IMPORTANT: Time must be in 24-hour format, for example, 17:00 for 5 p.m.

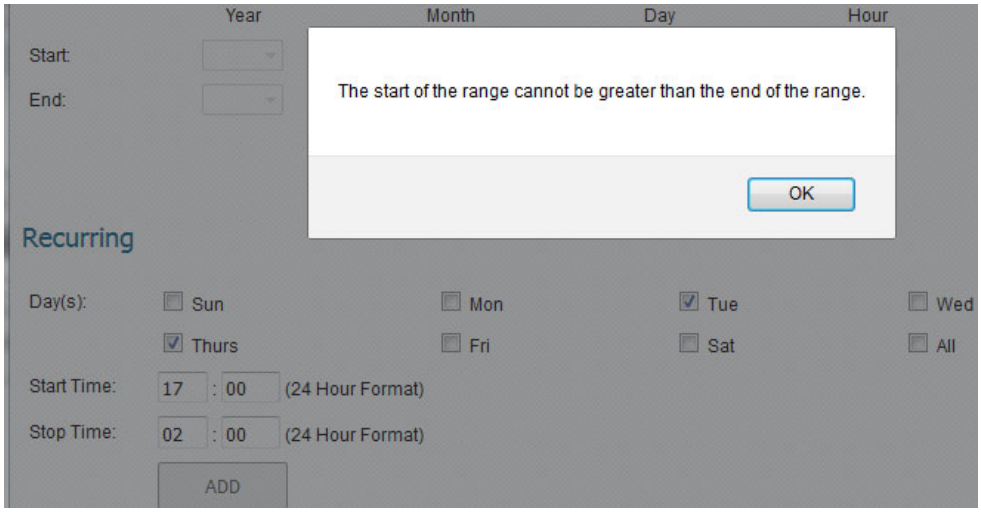
5 If the fields under **Once** are available, configure the:

- Starting date and time by selecting the **Year**, **Month**, **Date**, **Hour**, and **Minute** from the drop-down menus in the **Start** row. The hour is represented in 24-hour format.

- Ending date and time by selecting the **Year**, **Month**, **Date**, **Hour**, and **Minute** from the drop-down menus in the **End** row. The hour is represented in 24-hour format.

6 If the fields under **Recurring** are available:

- Select the checkboxes for the days of the week to apply to the schedule or select **All**.
- Enter the time of day for the schedule to begin in the **Start Time** field.
- Enter the time of day for the schedule to stop in the **Stop Time** field.



7 Click **ADD** to add the schedule to the **Schedule List**.

8 To delete:

- An existing schedule from the **Schedule List**:
  1) Select the schedule.
  2) Click **DELETE**.
- All existing schedules, click **DELETE ALL**.

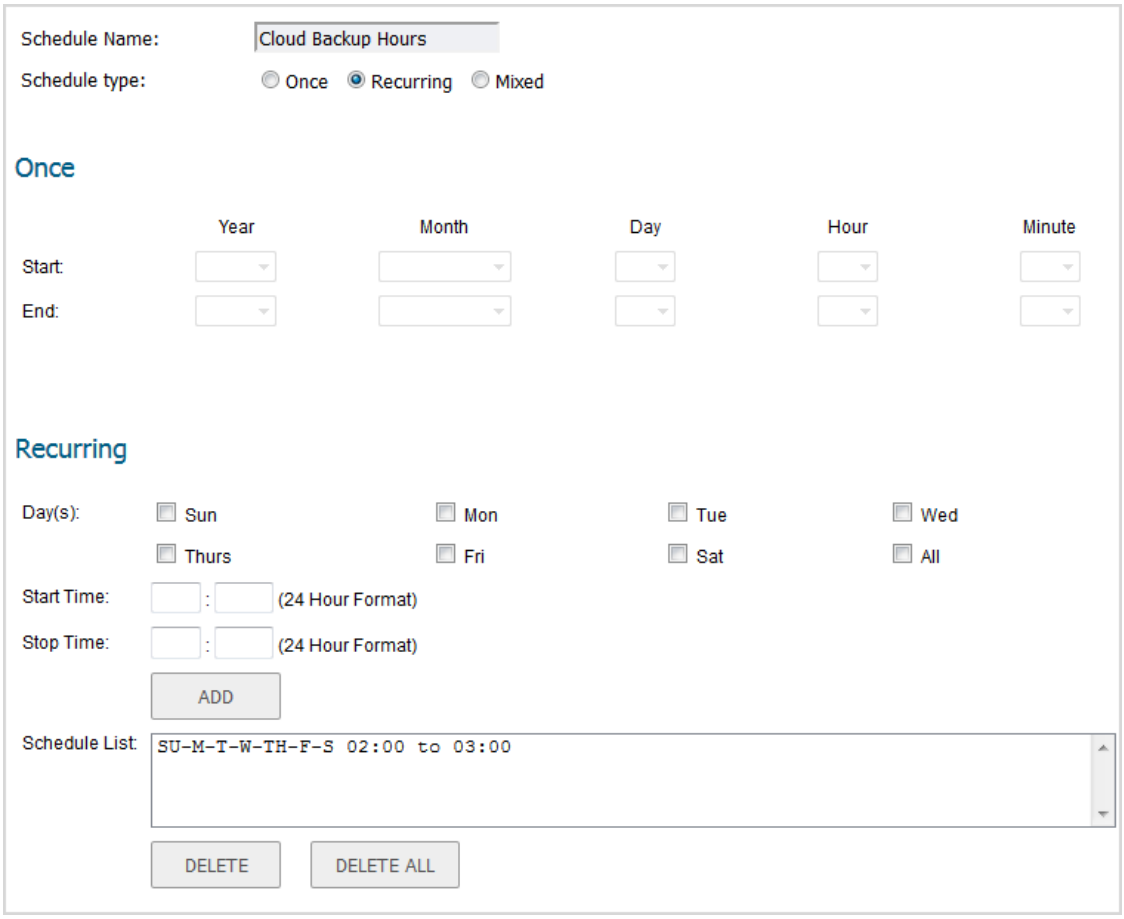
9 Click **OK**. The **Schedules** table is updated.

Modifying Schedules

To modify both default and custom schedules:

1 Navigate to **MANAGE | System Setup > Appliance > Schedules**.

2 Click the **Edit** icon for the schedule to be modified. The **Edit Schedule** dialog displays.



3 You can change any component of the schedule, such as time(s), type, and/or days, except the name of default schedules cannot be changed and the field is dimmed. To make changes, follow the procedure in Adding a Custom Schedule on page 78.

4 Click **OK**.

Deleting Custom Schedules

You can delete custom schedules, but you cannot delete default schedules.

Deleting Individual Schedules

To delete individual schedule objects that you created:

1 Navigate to **MANAGE | System Setup > Appliance > Schedules**.

2 In the **Schedules** table, to delete:

- A custom schedule, click its **Delete** icon.
- Multiple custom schedules:
    1) Select the checkboxes next to the custom schedules to delete. **DELETE** becomes available.
    2) Click **DELETE**.

Deleting All Schedules

To delete all schedule objects you created:

1. Navigate to **MANAGE | System Setup > Appliance > Schedules**.
2. In the **Schedules** table, select the checkbox next to the **Name** column header to select all custom schedules. **DELETE** becomes available.
3. Click **DELETE**.

Part 3

System Setup | User Management

- About Managing Users
- Configuring Settings for Managing Users
- Managing Authentication Partitions
- Configuring Local Users and Groups
- Managing Guest Services
- Managing Guest Accounts

7

About Managing Users

- About User Management on page 83
  - Using Local Users and Groups for Authentication on page 84
  - Using RADIUS for Authentication on page 87
  - Using LDAP/Active Directory/eDirectory Authentication on page 87
  - Using TACACS+ on page 91
  - About Single Sign-On on page 92
  - Installing the Single Sign-On Agent and/or Terminal Services Agent on page 102
  - About Multiple Administrator Support on page 112
  - Configuring Multiple Administrator Support on page 114

About User Management

NOTE: This topic provides an overview of the management capabilities of your SonicWall Security Appliance.

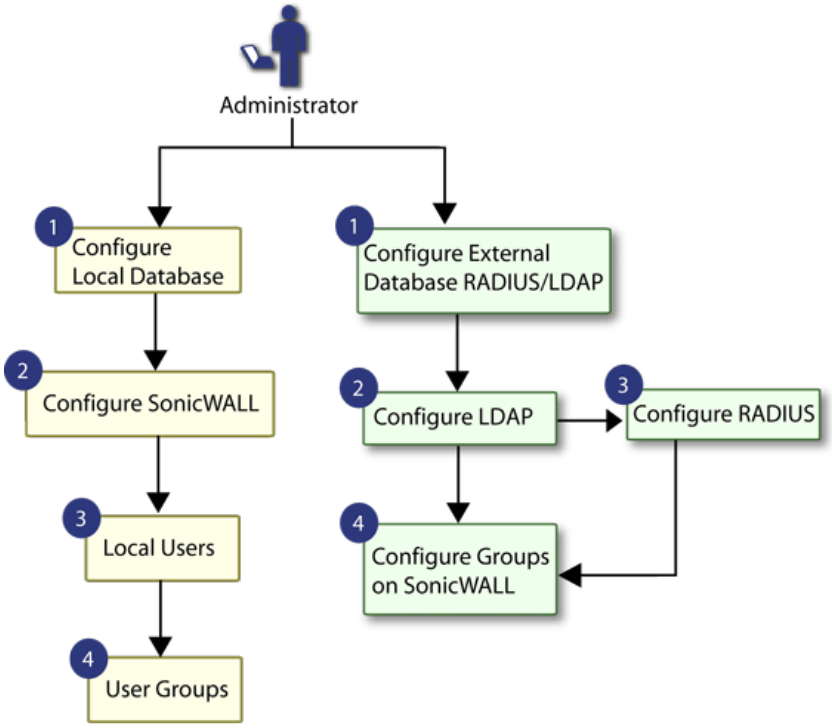
For detailed information and procedures about	See these topics
Setting up user authentication, web login, session management, RADIUS accounting, and policies	Configuring Settings for Managing Users on page 116.
Creating partitions for user authentication in environments with multiple non-interconnected domains	Managing Authentication Partitions on page 190
Creating and managing local users and local groups	Configuring Local Users and Groups on page 221.
Setting up guest services and accounts	Managing Guest Services on page 242 and Managing Guest Accounts on page 247

The SonicWall Security Appliance (firewall) provides a mechanism for managing locally and remotely authenticated users. User-level authentication gives users access to the LAN from remote locations on the Internet as well as a means to enforce or bypass content filtering policies for LAN users attempting to access the Internet. You can also permit only authenticated users to access VPN tunnels and send data across the encrypted connection.

The firewall authenticates all users as soon as they attempt to access network resources in a different zone (such as WAN, VPN, WLAN), which causes the network traffic to pass through the firewall. Users who log into a computer on the LAN, but perform only local tasks are not authenticated by the firewall. User-level authentication can be performed using a local user database, LDAP, RADIUS, or a combination of a local database with either LDAP or RADIUS. For networks with a large numbers of users, user authentication using LDAP or RADIUS servers can be more efficient.

SonicOS also provides Single Sign-On (SSO) capability. SSO can be used in conjunction with LDAP. See User management topology.

User management topology



Topics:

- Using Local Users and Groups for Authentication on page 84
- Using RADIUS for Authentication on page 87
- Using LDAP/Active Directory/eDirectory Authentication on page 87
- About Single Sign-On on page 92
- Installing the Single Sign-On Agent and/or Terminal Services Agent on page 102
- About Multiple Administrator Support on page 112
- Configuring Multiple Administrator Support on page 114

Using Local Users and Groups for Authentication

Topics:

- About User Databases on page 84
- About User Groups on page 85

About User Databases

The firewall provides a local database for storing user and group information. You can configure the firewall to use this local database to authenticate users and control their access to the network. The local database is a

good choice over LDAP or RADIUS when the number of users accessing the network is relatively small. Creating entries for dozens of users and groups takes time, although when the entries are in place they are not difficult to maintain.

The number of users supported by the local database on the firewall varies by platform is shown in Maximum Number of Supported Users by Platform. The maximum overall user limit is equal to the maximum number of SSO users and the maximum number of native users is equal to the maximum number of SSO users. The maximum web users is the maximum combined user logins from the web and the GVC, SSL-VP, and L2TP clients.

Maximum Number of Supported Users by Platform

Platform	SSO Users	Web Users	Web Server Threads	Platform	SSO Users	Web Users	Web Server Threads
NS*a* 6650	70,000	5,000	20	NS*a* 9650	100,000	5,000	30
NS*a* 5650	60,000	3,000	16	NS*a* 9450	90,000	5,000	30
NS*a* 4650	50,000	2,000	10	NS*a* 9250	80,000	5,000	20
NS*a* 3650	40,000	1,500	8				
NS*a* 2650	30,000	1,000	8	Super Massive 9600	100,000	5,000	30
				Super Massive 9400	90,000	5,000	30
NSA 6600	70,000	5,000	20	Super Massive 9200	80,000	5,000	20
NSA 5600	60,000	3,000	16				
NSA 4600	50,000	2,000	10	TZ600/TZ600P	500	500	8
NSA 3600	40,000	1,500	8	TZ500/TZ500W	500	500	8
NSA 2600	30,000	1,000	8	TZ400/TZ400W	500	150	8
				TZ350/TZ350W	500	150	8
				TZ300/TZ300W/TZ300P	500	150	8
				SOHO 250/250W	350	150	8
				SOHO W	250	150	8

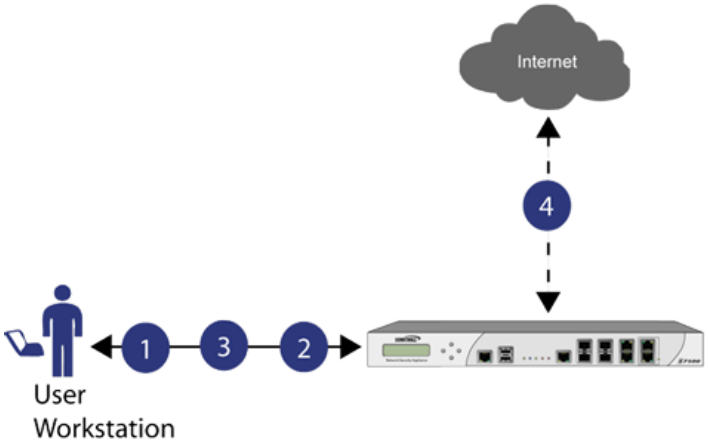
IMPORTANT: To achieve the maximum efficiency in handling these numbers, SonicWall recommends:

- For wireless users, use RADIUS Accounting as much as possible.
- Use SSO Agent version 4 or higher; do not use any SSO Agent older than version 3.6.10.
- Use the SSO Agent in DC logs mode with LogWatcher wherever possible.
- If NetAPI or WMI is needed to identify non-domain users, then do it in separate agents.
- Where possible, set exclusions to prevent anything that cannot be identified by SSO from triggering it.

About User Groups

To apply Content Filtering Service (CFS) policies to users, the users must be members of local groups and the CFS policies are then applied to the groups. To use CFS, you cannot use LDAP or RADIUS without combining that method with local authentication. When using the combined authentication method to use CFS policies, the local group names must be an exact match with the LDAP or RADIUS group names. When using the **LDAP + Local Users** authentication method, you can import the groups from the LDAP server into the local database on the firewall. This greatly simplifies the creation of matching groups, to which CFS policies can then be applied. See User management: Using local users and groups for authentication.

User management: Using local users and groups for authentication



1. User attempts to access the web.
2. SNWL requires authentication of the User: redirects workstation to authenticate.
3. User authenticates with credentials.
4. SNWL Local Database authorizes or denies access based on User privileges.

The SonicOS Management Interface provides a way to create local user and group accounts. You can add users and edit the configuration for any user, including settings for:

Group membership	Users can belong to one or more local groups. By default, all users belong to the groups **Everyone** and **Trusted Users**. You can remove these group memberships for a user and can add memberships in other groups.
VPN access	You can configure the networks that are accessible to a VPN client started by a user. When configuring VPN access settings, you can select from a list of networks. The networks are designated by their Address Group or Address Object names. **NOTE:** The VPN access configuration for users and groups affects the ability of remote clients using GVC, NetExtender, and SSL VPN Virtual Office bookmarks to access network resources. To allow GVC, NetExtender, or Virtual Office users to access a network resource, the network address objects or groups must be added to the “allow” list on the VPN Access tab.

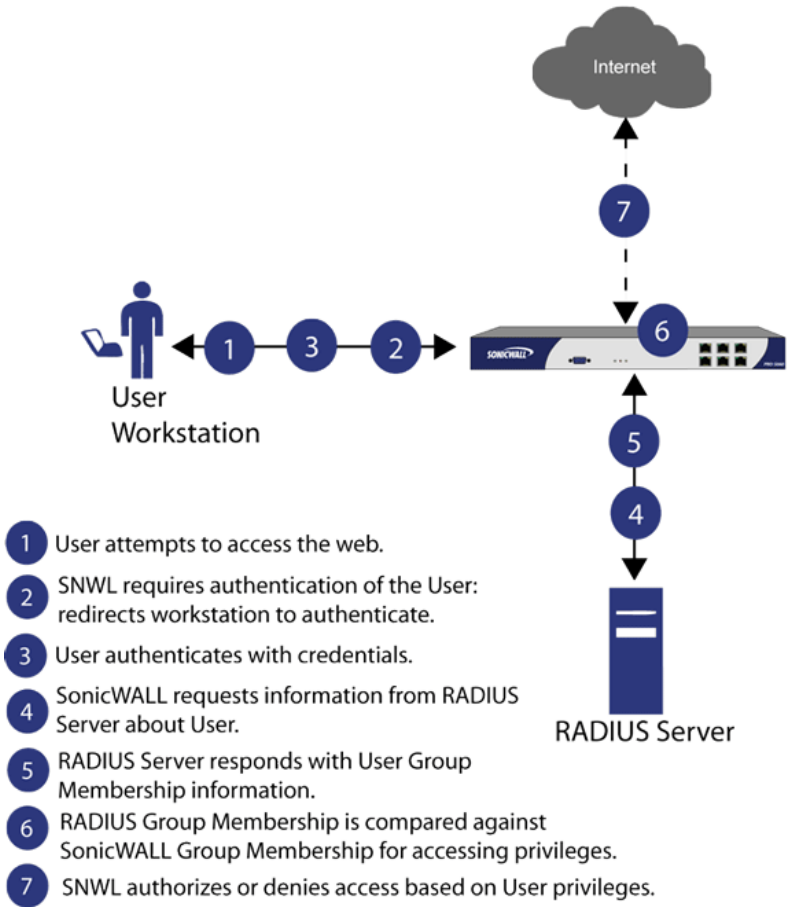
You can also add or edit local groups. Here are the configurable settings for groups:

Group settings	For administrator groups, you can configure SonicOS to allow login to the Management Interface without activating the login status popup window.
Group members	Groups have members that can be local users or other local groups.
VPN access	VPN access for groups is configured in the same way as VPN access for users. You can configure the networks that are accessible to a VPN client started by a member of this group. When configuring VPN access settings, you can select from a list of networks. The networks are designated by their **Address Group** or **Address Object** names.
CFS policy	You can apply a content filtering (CFS) policy to group members. The CFS policy setting is only available if the firewall is currently licensed for Premium Content Filtering Service.

Using RADIUS for Authentication

Remote Authentication Dial In User Service (RADIUS) is a networking protocol that provides centralized authentication, authorization, and accounting for SonicWall Security Appliances and SonicWave appliances to authenticate users attempting to access the network. The RADIUS server contains a database with user information and checks a user's credentials using authentication schemes such as Password Authentication Protocol (PAP), Challenge-handshake authentication protocol (CHAP), Microsoft CHAP (MSCHAP), or MSCHAPv2. See User management: Using RADIUS for authentication. For SonicWave appliances providing RADIUS authentication, see *SonicOS 6.5 Connectivity*.

User management: Using RADIUS for authentication



While RADIUS is very different from LDAP, primarily providing secure authentication, it can also provide numerous attributes for each entry, including a number of different ones that can be used to pass back user group memberships. RADIUS can store information for thousands of users, and is a good choice for user authentication purposes when many users need access to the network.

Using LDAP/Active Directory/eDirectory Authentication

Lightweight Directory Access Protocol (LDAP) defines a directory services structure for storing and managing information about elements in your network, such as user accounts, user groups, hosts, and servers. Several different standards exist that use LDAP to manage user account, group, and permissions. Some are proprietary systems like Microsoft Active Directory (AD), which you can manage using LDAP, or Novell eDirectory, which

provides an LDAP API for managing the user repository information. Some are open standards like SAMBA, which are implementations of the LDAP standards.

In addition to RADIUS and the local user database, SonicOS supports LDAP for user authentication, with support for numerous schemas including Microsoft Active Directory, Novell eDirectory directory services, and a fully configurable user-defined option that should allow SonicOS to interact with any schema.

Microsoft Active Directory also works with SonicWall Single Sign-On and the SonicWall SSO Agent. For more information, see About Single Sign-On on page 92.

Topics:

- LDAP Terms on page 88
- LDAP Directory Services Supported in SonicOS on page 89
- LDAP User Group Mirroring on page 89

LDAP Terms

These terms are useful when working with LDAP and its variants.

Active Directory (AD)	Microsoft directory service, commonly used with Windows-based networking. Microsoft Active Directory is compatible with LDAP.
Attribute	Data item stored in an object in an LDAP directory. Objects can have required attributes or allowed attributes. For example, the `dc` attribute is a required attribute of the `dcObject` (domain component) object.
cn	Common name attribute, a required component of many object classes throughout LDAP.
dc	Domain component attribute, commonly found at the root of a distinguished name and is commonly a required attribute.
dn	Distinguished name, which is a globally unique name for a user or other object. It is made up of a number of components, usually starting with a common name (`cn`) component and ending with a domain specified as two or more domain components (`dc`). For example, `cn=john`, `cn=users`, `dc=domain`, `dc=com`.
eDirectory	Novell directory service, used for Novell NetWare-based networking. Novell eDirectory has an LDAP gateway that can be used for management.
Entry	Data stored in the LDAP directory. Entries are stored in attribute/value (or name/value) pairs, where the attributes are defined by object classes. A sample entry would be `cn=john` where `cn` (common name) is the attribute and `john` is the value.
Object	In LDAP terminology, the entries in a directory are referred to as objects. For the purposes of the SonicOS implementation of the LDAP client, the critical objects are User and Group objects. Different implementations of LDAP can refer to these object classes in different fashions, for example, Active Directory refers to the user object as `user` and the group object as `group`, while RFC2798 refers to the user object as `inetOrgPerson` and the group object as `groupOfNames`.
Object class	Defines the type of entries that an LDAP directory might contain. A sample object class, as used by AD, would be `user` or `group`. Microsoft Active Directory's Classes can be browsed at http://msdn.microsoft.com/library/.
ou	Organizational unit attribute, a required component of most LDAP schema implementations.

Schema	Set of rules or the structure that defines the types of data that can be stored in a directory and how that data can be stored. Data is stored in the form of entries.
TLS	Transport Layer Security, the IETF standardized version of SSL (Secure Sockets Layer). TLS 1.1 and 1.2 are supported.

LDAP Directory Services Supported in SonicOS

To integrate with the most common directory services used in company networks, SonicOS supports integration with these LDAP schemas:

Microsoft Active Directory	Samba SMB
RFC2798 InetOrgPerson	Novell eDirectory
RFC2307 Network Information Service	User-defined schemas

SonicOS provides support for directory servers running these protocols:

LDAPv3 (RFC2251-2256, RFC3377)	LDAPv2 (RFC3494)
LDAPv3 over TLS (RFC2830)	LDAP Referrals (RFC2251)
LDAPv3 with STARTTLS (RFC2830)	

LDAP User Group Mirroring

LDAP User Group Mirroring provides automatic duplication of LDAP User Group configurations from an LDAP server to a SonicWall Security Appliance. You can manage LDAP User Groups exclusively on the LDAP server and do not need to manually duplicate configurations on the firewall. User group configurations are periodically read from the LDAP server and copied to the firewall.

LDAP User Group names that are copied to the firewall include the domain name in the format, `name@domain.com`. This ensures that user group names from various domains are unique.

These features and restrictions apply to mirrored LDAP User Groups:

- You can delete LDAP User Groups only on the LDAP server. They cannot delete the mirrored LDAP User Groups on the SonicWall Security Appliance. When a user group is deleted on the LDAP server, its mirrored group on the firewall is also deleted automatically.
- You can edit LDAP User Group names (and their comment fields) only on the LDAP server. They cannot edit the mirrored LDAP User Group name or its comment field on the firewall. The comment field displays `Mirrored from LDAP` on the firewall.
- You can add users as members to an LDAP User Group on the SonicWall Security Appliance.
- You cannot add groups to other groups on the SonicWall Security Appliance. Default user groups can only be configured on the LDAP server.
- You can configure things such as VPNs, SSL VPNs, CFS policies, and ISP policies for LDAP User Groups on the SonicWall Security Appliance (for more information about policies, see *SonicOS 6.5 Policies.*

  NOTE: LDAP User Groups are not deleted if they are configured in any Access Rules, App Control Rules, or other policies.

- When you disable LDAP User Group Mirroring, the mirrored user groups on the SonicWall Security Appliance are not deleted. They are changed so you can delete them manually. Local mirrored user groups can be re-enabled if they have not been deleted manually.
- When the system creates a mirrored group on the SonicWall Security Appliance, and the name of the mirrored group matches the name of an already existing, user-created (non-mirrored) local group, the

local group is not replaced. The local group memberships are updated to reflect the group settings that are configured on the LDAP server.

- If the system finds a user group on the LDAP server with a name that is the same as one of the default user groups on the SonicWall Security Appliance, no mirrored user group is created on the SonicWall Security Appliance. The memberships in the default user group are updated to reflect the group settings that are configured on the LDAP server.
- For groups created before SonicOS 6.2, if a local user group exists on the SonicWall Security Appliance with a simple name only (no domain) and that name matches the name of a user group on the LDAP server (which includes a domain), a new local user group is created on the SonicWall Security Appliance and is given the same domain as the corresponding user group on the LDAP server. The original local user group is retained with no domain. Users of the original group are given memberships in the LDAP group, the new local mirrored group, and the original local group (with no domain).

Integrating LDAP into the SonicWall Security Appliance

Integrating your firewall with an LDAP directory service requires configuring your LDAP server for certificate management, installing the correct certificate on your firewall, and configuring the firewall to use the information from the LDAP Server. For an introduction to LDAP, see Using LDAP/Active Directory/eDirectory Authentication on page 87.

Topics:

- Preparing Your LDAP Server for Integration on page 90
- Configuring the CA on the Active Directory Server on page 90

Preparing Your LDAP Server for Integration

Before beginning your LDAP configuration, you should prepare your LDAP server and your SonicWall for LDAP over TLS support. This requires:

- Installing a server certificate on your LDAP server.
- Installing a CA (Certificate Authority) certificate for the issuing CA on your firewall.

The following procedures describe how to perform these tasks in an Active Directory environment.

Configuring the CA on the Active Directory Server

To configure the CA on the Active Directory server:

TIP: Skip the first five steps if Certificate Services are already installed.

1. Navigate to **Start > Settings > Control Panel > Add/Remove Programs**
2. Select **Add/Remove Windows Components**
3. Select **Certificate Services**
4. Select **Enterprise Root CA** when prompted.
5. Enter the requested information. For information about certificates on Windows systems, see http://support.microsoft.com/kb/931125.
6. Launch the **Domain Security Policy** application: Navigate to **Start > Run** and run the command: **dompol.msc**.
7. Open **Security Settings > Public Key Policies**.
8. Right click **Automatic Certificate Request Settings**.

9 Select **New > Automatic Certificate Request**.

10 Step through the wizard, and select **Domain Controller** from the list.

Exporting the CA Certificate from the Active Directory Server

To export the CA certificate from the AD server:

1 Launch the **Certification Authority** application: **Start > Run > certsrv.msc**.

2 Right click on the CA you created, and select **properties**.

3 On the **General** tab, click **View Certificate**.

4 On the **Details** tab, select **Copy to File**.

5 Step through the wizard, and select the **Base-64 Encoded X.509 (.cer)** format.

6 Specify a path and filename to which to save the certificate.

Importing the CA Certificate in to SonicOS

To import the CA certificate in to SonicOS:

1 Browse to **System > CA Certificates**.

2 Select **Add new CA certificate**. Browse to and select the certificate file you just exported.

3 Click **Import certificate**.

LDAP Group Membership by Organizational Unit

The LDAP Group Membership by Organizational Unit feature provides the ability to set LDAP rules and policies for users located in certain Organizational Units (OUs) on the LDAP server.

When a user logs in, if user groups are set to grant memberships by LDAP location, the user is made a member of any groups that match its LDAP location.

You can set any local group, including default local groups (except for the **Everyone** group and the **Trusted Users** group) as a group with members that are set by their location in the LDAP directory tree.

When a user is a member of any local groups that are configured for LDAP location:

- The location of those local groups in the LDAP tree is learned.
- The location of the user's local groups is checked against all other local groups. If any other groups have the same LDAP location as that of the user's membership groups, the user is automatically set as a member of those groups for that login session.

When a user attempts to log in, whether with success or failure, the user's distinguished name is logged in the event log. This helps with troubleshooting if a user fails to get memberships to the expected groups.

Using TACACS+

SonicOS supports TACACS+ (Terminal Access Controller Access-Control System latest generation) for user authentication. The main characteristics of TACACS+ are:

- Provides separate authentication, authorization and accounting (AAA) services.
- Uses TCP for its transport.
- Entire TACACS+ body might be protected by the encryption.

About Single Sign-On

Topics:

- What Is Single Sign-On? on page 92
- Benefits of SonicWall SSO on page 92
- Platforms and Supported Standards on page 93
- How Does Single Sign-On Work? on page 94
- How Does SSO Agent Work? on page 96
- How Does Terminal Services Agent Work? on page 97
- How Does Browser NTLM Authentication Work? on page 99
- How Does RADIUS Accounting for Single-Sign-On Work? on page 100

What Is Single Sign-On?

Single Sign-On (SSO) is a transparent user-authentication mechanism that provides privileged access to multiple network resources with a single domain login to a workstation or through a Windows Terminal Services or Citrix server.

SonicWall Security Appliances provide SSO functionality using the Single Sign-On Agent (SSO Agent) and SonicWall Terminal Services Agent (TSA) to identify user activity. The SSO Agent identifies users based on workstation IP address. The TSA identifies users through a combination of server IP address, user name, and domain.

SonicWall SSO is also available for Mac and Linux users when used with Samba. Additionally, browser NTLM authentication allows SonicWall SSO to authenticate users who send HTTP traffic without involving the SSO Agent or Samba.

SonicWall SSO is configured in the **MANAGE | System Setup > Users > Settings** page of the SonicOS management interface. SSO is separate from the **Authentication method for login** settings, which can be used at the same time for authentication of VPN/L2TP client users or administrative users.

Based on data from SonicWall SSO Agent or TSA, the Security Appliance queries LDAP or the local database to determine group membership. Memberships are optionally checked by firewall policies to control who is given access, and can be used in selecting policies for Content Filtering and Application Control to control what they are allowed to access. User names learned UNIX SSO are reported in logs of traffic and events from the users, and in AppFlow Monitoring.

The configured inactivity timer applies with SSO but the session limit does not, though users who are logged out are automatically and transparently logged back in when they send further traffic.

Users logged into a workstation or Terminal Services/Citrix server directly, but not logged into the domain, are not authenticated unless they send HTTP traffic and browser NTML authentication is enabled (although they can optionally be authenticated for limited access). For users who are not authenticated by SonicWall SSO, a message displays indicating that a manual login to the Security Appliance is required for further authentication.

Users that are identified but lack the group memberships required by the configured policy rules are redirected to the Access Barred page.

Benefits of SonicWall SSO

SonicWall SSO is a reliable and time-saving feature that utilizes a single login to provide access to multiple network resources based on administrator-configured group memberships and policy matching. SonicWall SSO is transparent to end users and requires minimal administrator configuration.

By automatically determining when users have logged in or out based on workstation IP address traffic, or, for Terminal Services or Citrix, traffic from a particular user at the server IP address, SonicWall SSO is secure and hands-free. SSO authentication is designed to operate with any external agent that can return the identity of a user at a workstation or Terminal Services/Citrix server IP address using a SonicWall Directory Connector-compatible protocol.

SonicWall SSO works for any service on the firewall that uses user-level authentication, including Content Filtering Service (CFS), Access Rules, group membership and inheritance, and security services (IPS, GAV, and Anti-Spyware) inclusion/exclusion lists.

SonicWall SSO Agent can be installed on any Windows server on the LAN, and TSA can be installed on any terminal server. Other benefits of SonicWall SSO include:

Ease of use	Users only need to sign in once to gain automatic access to multiple resources.
Improved user experience	Windows domain credentials can be used to authenticate a user for any traffic type without logging into the appliance using a Web browser.
Transparency to users	Users are not required to re-enter user name and password for authentication.
Secure communication	Shared key encryption for data transmission protection.
Multiple SSO Agents	Up to 8 agents are supported to provide capacity for large installations
Multiple TSAs	Multiple terminal services agents (one per terminal server) are supported. The number depends on the model of the SonicWall Security Appliance and ranges from 8 to 512.
Login mechanism	Works with any protocol, not just HTTP.
Browser NTLM authentication	SonicWall SSO can authenticate users sending HTTP traffic without using the SSO Agent.
Mac and Linux support	With Samba 3.5 and higher, SonicWall SSO is supported for Mac and Linux users.
Per-zone enforcement	SonicWall SSO can be triggered for traffic from any zone even when not automatically initiated by firewall access rules or security services policies, providing user identification in event logging or AppFlow Monitoring.

Platforms and Supported Standards

The SSO Agent is compatible with all versions of SonicOS that support SonicWall SSO. SonicWall TSA also is supported.

The SSO feature supports LDAP and local database protocols. SonicWall SSO supports SonicWall Directory Connector. For all features of SonicWall SSO to work properly, SonicOS should be used with Directory Connector 3.1.7 or higher.

To use SonicWall SSO with Windows Terminal Services or Citrix, SonicOS 6.0 or higher is required, and SonicWall TSA must be installed on the server.

To use SonicWall SSO with browser NTLM authentication, SonicOS 6.0 or higher is required. The SSO Agent is not required for browser NTLM authentication.

Except when using only browser NTLM authentication, using SonicWall SSO requires that the SSO Agent be installed on a server within your Windows domain that can reach clients and can be reached from the appliance, either directly or through a VPN path, and/or TSA be installed on any terminal servers in the domain.

The following requirements must be met to run the SSO Agent:

- UDP port 2258 (by default) must be open; the firewall uses UDP port 2258 by default to communicate with SonicWall SSO Agent; if a custom port is configured instead of 2258, then this requirement applies to the custom port
- Windows Server, with latest service pack
- .NET Framework 2.0
- Net API or WMI

NOTE: Mac and Linux PCs do not support the Windows networking requests that are used by the SSO Agent, and hence require Samba 3.5 or newer to work with SonicWall SSO. Without Samba, Mac and Linux users can still get access, but you need to log in to do so. You can be redirected to the login prompt if policy rules are set to require authentication. For more information, see Accommodating Mac and Linux Users on page 108.

To run the TSA, the following requirements must be met:

- UDP port 2259 (by default) must be open on all terminal servers on which TSA is installed; the firewall uses UDP port 2259 by default to communicate with SonicWall TSA; if a custom port is configured instead of 2259, then this requirement applies to the custom port
- Windows Server, with latest service pack
- Windows Terminal Services or Citrix installed on the Windows Terminal Server system(s)

How Does Single Sign-On Work?

SonicWall SSO requires minimal administrator configuration and is transparent to the user.

SSO is triggered in these situations:

- If firewall access rules requiring user authentication apply to traffic that is not incoming from the WAN zone
- When no user groups are specified in access rules, but any of the following conditions exist, SSO is triggered for all traffic on the zone and not just for traffic subject to these conditions:
    - CFS is enabled on the zone and multiple CFS policies are set
    - IPS is enabled on the zone and there are IPS policies that require authentication
    - Anti-Spyware is enabled on the zone and there are Anti-Spyware policies that require authentication
    - Application Control policies that require authentication apply to the source zone
    - Per-zone enforcement of SSO is set for the zone

The SSO user table is also used for user and group identification needed by security services, including Content Filtering, Intrusion Prevention, Anti-Spyware, and Application Control.

SonicWall SSO Authentication Using the SSO Agent

For users on individual Windows workstations, the SSO Agent (on the SSO workstation) handles the authentication requests from the firewall. There are six steps involved in SonicWall SSO authentication using the SSO Agent.

The SSO authentication process is initiated when user traffic passes through a firewall. For example, when a user accesses the Internet. The sent packets are temporarily blocked and saved while the firewall sends a “User Name” request and workstation IP address to the authorization agent running the SSO Agent (the SSO workstation).

The authorization agent running the SSO Agent provides the firewall with the user name currently logged into the workstation. A User IP Table entry is created for the logged in user, similarly to RADIUS and LDAP.

SonicWall SSO Authentication Using the Terminal Services Agent

For users logged in from a Terminal Services or Citrix server, the TSA takes the place of the SSO Agent in the authentication process. The process is different in several ways:

- The TSA runs on the same server that the user is logged into, and includes the user name and domain along with the server IP address in the initial notification to the firewall.
- Users are identified by a user number as well as the IP address (for non-Terminal Services users, there is only one user at any IP address and so no user number is used). A non-zero user number is displayed in the SonicOS Management Interface using the format *x.x.x.x* `user` *n*, where *x.x.x.x* is the server IP address and *n* is the user number.
- The TSA sends a close notification to SonicOS when the user logs out, so no polling occurs.

After a user has been identified, the Security Appliance queries LDAP or a local database (based on administrator configuration) to find user group memberships, match the memberships against policy, and grant or restrict access to the user accordingly. Upon successful completion of the login sequence, the saved packets are sent on. If packets are received from the same source address before the sequence is completed, only the most recent packet is saved.

User names are returned from the authorization agent running the SSO Agent in the format `<domain>/<user-name>`. For locally configured user groups, the user name can be configured to be:

- The full name returned from the authorization agent running the SSO Agent (configuring the names in the firewall local user database to match).
- A simple user name with the domain component stripped off (default).

For the LDAP protocol, the `<domain>/<user-name>` format is converted to an LDAP distinguished name by creating an LDAP search for an object of class `domain` with a `dc` (domain component) attribute that matches the domain name. If one is found, then its distinguished name is used as the directory sub-tree to search for the user's object. For example, if the user name is returned as `SV/bob`, then a search for an object with `objectClass=domain` and `dc=SV` i performed. If that returns an object with distinguished name `dc=sv,dc=us,dc=sonicwall,dc=com`, then a search under that directory sub-tree is created for (in the Active Directory case) an object with `objectClass=user` and `sAMAccountName=bob`. If no domain object is found, then the search for the user object is made from the top of the directory tree.

When a domain object has been found, the information is saved to avoid searching for the same object. If an attempt to locate a user in a saved domain fails, the saved domain information is deleted and another search for the domain object is made.

User logout is handled slightly differently by SonicWall SSO using the SSO Agent as compared to SSO with TSA. The Security Appliance polls the authorization agent running the SSO Agent at a configurable rate to determine when a user has logged out. Upon user logout, the authentication agent running the SSO Agent sends a User Logged Out response to the firewall, confirming that the user has been logged out and terminating the SSO session. Rather than being polled by the Security Appliance, the TSA itself monitors the Terminal Services/Citrix server for logout events and notifies the Security Appliance as they occur, terminating the SSO session. For both agents, configurable inactivity timers can be set, and for the SSO Agent the user name request polling rate can be configured (set a short poll time for quick detection of logouts, or a longer polling time for less overhead on the system).

SonicWall SSO Authentication Using Browser NTLM Authentication

For users who are browsing using Mozilla-based browsers (including Internet Explorer, Firefox, Chrome, and Safari) the firewall supports identifying them through NTLM (NT LAN Manager) authentication. NTLM is part of a browser authentication suite known as "Integrated Windows Security" and is supported by all Mozilla-based browsers. NTLM allows a direct authentication request from the appliance to the browser without involving the SSO agent. NTLM is often used when a domain controller is not available, such as when the user is remotely authenticating over the Web.

NTLM Authentication is currently available for HTTP; it is not available for use with HTTPS traffic.

Browser NTLM authentication can be tried before or after the SSO agent attempts to acquire the user information. For example, if the SSO agent is tried first and fails to identify the user, then, if the traffic is HTTP, NTLM is tried.

To use this method with Linux or Mac clients as well as Windows clients, you can also enable SSO to probe the client for either **NetAPI** or **WMI**, depending on which is configured for the SSO Agent. This causes the firewall to probe for a response on the NetAPI/WMI port before requesting that the SSO Agent identify a user. If no response occurs, these devices fail SSO immediately. For a:

- Windows PC, the probe generally works (unless blocked by a personal firewall) and the SSO agent is used.
- Linux/Mac PC (assuming it is not set up to run Samba server), the probe fails, the SSO agent is bypassed, and NTLM authentication is used when HTTP traffic is sent.

NTLM cannot identify the user until they browse with HTTP, so any traffic sent before that is treated as unidentified. The default CFS policy is applied, and any rule requiring authenticated users does not allow the traffic to pass.

If NTLM is configured to be used before the SSO agent, then if HTTP traffic is received first, the user is authenticated with NTLM. If non-HTTP traffic is received first, the SSO agent is used for authentication.

How Does SSO Agent Work?

The SSO Agent can be installed on any workstation or server with a Windows domain that can communicate with clients and the firewall directly using the IP address or using a path, such as VPN. It is recommended, however, that the SSO Agent be installed on separate, standalone workstations or servers. For installation instructions for the SSO Agent, see Installing the SonicWall SSO Agent on page 103.

Multiple SSO agents are supported to accommodate large installations with thousands of users. You can configure up to eight SSO agents, each running on a dedicated, high-performance PC in your network.

NOTE: When using NetAPI or WMI, one SSO Agent can support up to approximately 2500 users, depending on the performance level of the hardware that it is running on, how it is configured on the firewall, and other network-dependent factors. Depending on similar factors, when configured to read from domain controller security logs, one SSO Agent can support a much larger number of users identified through that mechanism, potentially up to 50,000+ users

The SSO Agent only communicates with clients and the firewall. The SSO Agent uses a shared key for encryption of messages between the SSO Agent and the firewall.

NOTE: The shared key is generated in the SSO Agent and the key entered in the firewall during SSO configuration must match the SSO Agent-generated key exactly.

The firewall queries the SSO Agent over the default port 2258. The SSO Agent then communicates between the client and the firewall to determine the client's user ID. The SSO Agent is polled, at a rate that is configurable by the administrator, by the firewall to continually confirm a user's login status.

Logging

The SSO Agent sends log event messages to the Windows Event Log based on administrator-selected logging levels.

The firewall also logs SSO Agent-specific events in its event log:

NOTE: The **Notes** field of log messages specific to the SSO Agent contain the text `<domain/user-name>, authentication by SSO Agent`. For more information about log messages, see *SonicOS 6.5 Investigate*.

User login denied - not allowed by policy rule	User has been identified but does not belong to any user groups allowed by the policy blocking the user's traffic.
User login denied - not found locally	User has not been found locally and **Allow only users listed locally** is selected in the firewall.
User login denied - SSO Agent agent timeout	Attempts to contact the SSO Agent have timed out.
User login denied - SSO Agent configuration error	SSO Agent is not properly configured to allow access for this user.
User login denied - SSO Agent communication problem	Problem communicating with the workstation running the SSO Agent.
User login denied - SSO Agent agent name resolution failed	SSO Agent is unable to resolve the user name.
SSO Agent returned user name too long	User name is too long.
SSO Agent returned domain name too long	Domain name is too long.

How Does Terminal Services Agent Work?

The TSA can be installed on any Windows Server machine with Terminal Services or Citrix installed. The server must belong to a Windows domain that can communicate with the firewall directly using the IP address or using a path, such as VPN.

For installation instructions for the TSA, refer to Installing the SonicWall Terminal Services Agent on page 103.

Topics:

- Multiple TSA Support on page 97
- Encryption of TSA Messages and Use of Session IDs on page 98
- Connections to Local Subnets on page 98
- Non-Domain User Traffic from the Terminal Server on page 98
- Non-User Traffic from the Terminal Server on page 98

Multiple TSA Support

To accommodate large installations with thousands of users, firewalls are configurable for operation with multiple terminal services agents (one per terminal server). The number of agents supported depends on the model, as shown in Terminal Services Agents Supported Per Model.

Terminal Services Agents Supported Per Model

SonicWall Network Security Appliance	TS Agents Supported	SonicWall Network Security Appliance	TS Agents Supported	SonicWall Network Security Appliance	TS Agents Supported
NS*a* 9650	512	NS*a* 6650	256	TZ600/TZ600 P	4
NS*a* 9450	512	NS*a* 5650	128	TZ500/TZ500 W	4
NS*a* 9250	512	NS*a* 4650	64	TZ400/TZ400 W	4
NSA 9600	512	NS*a* 3650	16	TZ350/350 W	4
NSA 9400	512	NS*a* 2650	8	TZ300/TZ300 W/TZ300 P	4
NSA 9200	512	NSA 6600	256		
SM 9600	512	NSA 5600	128	SOHO 250/250W	4
SM 9400	512	NSA 4600	64	SOHO W	4

Terminal Services Agents Supported Per Model (Continued)

SonicWall Network Security Appliance	TS Agents Supported	SonicWall Network Security Appliance	TS Agents Supported	SonicWall Network Security Appliance	TS Agents Supported
SM 9200	512	NSA 3600	16		
		NSA 2600	8		

For all SonicWall Security Appliances, a maximum of 32 IP addresses is supported per terminal server, where the server might have multiple NICs (network interface controllers). There is no limit on users per terminal server.

Encryption of TSA Messages and Use of Session IDs

The TSA uses a shared key for encryption of messages between the TSA and the firewall when the user name and domain are contained in the message. The first open notification for a user is always encrypted, because the TSA includes the user name and domain.

NOTE: The shared key is created in the TSA, and the key entered in the firewall during SSO configuration must match the TSA key exactly.

The TSA includes a user session ID in all notifications rather than including the user name and domain every time. This is efficient, secure, and allows the TSA to re-synchronize with Terminal Services users after the agent restarts.

Connections to Local Subnets

The TSA dynamically learns network topology based on information returned from the appliance and, after being learned, it does not send notifications to the appliance for subsequent user connections that do not go through the appliance. As there is no mechanism for the TSA to “unlearn” these local destinations, the TSA should be restarted if a subnet is moved between interfaces on the appliance.

Non-Domain User Traffic from the Terminal Server

The firewall has the **Allow limited access for non-domain users** setting for optionally giving limited access to non-domain users (those logged into their local machine and not into the domain), and this works for terminal services users as it does for other SSO users.

If your network includes non-Windows devices or Windows computers with personal firewalls running, select **Probe user for** and choose the radio button for either **NetAPI** or **WMI** depending on which is configured for the SSO Agent. This causes the firewall to probe for a response on the NetAPI/WMI port before requesting that the SSO Agent identify a user. If no response occurs, these devices fail SSO immediately. Such devices do not respond to, or could block, the Windows networking messages used by the SSO Agent to identify a user.

Non-User Traffic from the Terminal Server

Non-user connections are opened from the Terminal Server for Windows updates and anti-virus updates. The TSA can identify a connection from a logged-in service as being a non-user connection, and indicates this in the notification to the appliance.

To control handling of these non-user connections, an **Allow Terminal Server non-user traffic to bypass user authentication in access rules** checkbox is available in the TSA configuration on the appliance. When selected, these connections are allowed. If this checkbox is not selected, then the services are treated as local users and can be given access by selecting the **Allow limited access for non-domain users** setting and creating user accounts on the appliance with the corresponding service names.

NOTE: Ping (ICMP) traffic from the TSA is recognized as non-user traffic, but not as system service traffic. Therefore, it is not allowed to bypass user authentication and is dropped after the Agent times out. To prevent ICMP traffic from being dropped, add an access rule in the **Policies | Rules > Access Rule** page to allow ICMP from the terminal server(s) without requiring user authentication. For further information about access rules, see *SonicOS 6.5 Policies*.

How Does Browser NTLM Authentication Work?

Topics:

- NTLM Authentication of Domain Users on page 99
- NTLM Authentication of Non-Domain Users on page 99
- Credentials for NTLM Authentication in the Browser on page 99

NTLM Authentication of Domain Users

For domain users, the NTLM response is authenticated through the MSCHAP mechanism in RADIUS. RADIUS must be enabled on the appliance. For more information about NTLM authentication, see Configuring Settings for Managing Users on page 116.

NTLM Authentication of Non-Domain Users

With NTLM, non-domain users could be users who are logged into their PC rather than into the domain, or could be users who were prompted to enter a user name and password and entered something other than their domain credentials. In both cases, NTLM allows for distinguishing these from domain users.

If the user name matches a local user account on the firewall, then the NTLM response is validated locally against the password of that account. If successful, the user is logged in and given privileges based on that account. User group memberships are set from the local account, not from LDAP, and (as the password has been validated locally) include membership of the Trusted Users group.

If the user name does not match a local user account, the user is not logged in. The **Allow limited access for non-domain users** option does not apply for users authenticated through NTLM.

Credentials for NTLM Authentication in the Browser

For NTLM authentication, the browser either uses the domain credentials (if the user is logged into the domain), thus providing full single-sign-on functionality, or prompts the user to enter a name and password for the website being accessed (the firewall in this case). Different factors affect the browser's ability to use the domain credentials when the user is logged into the domain. These factors depend on the type of browser being used:

Internet Explorer (9.0 or above)	Uses the user's domain credentials and authenticates transparently if the website that it is logging into the firewall (the SonicWall Security Appliance) is in the local intranet, according to the Security tab in its Internet Options. This requires adding the firewall to the list of websites in the Local Intranet zone in the Internet Options. This can be done through the domain's group policy in the Site to Zone Assignment List under Computer Configuration, Administrative Templates, Windows Components, Internet Explorer, Internet Control Panel, Security Page.
Google Chrome	Behaves the same as Internet Explorer, including requiring that the firewall be added to the list of websites in the Local Intranet zone in the Internet Options.
Firefox	Uses the user's domain credentials and authenticates transparently if the website that it is logging into the firewall is listed in the **network.automatic-ntlm-auth.trusted-uris** entry in its configuration (accessed by entering **about:config** in the Firefox address bar).
Safari	Although Safari does support NTLM, it does not currently support fully transparent log on using the user's domain credentials. **NOTE:** Safari does not operate on Windows platforms.
Browsers on Non-PC Platforms	Non-PC platforms, such as Linux and Mac, can access resources in a Windows domain through Samba, but do not have the concept of "logging the PC into the domain" as Windows PCs do. Hence, browsers on these platforms do not have access to the user's domain credentials and cannot use them for NTLM.

When a user is not logged into the domain or the browser cannot use their domain credentials, it prompts for a name and password to be entered, or uses cached credentials if the user has previously opted to have it save them.

In all cases, should authentication fail when using the user's domain credentials (which could be because the user does not have the privileges necessary to get access), then the browser prompts the user to enter a name and password. This allows the user to enter credentials different from the domain credentials to get access.

NOTE: When NTLM is enabled for Single Sign-On enforcement, an HTTP/HTTPS access rule with **Trusted Users** as **Users Allowed** must be added to the **LAN to WAN** rules in the **MANAGE | Policies > Rules > Access Rules** page (for more information, see *SonicOS 6.5 Policies*). This rule triggers an NTLM authentication request to the user. Without the access rule, other configurations, such as restrictive Content Filter policies, might block the user from Internet access and prevent the authentication request.

How Does RADIUS Accounting for Single-Sign-On Work?

RADIUS Accounting is specified by RFC 2866 as a mechanism for a network access server (NAS) to send user login session accounting messages to an accounting server. These messages are sent at user login and logoff. Optionally, they can also be sent periodically during the user's session.

When a customer uses an external or third-party network access appliance to perform user authentication (typically for remote or wireless access) and the appliance supports RADIUS accounting, a SonicWall appliance can act as the RADIUS Accounting Server, and can use RADIUS Accounting messages sent from the customer's network access server for single sign-on (SSO) in the network.

NOTE: A SonicWall SMA 1000 Series appliance running SMA 11.4 or higher can be configured as an external RADIUS Accounting client, with the SonicWall firewall as the RADIUS Accounting server.

When a remote user connects through a SonicWall SMA or third-party appliance, the SMA or third-party appliance sends an accounting message to the SonicWall appliance (configured as a RADIUS accounting server). The SonicWall appliance adds the user to its internal database of logged in users based on the information in the accounting message.

When the user logs out, the SonicWall SMA or third-party appliance sends another accounting message to the SonicWall Security Appliance, which then logs the user out.

NOTE: When a network access server (NAS) sends RADIUS accounting messages, it does not require the user to be authenticated by RADIUS. The NAS can send RADIUS accounting messages even when the third-party appliance is using LDAP, its local database, or any other mechanism to authenticate users.

RADIUS accounting messages are not encrypted. RADIUS accounting is inherently secure against spoofing because it uses a request authenticator and a shared secret. RADIUS accounting requires that a list of the network access servers (NASs), that can send RADIUS Accounting messages, be configured on the appliance. This configuration supplies the IP address and shared secret for each NAS.

Topics:

- RADIUS Accounting Messages on page 100
- SonicWall Compatibility with Third Party Network Appliances on page 101
- Proxy Forwarding on page 101
- Non-Domain Users on page 102
- IPv6 Considerations on page 102
- RADIUS Accounting Server Port on page 102

RADIUS Accounting Messages

RADIUS accounting uses two types of accounting messages:

- **Accounting-Request**
- **Accounting-Response**

An **Accounting-Request** can send one of three request types specified by the **Status-Type** attribute:

This request	Is sent
Start	When a user logs in.
Stop	When a user logs out.
Interim-Update	Periodically during a user login session.

Accounting messages follow the RADIUS standard specified by RFC 2866. Each message contains a list of attributes and an authenticator that is validated by a shared secret.

These SSO-relevant attributes are set in **Accounting-Requests**:

Status-Type	Type of accounting request (**Start**, **Stop**, or **Interim-Update**).
User-Name	User's login name. The format is not specified by the RFC and can be a simple login name or a string with various values such as login name, domain, or distinguished name (DN).
Framed-IP-Address	User's IP address. If NAT is used, this must be the user's internal IP address.
Calling-Station-Id	String representation of the user's IP address, used by some appliances such as SMA.
Proxy-State	Pass-though state used for forwarding requests to another RADIUS accounting server.

SonicWall Compatibility with Third Party Network Appliances

For SonicWall Security Appliances to be compatible with third-party network appliances for SSO through RADIUS Accounting, the third-party appliance must be able to:

- Support RADIUS Accounting.
- Send both **Start** and **Stop** messages. Sending **Interim-Update** messages is not required.
- Send the user's IP address in either the **Framed-IP-Address** or **Calling-Station-Id** attribute in both **Start** and **Stop** messages.

NOTE: In the case of a remote access server using NAT to translate a user's external public IP address, the attribute must provide the internal IP address that is used on the internal network, and it must be a unique IP address for the user. If both attributes are being used, the **Framed-IP-Address** attribute must use the internal IP address, and the **Calling-Station-Id** attribute should use the external IP address.

The user's login name should be sent in the **User-Name** attribute of **Start** messages and **Interim-Update** messages. The user's login name can also be sent in the **User-Name** attribute of **Stop** messages, but is not required. The **User-Name** attribute must contain the user's account name and might include the domain also, or it must contain the user's distinguished name (DN).

Proxy Forwarding

A SonicWall Security Appliance acting as a RADIUS accounting server can proxy-forward requests to up to four other RADIUS accounting servers for each network access server (NAS). Each RADIUS accounting server is separately configurable for each NAS.

To avoid the need to re-enter the configuration details for each NAS, SonicOS allows you to select the forwarding for each NAS from a list of configured servers.

The proxy forwarding configuration for each NAS client includes timeouts and retries. How to forward requests to two or more servers can be configured by selecting these options:

- **try the next server on a timeout**

- **forward each request to all the servers**

Non-Domain Users

Users reported to a RADIUS accounting server are determined to be local (non-domain) users in these cases:

- The user name was sent without a domain, and it is not configured to look up domains for the server through LDAP.
- The user name was sent without a domain, and it is configured to look up domains for the server through LDAP, but the user name was not found.
- The user name was sent with a domain, but the domain was not found in the LDAP database.
- The user name was sent with a domain, but the user name was not found in the LDAP database.

A non-domain user authenticated by RADIUS accounting is subject to the same constraints as one authenticated by the other SSO mechanisms, and the following restrictions apply:

- The user is logged in only if **Allow limited access for non-domain users** is set.
- The user is not made a member of the Trusted Users group.

IPv6 Considerations

In RADIUS accounting, these attributes are used to contain the user's IPv6 address:

- Framed-Interface-Id / Framed-IPv6-Prefix
- Framed-IPv6-Address

Currently, all these IPv6 attributes are ignored.

Some devices pass the IPv6 address as text in the **Calling-Station-ID** attribute.

The **Calling-Station-ID** is also ignored if it does not contain a valid IPv4 address.

RADIUS accounting messages that contain an IPv6 address attribute and no IPv4 address attribute are forwarded to the proxy server. If no proxy server is configured, IPv6 attributes discarded.

RADIUS Accounting Server Port

RADIUS accounting normally uses UDP port:

1813	IANA-specified port. The SonicWall Security Appliance listens on port 1813 by default.
1646	An older, unofficial, standard port.

Other port numbers can be configured for the RADIUS accounting port, but the SonicWall Security Appliance can listen on only one port. So, if you are using multiple network access servers (NASs), they must all be configured to communicate on the same port number.

Installing the Single Sign-On Agent and/or Terminal Services Agent

Configuring SSO is a process that includes installing and configuring the SonicWall SSO Agent and/or the SonicWall Terminal Services Agent (TSA), and configuring a firewall running SonicOS to use the SSO Agent or TSA. For an introduction to SonicWall SSO, see About Single Sign-On on page 92.

Topics:

- Installing the SonicWall SSO Agent on page 103
- Installing the SonicWall Terminal Services Agent on page 103

Installing the SonicWall SSO Agent

The SonicWall SSO Agent is part of the SonicWall Directory Connector. The SonicWall SSO Agent must be installed on at least one, and up to eight, workstations or servers in the Windows domain that have access to the Active Directory server using VPN or IP. It is recommended that these workstations or servers be separate, standalone workstations or servers. The SonicWall SSO Agent must have access to your firewall.

To install the SonicWall SSO Agent, see the procedure in the *SonicWall Directory Services Connector Administration Guide*. You can download this guide from https://www.sonicwall.com/support/technical-documentation.

Installing the SonicWall Terminal Services Agent

Install the SonicWall TSA on one or more terminal servers on your network within the Windows domain. The SonicWall TSA must have access to your SonicWall Security Appliance, and the Security Appliance must have access to the TSA. If you have a software firewall running on the terminal server, you might need to open up the UDP port number for incoming messages from the Security Appliance.

SonicWall TSA is available for download without charge from MySonicWall.

For instructions on installing the SonicWall TSA, see the *SonicWall Directory Services Connector Administration Guide*.

Topics:

- Accessing the SonicWall Terminal Services Agent on page 103
- Creating a SonicWall TSA Trouble Shooting Report on page 103

Accessing the SonicWall Terminal Services Agent

After installing the SonicWall TSA and restarting your Windows Server system, you can double click the SonicWall TSA desktop icon created by the installer to launch it for configuration, to generate a trouble shooting report (TSR), or to see the status and version information.



For more information, see the *SonicWall Directory Services Connector Administration Guide*.

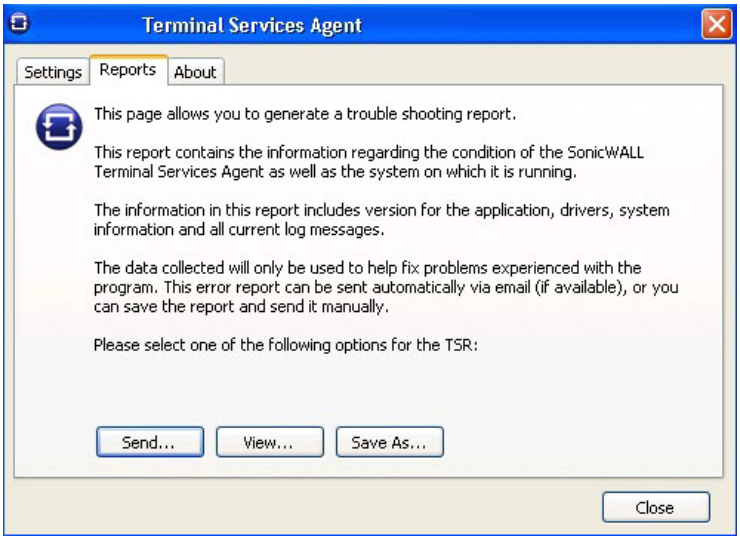
Creating a SonicWall TSA Trouble Shooting Report

You can create a trouble shooting report (TSR) containing all current log messages and information about the agent, driver, and system settings to examine or to send to SonicWall Technical Support for assistance.

To create a TSR for the SonicWall TSA:

1 Double-click the **SonicWall TSA** desktop icon. The **SonicWall Terminal Services Agent** window displays.

2 Click the **Reports** tab.



3 To generate the TSR and:

- Automatically email it to SonicWall Technical Support, click **Send**.
- Examine it in your default text editor, click **View**.
- Save it as a text file, click **Save As**.

4 When finished, click **Close**.

Single Sign-On Advanced Features

Topics:

- About Single Sign-On on page 104
- About the Advanced Settings on page 105
- Viewing SSO Mouseover Statistics on page 105
- Using the Single Sign-On Statistics in the TSR on page 106
- Examining the Agent on page 107
- Remedies on page 107

About Single Sign-On

When a user first tries to send traffic through a SonicWall Security Appliance that is using Single Sign-On (SSO), the Security Appliance sends a “who is this” request to SonicWall SSO Agent. The agent queries the user’s PC through Windows networking, and returns the user name to the firewall. If the user name matches any criteria set in the policies, then the user is considered as “logged on” by the SonicWall. When users are logged into the SonicWall using SSO, the SSO feature also provides detection of logouts. To detect logouts, the Security Appliance repeatedly polls the agent to check if each user is still logged in. This polling, along with the initial identification requests, could potentially result in a large loading on the SonicWall SSO Agent application and the PC on which it is running, especially when very large numbers of users are connecting.

The SonicWall SSO feature utilizes a rate-limiting mechanism to prevent the appliance from swamping the agent with these user requests. Both automatic calculations and a configurable setting on the appliance govern how this rate-limiting operates. The SonicWall SSO feature automatically calculates the maximum number of user

requests contained in each message to the agent that can be processed in the poll period, based on recent polling response times. Also, the timeout on a multi-user request is automatically set to be long enough to reduce the likelihood of an occasional long timeout during polling. The configurable setting controls the number of requests to send to the agent at a time, and can be tuned to optimize SSO performance and prevent potential problems. This section provides a guide to choosing suitable settings.

The potential for problems resulting from overloading the agent can be reduced by running the agent on a dedicated high-performance PC, and possibly also by using multiple agents on separate PCs, in which case the load is shared between them. The latter option also provides redundancy in case one of the agent PCs fails. The agent should run on a Windows Server PC.

About the Advanced Settings

The **Maximum requests to send at a time** setting is available when configuring SSO agents (for more information about configuring SSO agents, see Configuring SonicOS to Use the SonicWall SSO Agent on page 167).

This setting controls the maximum number of requests that can be sent from the appliance to the agent at the same time. The agent processes multiple requests concurrently, spawning a separate thread in the PC to handle each. Sending too many requests at a time can overload the PC on which the agent is running. If the number of requests to send exceeds the maximum, then some are placed on an internal “ring buffer” queue (see Using the Single Sign-On Statistics in the TSR on page 106 and Viewing SSO Mouseover Statistics on page 105). Requests waiting on the ring buffer for too long could lead to slow response times in SSO authentication.

This setting works in conjunction with the automatically calculated number of user requests per message to the agent when polling to check the status of logged in users. The number of user requests per message is calculated based on recent polling response times. SonicOS adjusts this number as high as possible to minimize the number of messages that need to be sent, which reduces the load on the agent and helps reduce network traffic between the appliance and the agent. However, the number is kept low enough to allow the agent to process all of the user requests in the message within the poll period. This avoids potential problems such as timeouts and failures to quickly detect logged out users.

Viewing SSO Mouseover Statistics

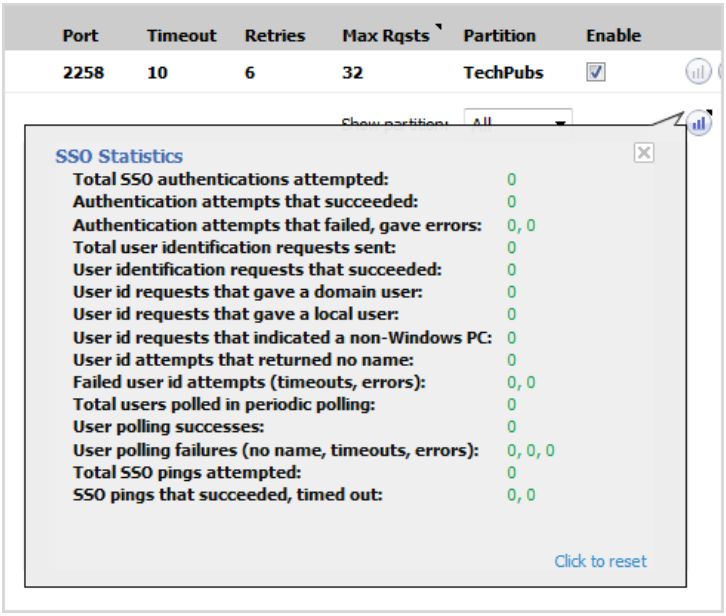
The **SSO Authentication Configuration** dialog provides mouseover statistics about each agent and for all SSO agents. On the **SSO Agents** page, a green LED-style icon next to an agent indicates the agent is up and running. A red LED icon indicates the agent is down.

To view the statistics for:

- A particular agent, hover your mouse over the **Statistics** icon for the SSO agent.

- All SSO agents, hover your mouse over the **Statistics** icon under the table.

  TIP: This also works for individual TSAs on **Terminal Services**.



To close the statistics display, click the **Close** icon.

To clear all the displayed values, click **Click to reset**.

Using the Single Sign-On Statistics in the TSR

A rich set of SSO performance and error statistics is included in the Tech Support Report (TSR). These can be used to gauge how well SSO is performing in your installation. Download the TSR on the **INVESTIGATE | Tools > System Diagnostics** page and search for the title, SSO operation statistics. Here are the counters to look at in particular:

1. Under **SSO ring buffer statistics**, look at **Ring buffer overflows** and **Maximum time spent on ring**. If the latter approaches or exceeds the polling rate, or if any ring buffer overflows are shown, then requests are not being sent to the agent quickly enough. Also, if the **Current requests waiting on ring** is constantly increasing, that would indicate the same. This means that the **Maximum requests to send at a time** value should be increased to send requests faster. However, that increases the load on the agent, and if the agent cannot handle the additional load, then problems result, in which case it might be necessary to consider moving the agent to a more powerful PC or adding additional agents.
2. Under **SSO operation statistics**, look at **Failed user id attempts with time outs** and **Failed user id attempts with other errors**. These should be zero or close to it – significant failures shown here indicate a problem with the agent, possibly because it cannot keep up with the number of user authentications being attempted.
3. Also under **SSO operation statistics**, look at the **Total users polled in periodic polling**, **User polling failures with time outs**, and **User polling failures with other errors**. Seeing some timeouts and errors here is acceptable and probably to be expected, and occasional polling failures do not cause problems. However, the error rate should be low (an error rate of about 0.1% or less should be acceptable). Again, a high failure rate here would indicate a problem with the agent, as above.
4. Under **SSO agent statistics**, look at the **Avg user ID request time** and **Avg poll per-user resp time**. These should be in the region of a few seconds or less – something longer indicates possible problems on the network. Note, however, that errors caused by attempting to authenticate traffic from non-Windows PCs through SSO (which can take a significantly long time) can skew the **Avg user ID request time** value, so if

this is high but **Avg poll per-user resp time** looks correct, that would indicate the agent is probably experiencing large numbers of errors, likely due to attempting to authenticate non-Windows devices – see Step 6.

5 If using multiple agents, then also under **SSO agent statistics** look at the error and timeout rates reported for the different agents, and also their response times. Significant differences between agents could indicate a problem specific to one agent that could be addressed by upgrading or changing settings for that agent in particular.

6 Traffic from devices other than PCs can trigger SSO identification attempts and that can cause errors and/or timeouts to get reported in these statistics. This can be avoided by configuring an address object group with the IP addresses of such devices, and doing one or both of the following:

- If using Content Filtering, select that address object with the **Bypass the Single Sign On process for traffic from** setting on the **Enforcement** tab of the SSO configuration dialog.
- If access rules are set to allow only authenticated users, set separate rules for that address object with **Users Allowed** set to **All**.

To identify the IP addresses concerned, look in the TSR and search for "IP addresses held from SSO attempts". This lists SSO failures in the preceding period set by the **Hold time after failure** setting.

NOTE: If any of the listed IP addresses are for are Mac/Linux PCs, see Accommodating Mac and Linux Users on page 108.

To limit the rate of errors due to this, you can also extend the **Hold time after failure** setting on the Users tab.

Examining the Agent

If the statistics in the TSR report indicate a possible problem with the agent, a good next step would be to run Windows Task Manager on the PC on which the agent is running and look at the CPU usage on the **Performance** tab, plus the CPU usage by the `CIAService.exe` process on the Processes tab. If the latter is using a large percentage of the CPU time and the CPU usage is spiking close to 100%, this is an indication that the agent is getting overloaded. To try to reduce the loading, you can decrease the **Maximum requests to send at a time** setting; see Using the Single Sign-On Statistics in the TSR above, Step 1.

Remedies

If the settings cannot be balanced to avoid overloading the agent's PC while still being able to send requests to the agent fast enough, then one of the following actions should be taken:

- Consider reducing the polling rate configured in the **Users** section of the **SSO Authentication** dialog by increasing the poll time. This reduces the load on the agent, at the cost of detecting logouts less quickly.

  NOTE: In an environment with shared PCs, it is probably best to keep the poll interval as short as possible to avoid problems that could result from not detecting logouts when different users use the same PC, such as the initial traffic from the second user of a PC possibly being logged as sent by the previous user.

- Move the agent to a higher-performance, dedicated PC.
- Configure an additional agent or agents.

Configuring Access Rules

Enabling SonicWall SSO affects policies on the **Rules > Access Rules** page of the SonicOS Management Interface. Rules set under **Rules > Access Rules** are checked against the user group memberships returned from a SSO LDAP query, and are applied automatically.

Topics:

- Automatically Generated Rules for SonicWall SSO on page 108
- Accommodating Mac and Linux Users on page 108
- Allowing ICMP Pings from a Terminal Server on page 109
- About Access Rules on page 109

Automatically Generated Rules for SonicWall SSO

When a SonicWall SSO agent or TSA is configured in the SonicOS Management Interface, an access rule and corresponding NAT policy are created to allow the replies from the agent into the LAN. These rules use either a **SonicWall SSO Agents** or **SonicWall Terminal Services Agents** address group object, which has a member address object for each configured agent. The member address objects are automatically added to and deleted from the group object as agents are added or deleted. The member address objects are also updated automatically as an agent's IP address changes, including when an IP address is resolved through DNS (where an agent is given by DNS name).

If SonicWall SSO agents or TSAs are configured in different zones, the access rule and NAT policy are added to each applicable zone. The same **SonicWall SSO Agents** or **SonicWall Terminal Services Agents** address group is used in each zone.

NOTE: Do not enable Guest Services in the same zone where SonicWall SSO is being used. Enabling Guest Services disables SSO in that zone, causing users who have authenticated through SSO to lose access. Create a separate zone for Guest Services.

Accommodating Mac and Linux Users

Mac and Linux systems do not support the Windows networking requests that are used by the SonicWall SSO agent, but can use Samba 3.5 or newer to work with SonicWall SSO.

Using SSO on Mac and Linux With Samba

For Windows users, SonicWall SSO is used by a Security Appliance to automatically authenticate users in a Windows domain. It allows the users to get access through the Security Appliance with correct filtering and policy compliance without the need to identify themselves through any additional login process after their Windows domain login.

Samba is a software package used by Linux/UNIX or Mac machines to give their users access to resources in a Windows domain (through Samba's `smbclient` utility) and/or to give Windows domain users access to resources on the Linux or Mac machine (through a Samba server).

A user working on a Linux PC or Mac with Samba in a Windows domain can be identified by SonicWall SSO, but it requires proper configuration of the Linux/Mac machine, the SSO Agent, and possibly some reconfiguration of the appliance. For example, the following configuration is necessary:

- To use SonicWall SSO with Linux/Mac users, the SonicWall SSO Agent must be configured to use **NetAPI** rather than **WMI** to get the user login information from the user's machine.
- For Samba to receive and respond to the requests from the SonicWall SSO Agent, it must be set up as a member of the domain and the Samba server must be running and properly configured to use domain authentication.

SonicWall SSO is supported by Samba 3.5 or newer.

NOTE: If multiple users log into a Linux PC, access to traffic from that PC is granted based on the most recent login.

Using SSO on Mac and Linux Without Samba

Without Samba, Mac and Linux users can still get access, but you need to log in to the firewall to do so. This can cause the following problems:

- Traffic from Mac or Linux systems might keep triggering SSO identification attempts unless the user logs in. This could potentially be a performance overhead to the SSO system if there are a large number of such systems, although the effect would be somewhat mitigated by the “hold after failure” timeout.
- If per-user Content Filtering (CFS) policies are used without policy rules with user level authentication, the default CFS policy is applied to users of Mac and Linux systems unless they manually log in first.
- If policy rules are set requiring user level authentication, Web browser connections from users of Mac and Linux systems are redirected to the login page after the SSO failure, but the failure might initiate a timeout that would cause a delay for the user.

To avoid these problems, the **Don't invoke Single Sign On to Authenticate Users** option is available when configuring access rules on the **MANAGE | Policies** > **Rules > Access Rules** page (for more information about configuring access rules, see *SonicOS 6.5 Policies*). This option is visible only when SonicWall SSO is enabled. If this option is selected, SSO is not attempted for traffic that matches the rule, and unauthenticated HTTP connections that match it are directed straight to the login page. Typically, the **Source** drop-down menu would be set to an address object containing the IP addresses of Mac and Linux systems.

In the case of CFS, a rule with this option enabled can be added “in front of” CFS so that HTTP sessions from Mac and Linux systems are automatically redirected to log in, avoiding the need for these users to log in manually.

NOTE: Do not select the **Don't invoke Single Sign On to Authenticate Users** option for use with devices that are allowed to bypass the user authentication process entirely. Any devices that could be affected by an access rule when this option is enabled must be capable of logging in manually. A separate access rule should be added for such devices, with **Users Allowed** set to **All**.

Allowing ICMP Pings from a Terminal Server

In Windows, outgoing ICMP pings from users on the Terminal Server are not sent through a socket, so they are not seen by the TSA, and hence the Security Appliance receives no notifications for them. Therefore, if firewall rules are using user-level authentication and pings are to be allowed through, you must create separate access rules to allow them from All.

About Access Rules

Access rules provide you with the ability to control user access. Rules set from the **MANAGE | Policies > Rules > Access Rules** page are checked against the user group memberships returned from a SSO LDAP query, and are applied automatically. Access rules are network management tools that allow you to define inbound and outbound access policy, configure user authentication, and enable remote management of the Security Appliance. The **Rules > Access Rules** page provides a sortable access rule management interface.

NOTE: More specific policy rules should be given higher priority than general policy rules. The general specificity hierarchy is source, destination, service. User identification elements, for example, user name and corresponding group permissions, are not included in defining the specificity of a policy rule.

By default, the firewall’s stateful packet inspection allows all communication from the LAN to the Internet, and blocks all traffic to the LAN from the Internet.

Additional network access rules can be defined to extend or override the default access rules. For example, access rules can be created that block certain types of traffic such as IRC from the LAN to the WAN, or allow certain types of traffic, such as Lotus Notes database synchronization, from specific hosts on the Internet to specific hosts on the LAN, or restrict use of certain protocols such as Telnet to authorized users on the LAN.

CAUTION: The ability to define network access rules is a powerful tool. Using custom access rules can disable firewall protection or block all access to the Internet. Use caution when creating or deleting network access rules.

For detailed information about access rules, see *SonicOS 6.5 Policies*.

Managing SonicOS with HTTP Login from a Terminal Server

The SonicWall Security Appliance normally grants access through policies based on authentication credentials supplied through an HTTP login for one user at an IP address. For users on a terminal server, this method of authenticating one user per IP address is not possible. However, HTTP login is still allowed from a terminal server only for the purpose of administration of the appliance, subject to the following limitations and requirements:

- Internet access from the terminal server is controlled from the TSA, and HTTP login does not override that — a user on a terminal server is not granted any access through the Security Appliance based on credentials supplied through an HTTP login.
- HTTP login from a terminal server is allowed only for the built-in **admin** account and other user accounts with administrator privileges. An attempt to log in with a non-administrative account fails with the error, `Not allowed from this location`.
- On successful HTTP login, an administrative user is taken straight to the Management Interface. The small **User Login Status** page is not displayed.
- The administrative user account used for HTTP login from the terminal server does not need to be the same user account that was used for login to the terminal server. It is shown on the Security Appliance as an entirely separate login session.
- Only one user at a time can manage the Security Appliance from a given terminal server. If two users attempt to do so simultaneously, the most recently logged in user takes precedence, and the other user sees the error, `This is not the browser most recently used to log in`.
- On a failure to identify a user due to communication problems with the TSA, an HTTP browser session is not redirected to the Web login page (as happens on a failure in the SSO case). Instead, it goes to a new page with the message, `The destination that you were trying to reach is temporarily unavailable due to network problems`.

Viewing and Managing SSO User Sessions

Topics:

- Logging Out SSO Users on page 110
- Configuring Additional SSO User Settings on page 111
- Viewing SSO and LDAP Messages with Packet Monitor on page 111
- Capturing SSO Messages on page 111
- Capturing LDAP Over TLS Messages on page 111

Logging Out SSO Users

The **MONITOR | Current Status > User Sessions > Active Users** page displays user sessions on the Security Appliance. For information about viewing the user's settings and how to log out users, see *SonicOS 6.5 Monitoring*.

NOTE: Changes in a user's settings, configured under **MANAGE | System Setup > Users > Settings**, are not reflected during that user's current session; you must manually log the user out for changes to take effect. The user is transparently logged in again, with the changes reflected.

Configuring Additional SSO User Settings

The **MANAGE | System Setup > Users > Settings** page provides configuration options for user session settings, global user settings, and acceptable use policy settings, in addition to SSO and other user login settings.

The options to limit user sessions under **User Session** apply to users logged in using SSO. SSO users are logged out according to session limit settings, but are automatically and transparently logged back in when they send further traffic.

NOTE: Do not set the login session limit interval too low. This could potentially cause performance problems, especially for deployments with many users.

Changes applied in the **Users > Settings** page during an active SSO session are not reflected during that session.

TIP: You must log the user out for changes to take effect. The user is immediately and automatically logged in again, with the changes made.

Viewing SSO and LDAP Messages with Packet Monitor

The Packet Monitor feature available on **Investigate | Tools > Packet Monitor** provides options to enable capture of decrypted messages to and from the SSO agent, and decrypted LDAP over TLS (LDAPS) messages. For further information, see *SonicOS 6.5 Investigate*.

Capturing SSO Messages

For further information about using the Packet Monitor, see *SonicOS 6.5 Investigate*.

To capture decrypted messages to or from the SSO authentication agent:

1. Navigate to **INVESTIGATE | Tools > Packet Monitor**.
2. Under the **Hex Dump** section, click **CONFIGURE**. The **Packet Monitor Configuration** dialog displays.
3. Click **Advanced Monitor Filter.**
4. Select **Monitor intermediate Packets**.
5. Select **Monitor intermediate decrypted Single Sign On agent messages**.
6. Click **OK**.

The packets are marked with **(sso)** in the ingress/egress interface field. They have dummy Ethernet, TCP, and IP headers, so some values in these fields might not be correct.

This enables decrypted SSO packets to be fed to the Packet Monitor, but any monitor filters are still applied to them.

Captured SSO messages are displayed fully decoded on the **Tools > Packet Monitor** page.

Capturing LDAP Over TLS Messages

To capture decrypted LDAP over TLS (LDAPS) packets:

1. Navigate to **INVESTIGATE | Tools > Packet Monitor**.
2. Under the **Hex Dump** section, click **CONFIGURE**. The **Packet Monitor Configuration** dialog displays.
3. Click **Advanced Monitor Filter.**
4. Select **Monitor intermediate Packets**.
5. Select **Monitor intermediate decrypted LDAP over TLS packets**.
6. Click **OK**.

The packets are marked with (**ldp**) in the ingress/egress interface field. They have dummy Ethernet, TCP, and IP headers, so some values in these fields might not be correct. The LDAP server port is set to 389 so that an external capture analysis program (such as Wireshark) knows to decode these packets as LDAP. Passwords in captured LDAP bind requests are obfuscated. The LDAP messages are not decoded in the Packet Monitor display, but the capture can be exported and displayed in WireShark to view them decoded.

This enables decrypted LDAPS packets to be fed to the packet monitor, but any monitor filters are still applied to them.

NOTE: LDAPS capture only works for connections from the firewall's LDAP client, and does not display LDAP over TLS connections from an external LDAP client that pass through the firewall.

About Multiple Administrator Support

You can configure multiple administrator profiles, as described in Configuring Local Users and Groups on page 221.

When using RADIUS or LDAP authentication, if you want to ensure that some or all administrative users are always able to manage the appliance, even when the RADIUS or LDAP server becomes unreachable, then you can use the **RADIUS + Local Users or LDAP + Local Users** option and configure the accounts for those particular users locally.

For users authenticated by RADIUS or LDAP, create user groups named **SonicWall Administrators** and/or **SonicWall Read-Only Admins** on the RADIUS or LDAP server (or its back-end) and assign the relevant users to those groups.

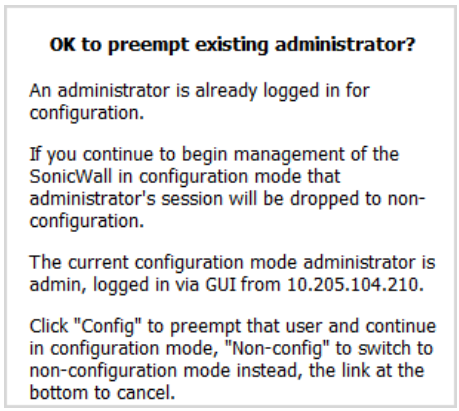
NOTE: For RADIUS, you probably need a special configuration of the RADIUS server to return the user group information.

Topics:

- Preempting Administrators on page 112
- Logging In with Administrator Rights on page 113

Preempting Administrators

When an administrator attempts to log in while another administrator is logged in, this message is displayed:



This message gives you three options:

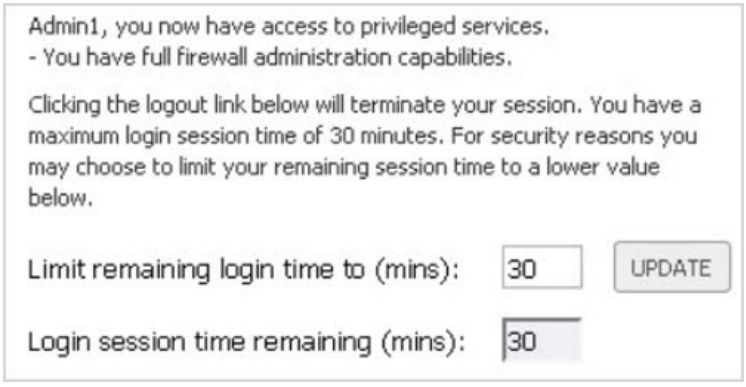
Config	Preempts the current administrator. The current administrator is dropped to non-config mode and you are given full administrator access.
Non-config	You are logged into the SonicWall Security Appliance in non-config mode. The current administrator's session is not disturbed.
do not begin management	Returns to the login screen.

Logging In with Administrator Rights

A user other than admin (that is, not the **admin** user) can log in with administrator rights.

To log in with Administrator rights:

1 Log in with your administrator credentials. The **User Login Status** message displays.



2 To

- Go to the SonicWall management interface, click **MANAGE**. You are prompted to enter your password again. This is a safeguard to protect against unauthorized access when administrators are away from their computers and do not log out of their sessions.
- Change your password click **CHANGE PASSWORD**. The dialog for changing your password displays.

Disabling the User Login Status Popup

You can disable the **User Login Status** popup, if you prefer to allow certain users to log in solely for the purpose of managing the SonicWall Security Appliance, rather than for privileged access through the Security Appliance. To disable the popup, select the **Members go straight to the management UI on web login** option when adding or editing the local group.

If you want some user accounts to be administrative only, while other users need to log in for privileged access through the appliance, but also with the ability to administer it (that is, some go straight to the management interface on login, while others get the **User Login Status** popup dialog with a **MANAGE** button), this can be achieved by:

1 Creating a local group with the **Members go straight to the management UI on web login** option selected.

2 Adding the group to the relevant administrative group, but do not select this option in the administrative group.

3 Adding those user accounts that are to be administrative-only to the new user group. The **User Login Status** popup is disabled for these users.

4 Adding the user accounts that are to have privileged and administrative access directly to the top-level administrative group.

Configuring Multiple Administrator Support

Topics:

- Configuring Additional Administrator User Profiles on page 114
- Configuring Administrators Locally when Using LDAP or RADIUS on page 114
- Preempting Administrators on page 112
- Logging In with Administrator Rights on page 113
- Verifying Multiple Administrators Support Configuration on page 115
- Viewing Multiple Administrator Related Log Messages on page 115

Configuring Additional Administrator User Profiles

Configuring additional administrators is the same as configuring additional local users and then adding them to the proper local group:

This group	Gives the user
Limited Administrators	Limited administrator configuration privileges.
SonicWall Administrators	Full administrator configuration privileges.
SonicWall Read-Only Admins	Viewing privileges only for the entire Management Interface.

For how to configure local users and local groups, see Configuring Local Users and Groups on page 221.

Configuring Administrators Locally when Using LDAP or RADIUS

When using RADIUS or LDAP authentication, if you want to ensure that some or all administrative users are always able to manage the SonicWall Security Appliance, even if the RADIUS or LDAP server becomes unreachable, then you can use the **RADIUS + Local Users** or **LDAP + Local Users** option and configure the accounts for those particular users locally.

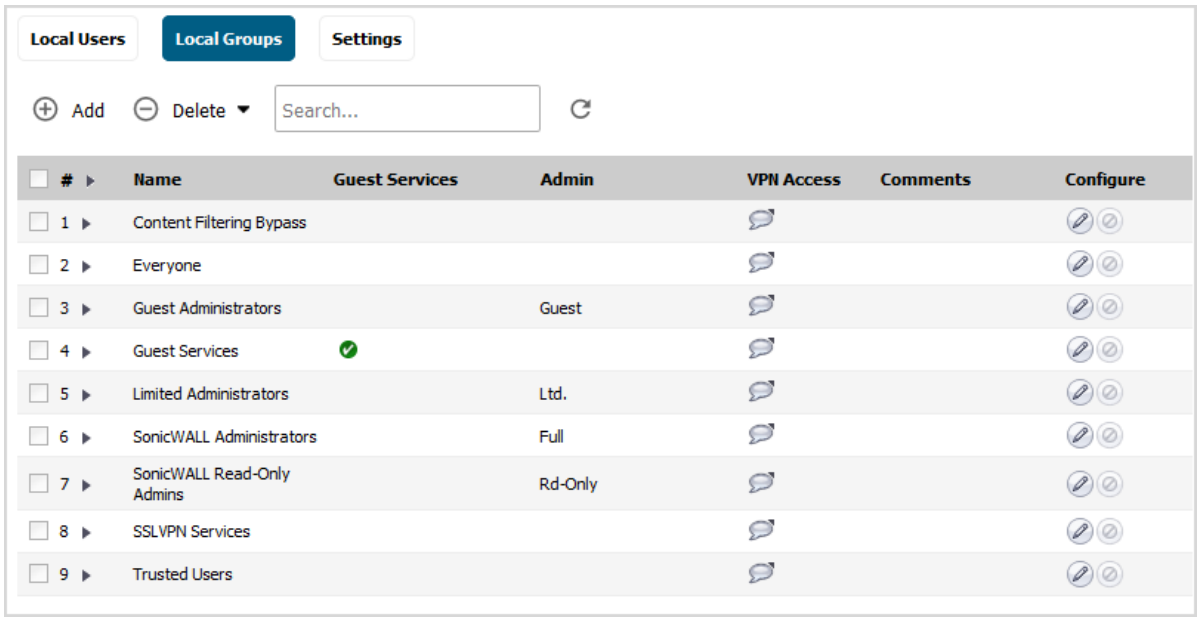
For users authenticated by RADIUS or LDAP, create user groups named **SonicWall Administrators** and/or **SonicWall Read-Only Admins** on the RADIUS or LDAP server (or its back-end) and assign the relevant users to those groups.

NOTE: For RADIUS, you probably need a special configuration of the RADIUS server to return the user group information.

For how to configure administrators when using LDAP or RADIUS, see Configuring Local Users and Groups on page 221.

Verifying Multiple Administrators Support Configuration

User accounts with administrator and read-only administrators can be viewed on the **MANAGE | System Setup > Users > Local Users & Groups > Local Groups** page.



You can determine which configuration mode you are in by looking at **Mode** in the top right corner of the Management Interface:



When changes are made, the status bar reads:





When changes are attempted, the status bar reads:

Viewing Multiple Administrator Related Log Messages

Log messages are generated for these events:

- A GUI or CLI user begins configuration mode (including when an admin logs in).
- A GUI or CLI user ends configuration mode (including when an admin logs out).
- A GUI user begins management in non-config mode (including when an admin logs in and when a user in configuration mode is preempted and dropped back to read-only mode).
- A GUI user begins management in read-only mode.

A GUI user terminates either of the above management sessions (including when an admin logs out).

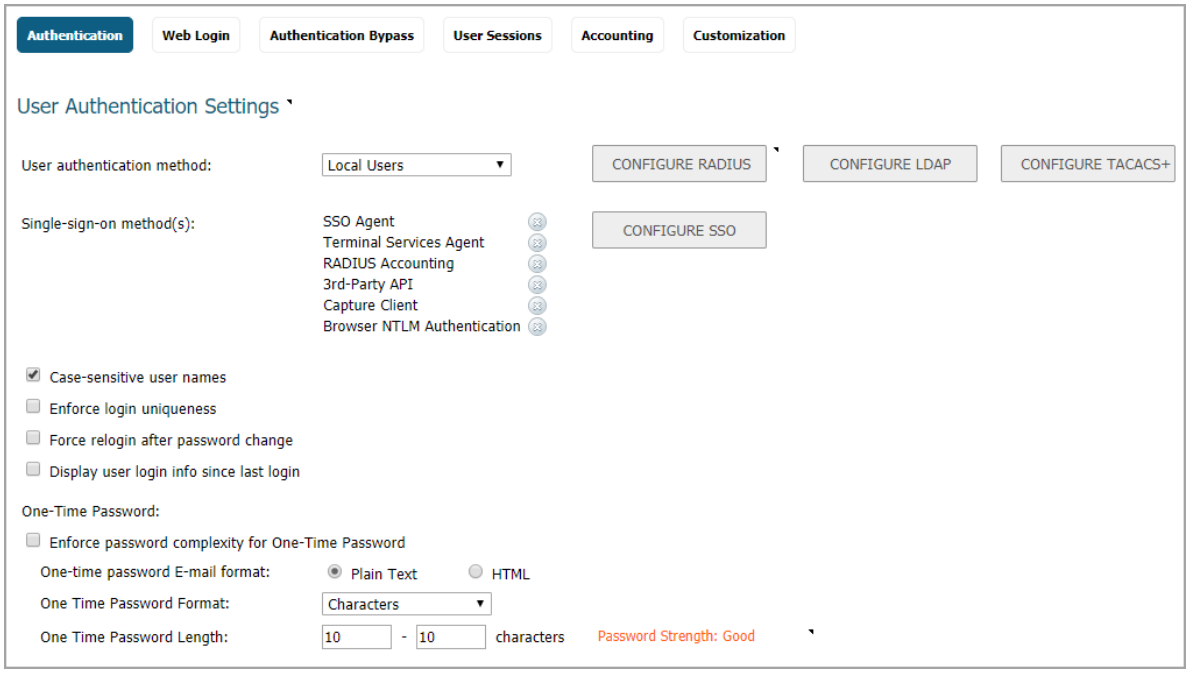
8

Configuring Settings for Managing Users

Topics:

- Users > Settings on page 116
  - Configuring User Authentication and Login Settings on page 117
  - Customization on page 132
  - Accounting on page 139
  - Configuring RADIUS Authentication on page 142
  - Configuring the SonicWall for LDAP on page 149
  - About LDAP Test on page 165
  - Configuring TACACS + for Authentication on page 166
  - Configuring SonicOS to Use the SonicWall SSO Agent on page 167

Users > Settings



On **MANAGE | System Setup | Users > Settings**, you can configure the authentication method required, global user settings, and an acceptable user policy that is displayed to users when logging onto your network.

Topics:

- Configuring User Authentication and Login Settings on page 117
- Configuring User Sessions on page 128
- Configuring RADIUS Authentication on page 142
- Configuring the SonicWall for LDAP on page 149
- Configuring TACACS + for Authentication on page 166
- Configuring SonicOS to Use the SonicWall SSO Agent on page 167

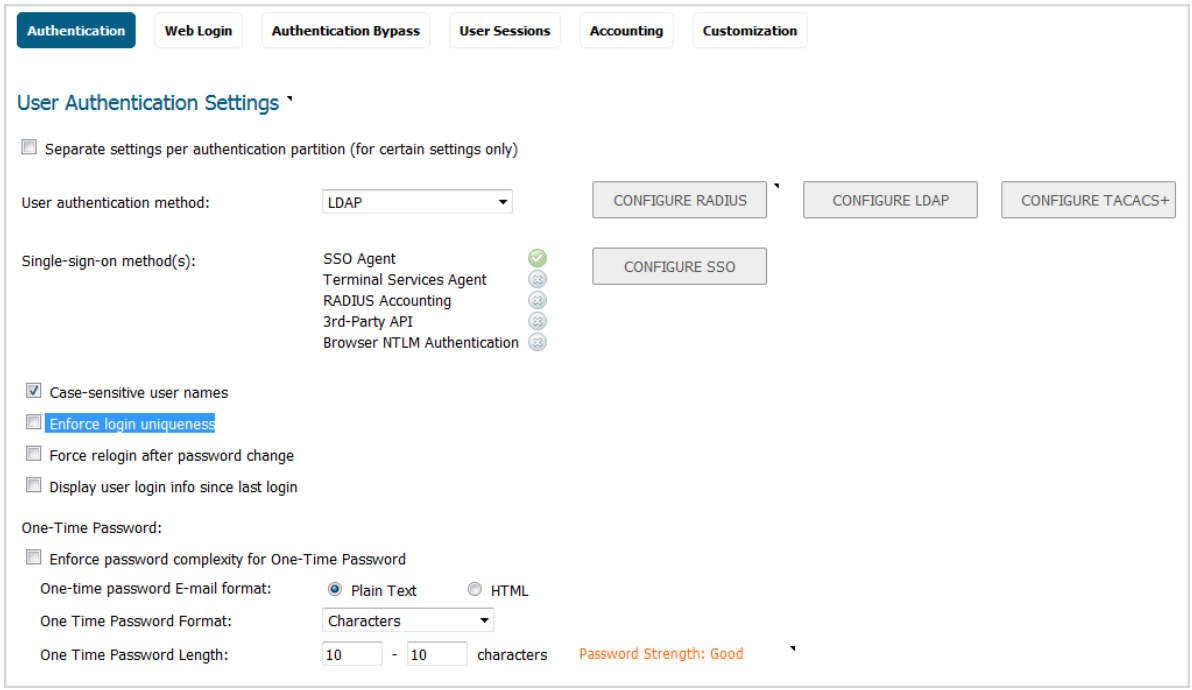
Configuring User Authentication and Login Settings

IMPORTANT: When you have finished configuring the **Users > Settings** page, click **ACCEPT**.

Topics:

- User Authentication Settings on page 118
- User Web Login Settings on page 121
- Captive Portal Authentication on page 123
- Authentication Bypass Settings on page 124
- User Session Settings on page 129
- User Session Settings for SSO Authenticated Users on page 130
- User Session Settings for Web Login on page 131
- Post-Login Acceptable Use Policy on page 134
- Customized Login Pages on page 136

User Authentication Settings



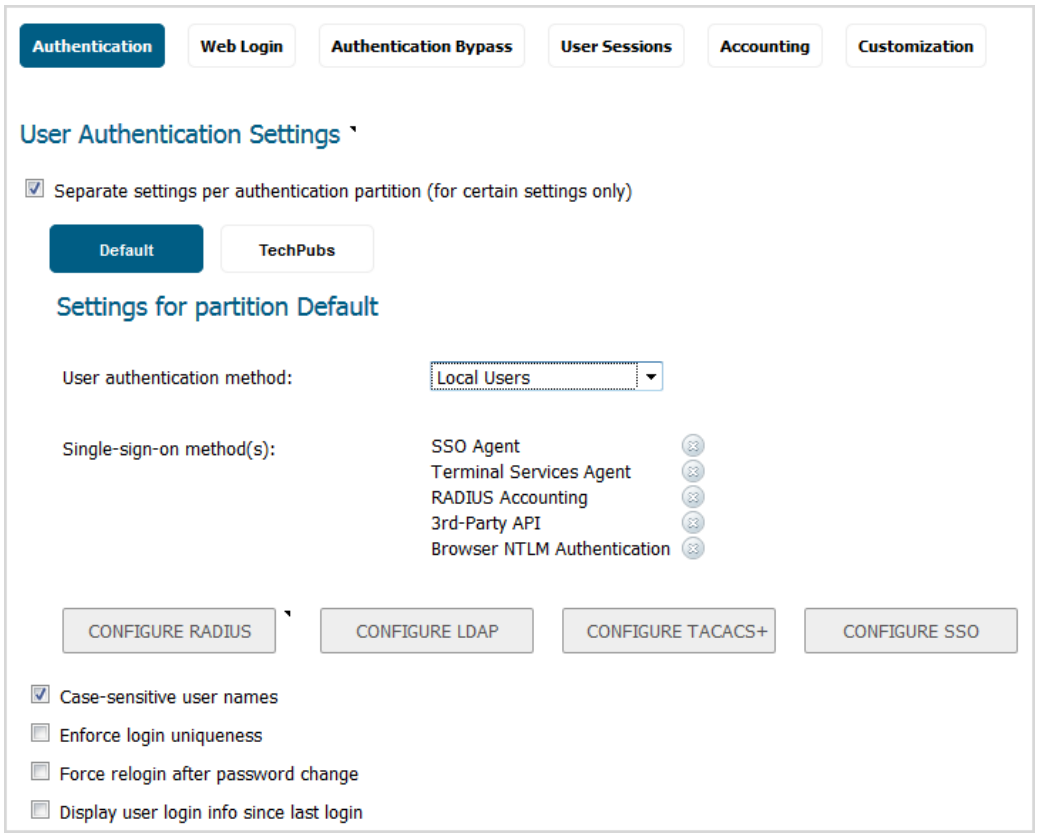
To configure user authentication settings:

1 Navigate to **MANAGE | System Setup | Users > Settings**.

2 If partitioning is:

- Not enabled, go to Step 4.

- Enabled, the **Separate settings per authentication partition (for certain settings only)** option displays. Select the option to display the **Settings for partition** options.



3 For each partition, perform Step 4 onward.

4 From **User Authentication method**, select the type of user account management your network uses:

Local Users

To configure users in the local database in the Security Appliance using the **MANAGE | System Setup | Users > Local Users & Groups** page.

For information about using the local database for authentication and detailed configuration instructions, see these sections, see Using Local Users and Groups for Authentication on page 84.

RADIUS

You have more than 1,000 users or want to add an extra layer of security for authenticating the user to the Security Appliance. If you select RADIUS for user authentication, users must log into the Security Appliance using HTTPS to encrypt the password sent to the Security Appliance. If a user attempts to log into the Security Appliance using HTTP, the browser is automatically redirected to HTTPS.

RADIUS might be required in addition to LDAP in a number of cases:

- LDAP does not usually support CHAP/MS-CHAP authentication (Microsoft Active Directory and Novell eDirectory do not), so the SonicWall authenticates CHAP/MS-CHAP through RADIUS if that is the case and RADIUS is configured.
- If NTLM is used for SSO, it can only be authenticated through RADIUS in MS-CHAP mode.

	RADIUS might be required for CHAP/MS-CHAP with L2TP servers or with VPN or SSL VPN clients, including NetExtender and Portal, or if it might be required for NTLM. **NOTE:** LDAP is generally still used for non-CHAP authentications when RADIUS is used for CHAP. For information about using a RADIUS database for authentication, see Using RADIUS for Authentication on page 87. For detailed configuration instructions, see Configuring RADIUS Authentication on page 142.
RADIUS + Local Users	You want to use both RADIUS and the Security Appliance local user database for authentication.
LDAP	You use a Lightweight Directory Access Protocol (LDAP) server, Microsoft Active Directory (AD) server, or Novell eDirectory to maintain all your user account data. For information about using an LDAP database for authentication, see Using LDAP/Active Directory/eDirectory Authentication on page 87. For detailed configuration instructions, see Integrating LDAP into the SonicWall Security Appliance on page 90.
LDAP + Local Users	You want to use both LDAP and the Security Appliance local user database for authentication.
TACACS +	You want to use TACACS + (Terminal Access Controller Access-Control System latest generation) for user authentication.
TACACS + + Local User	You want to use both TACACS + and the Security Appliance's local user database for authentication.

5 For **Single-sign-on method**, select one of the following:

NOTE: Do not select any of these options if you are not using Single Sign-On to authenticate users.

SSO Agent	You are using Active Directory for authentication and the SSO Agent is installed on a computer in the same domain. For detailed SSO configuration instructions, see About Single Sign-On on page 104.
Terminal Services Agent	You are using Terminal Services and the Terminal Services Agent (TSA) is installed on a terminal server in the same domain.
RADIUS Accounting	You want a network access server (NAS) to send user login session accounting messages to an accounting server.
3rd Party API	You want to use a third-party API.
Capture Client	You want to use SonicWall Capture Client for user authentication.
Browser NTLM authentication only	You want to authenticate Web users without using the SSO Agent or TSA. Users are identified as soon as they send HTTP traffic. NTLM requires RADIUS to be configured (in addition to LDAP, if using LDAP), for access to MSCHAP authentication. If LDAP is selected above, a separate **CONFIGURE** for RADIUS appears here when NTLM is selected.

6 Select **Case-sensitive user names** to enable matching based on capitalization of user account names.

7 Select **Enforce login uniqueness** to prevent the same user name from being used to log into the network from more than one location at a time. This option applies to both local users and RADIUS/LDAP users, but it does not apply to the default administrator with the username, **admin**. This option is not selected by default.

8 To make users log in after changing their passwords, select **Force relogin after password change**. This option is not selected by default.

9 To display user login information since the last log in on the **MONITOR | Current Status | System Status** page, select **Display user login info since last login**. This option is not selected by default.

If this option is enabled, user login information—including last successful login timestamp, number of all user successful login attempts, unsuccessful login attempts, and administrator privilege changes—are displayed in **INVESTIGATE | Logs | Event Logs**. For more information about logs, see *SonicOS 6.5 Investigate*.

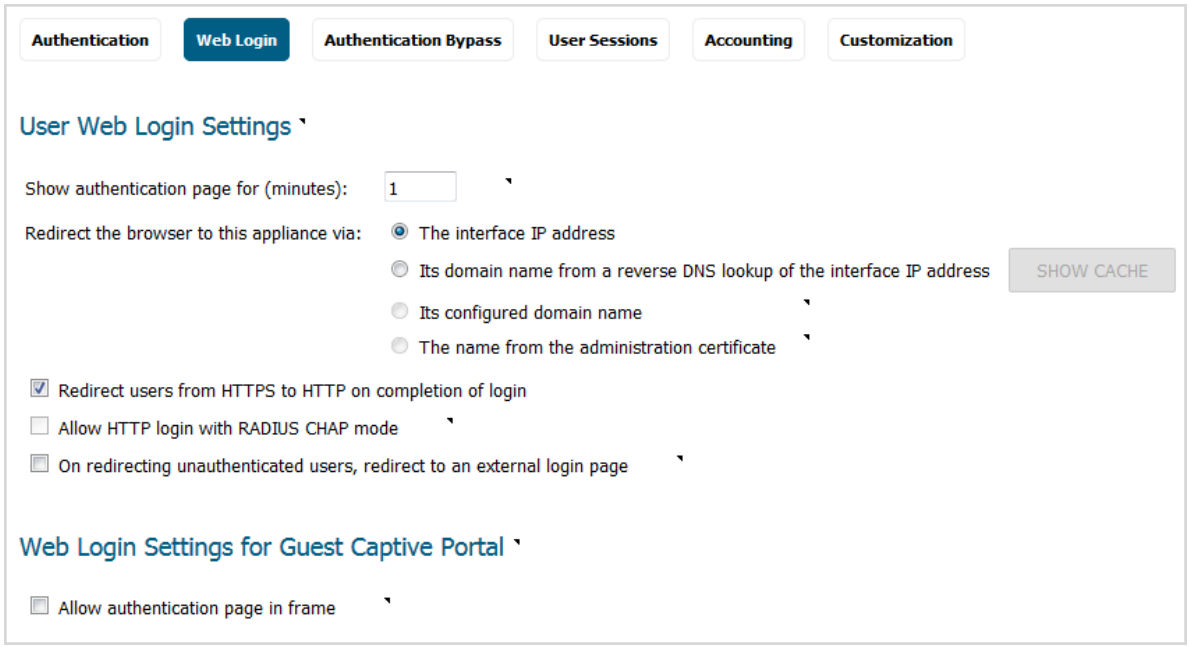
10 Configure the following **One-Time Password** options:

- **One-time password Email format** – Select either **Plain text** or **HTML**.
- **One Time Password Format** – Select **Characters** (default), **Characters+Numbers**, or **Numbers** from the drop-down menu.

  TIP: The format selection along with the two values for password length result in a password strength of Poor, Good, or Excellent. The strongest passwords have long lengths and either **Characters** or **Characters+Numbers** format; The weakest password strength is the **Numbers** format regardless of length.

- At **One Time Password Length**, enter the minimum length in the first field and the maximum length in the second field. The minimum and maximum must be within the range of 4 to 14, with a default value of **10** for each field. The minimum length cannot be greater than the maximum length.

User Web Login Settings



To configure user web login settings:

1 Navigate to **MANAGE | System Setup** | **Users > Settings**.

2 Click **Web Login**.

3 In the **Show user authentication page for (minutes)** field, enter the number of minutes that users have to log in with their username and password before the login page times out. If it times out, a message

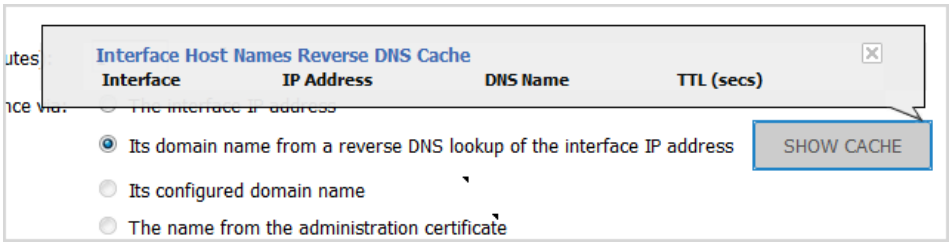
displays informing them what they must do before attempting to log in again. The default time is **1** minute.

While the login authentication page is displayed, it uses system resources. By setting a limit on how long a login can take before the login page is closed, you free up those resources.

4 From **Redirect the browser to this appliance via**, choose the option that determines how a user's browser is initially redirected to the SonicWall appliance's Web server:

- **The interface IP address** – Select this to redirect the browser to the IP address of the appliance Web server interface. This option is selected by default.
- **Its domain name from a reverse DNS lookup of the interface IP address** – Enables **Show Cache** which, when clicked, displays the appliance Web server's Interface, IP Address, DNS Name, and TTL (in seconds). This option is not selected by default.

  Click **Show Cache** to verify the domain name (DNS name) being used for redirecting the user's browser.



- **Its configured domain name** – Select to enable redirecting to a domain name configured on **System Setup | Appliance > Base Settings**. Redirecting to the name from the administration certificate is allowed when an imported certificate has been selected for HTTPS web management on that page.

  NOTE: This option is available only if a domain name has been specified on **System Setup | Appliance > Base Settings**. Otherwise, this option is dimmed.

- **The name from the administration certificate** – Select to enable redirecting to a domain name with a properly signed certificate. Redirecting to the name from this administration certificate is allowed when an imported certificate has been selected for HTTPS web management on that page. Configure the domain name on **System Setup | Appliance > Base Settings**.

  NOTE: This option is available only if a certificate has been imported for HTTPS management in the **Web Management Settings** section of **System Setup | Appliance > Base Settings**. See Configuring Base Settings on page 18.

  TIP: If you are using imported administration certificates, use this option. If you are not going to use an administration certificate, select the **Its configured domain name** option.

  To do HTTPS management without the browser displaying invalid-certificate warnings, you need to import a certificate properly signed by a certification authority (administration certificate) rather than use the internally generated self-signed one. This certificate must be generated for the appliance and its host domain name. A properly signed certificate is the best way to obtain an appliance's domain name.

  If you use an administration certificate, then to avoid certificate warnings, the browser needs to redirect to that domain name rather than to the IP address. For example, if you browse the Internet and are redirected to log in at `https://gateway.sonicwall.com/auth.html`, the administration certificate on the appliance says that the appliance really is `gateway.sonicall.com`, so the browser displays the login page. If you are redirected to `https://10.0.02/auth.html`, however, even though the certificate says it is `gateway.sonicall.com`, the browser has no way to tell if that is correct, so it displays a certificate warning instead.

5 Select **Redirect users from HTTPS to HTTP on completion of login** if you want users to be connected to the network through your Security Appliance through HTTP after logging in through HTTPS. If you have a large number of users logging in through HTTPS, you might want to redirect them to HTTP, because HTTPS consumes more system resources than HTTP. This option is selected by default. If you deselect this option, a warning dialog displays.

6 Select **Allow HTTP login with RADIUS CHAP mode** to have a CHAP challenge be issued when a RADIUS user attempts to log in using HTTP. This allows for a secure connection without using HTTPS. Be sure to check that the RADIUS server supports this option. This option is not selected by default.

> **NOTE:** If you log in using this method, you are restricted in the management operations you can perform because some operations require the appliance to know the administrator's password; with CHAP authentication by a remote authentication server, the appliance does not know the password.
>
> If this setting is checked, therefore, any users who are members of administrative user groups might need to manually log in through HTTPS if logging in for administration. This restriction does not apply to the built-in **admin** account.

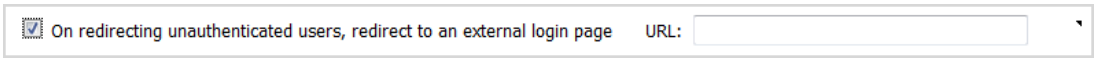
> **NOTE:** When using LDAP, this mechanism can normally be used by setting the **Authentication method for login** to **RADIUS** and then selecting LDAP as the mechanism for setting user group memberships in the RADIUS configuration.

7 To redirect HTTP/HTTPS traffic from unauthenticated users to a specified URL instead of the SonicWall's own login page, select **On redirecting unauthenticated users, redirect to an external login page**. This option allows users to be authenticated by an external authentication system. This option is not selected by default.

> **TIP:** To allow only unauthenticated users to be redirected, you need to create one or more access rules for this situation.

> **NOTE:** The external system can subsequently use the SSO third-party API or RADIUS Accounting to pass the user's name and credentials to the firewall so they are identified for such activities as access control and logging.

When you select this option, the **URL** field displays.



Enter the URL for redirection in the field.

8 To configure options elated to the captive portal configured in a zone's guest settings, scroll to **Web Login Settings for Guest Captive Portal**.

9 For captive portal guest authentication, to allow the authentication page to show in a portal host page as a frame, select **Allow authentication page in frame**. This option is not selected by default.

10 Click **ACCEPT**.

Captive Portal Authentication

Captive Portal Authentication helps a user gain access after being authenticated and authorized. The user who seeks web access to a network is redirected to the authentication web login page hosted on the captive portal server that is integrated with the RADIUS server.

This authentication method is an extension to SonicWall's existing LHM (Lightweight Hotspot Messaging). LHM deployment requires all authentications to be handled by the external LHM server. Captive portal authentication puts more leverage on the firewall itself to communicate with the RADIUS server to complete the authentication process.

To configure captive portal authentication:

1 Configure the RADIUS portal server:

   a Configure the user information and user group information. The user group name must:

      - Be returned to the firewall with an ACCEPT message.
      - Match a group name on the firewall.
      - Have guest privilege on the firewall.

   b Configure the **idle timeout** and **session timeout** attributes if the firewall requires they be returned with the ACCEPT message.

   c Define the welcome URL as a **Vendor Specific Attribute**; SonicWall's vendor code is 8741.

   d If RADIUS accounting is supported, set the interim interval.

2 Navigate to **MANAGE | System Setup | Users > Settings**.

3 Click **Authentication**.

4 Under **User Authentication Settings**, select **RADIUS + LOCAL Users** from **User authentication method**.

5 Click **CONFIGURE RADIUS**. The **RADIUS Configuration** dialog displays.

6 Click **RADIUS Users**.

7 For **Mechanism for looking up user group memberships for RADIUS users**:, choose **Use RADIUS Filter-Id attribute on RADIUS server**.

8 Click **OK**.

9 Navigate to **MANAGE | System Setup | Network > Zones**.

10 Click either the **Add** icon or the **Edit** icon for a wireless zone. The **Add/Edit Zone** dialog displays.

11 Ensure the **Security Type** is **Wireless**.

12 Follow the steps for configuring a zone for captive portal authentication with RADIUS.

Authentication Bypass Settings

SonicOS Guest Services allows guest users to have access through your network directly to the Internet without access to your protected network. To do this, SonicOS uses the IP address of the user's computer.

Using the IP address as the identifier is useful when guest user traffic passes through a network router, as this changes the source MAC address to that of the router. However, the user's IP address passes through unchanged.

If only the MAC address is used for identification, two clients behind the same router have the same MAC address upon reaching the Security Appliance. When one client gets authenticated, the traffic from the other client is also treated as authenticated and bypasses the guest service authentication.

By using the client IP address for identification, all guest clients behind the routed device are required to authenticate independently.

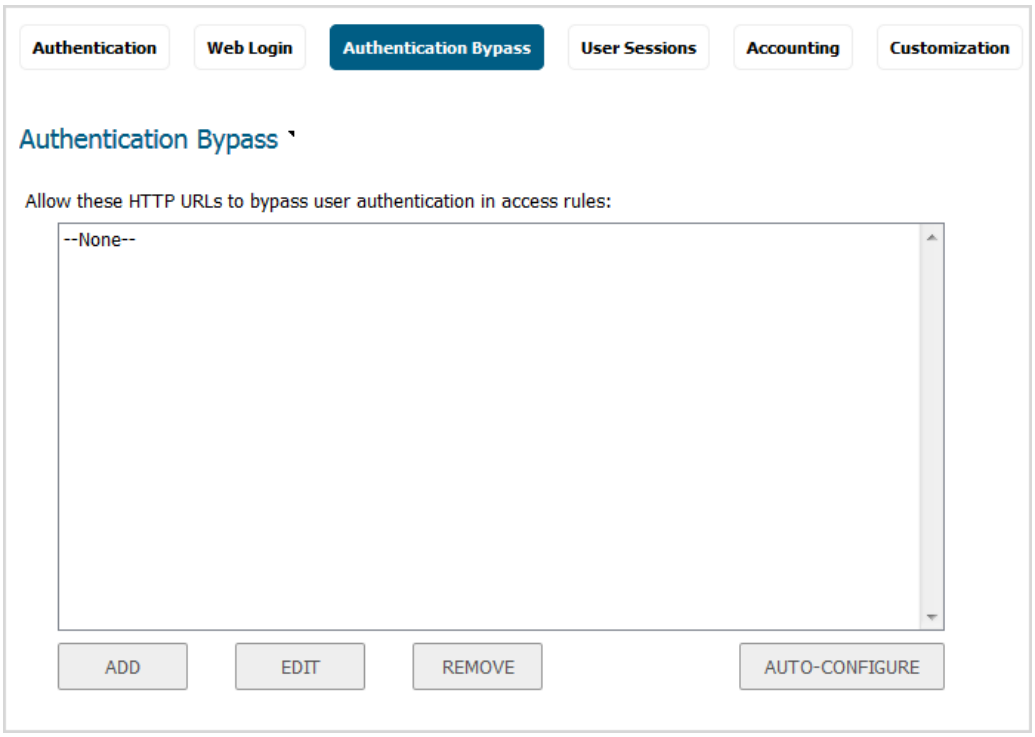
Topics:

- Adding URLs to Authentication Bypass on page 125
- Configuring Auto-configuration on page 125
- Converting URLs for Wildcard Matching on page 127
- Converting to Network(s) on page 127

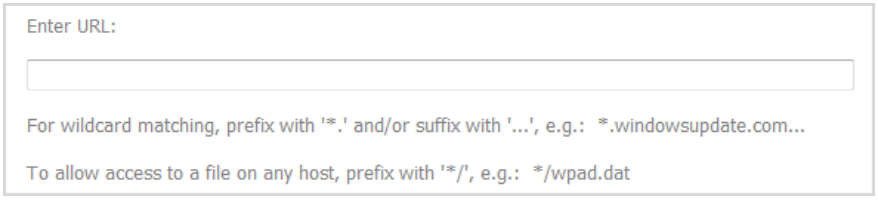
Adding URLs to Authentication Bypass

To add HTTP URLs user authentication bypass in Access Rules:

1 Navigate to **System Setup | Users > Settings > Authentication Bypass**.

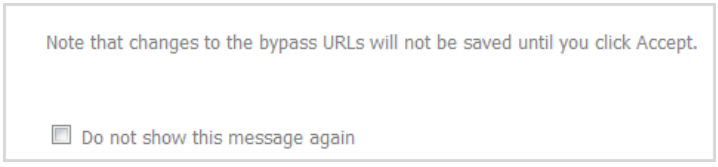


2 Click **ADD**. The **Add URL** popup displays.



3 Enter the URL in the **Enter URL** field.

4 Click **OK**. A popup confirmation message displays.



5 Click **OK**.

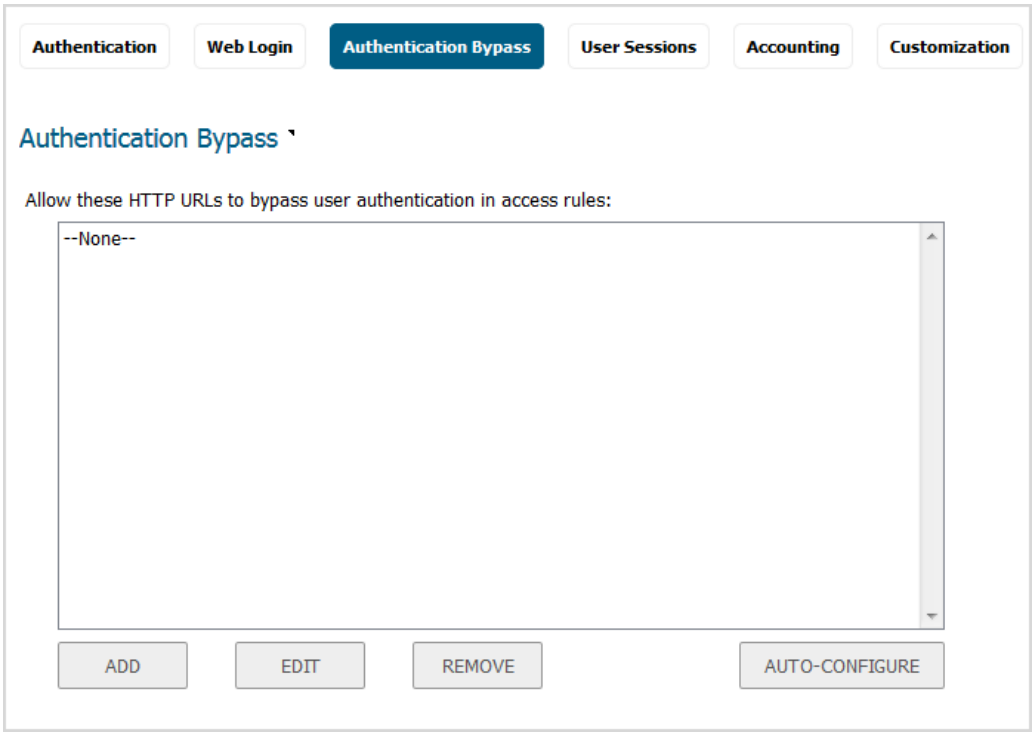
6 When finished adding URLs, click **ACCEPT**.

Configuring Auto-configuration

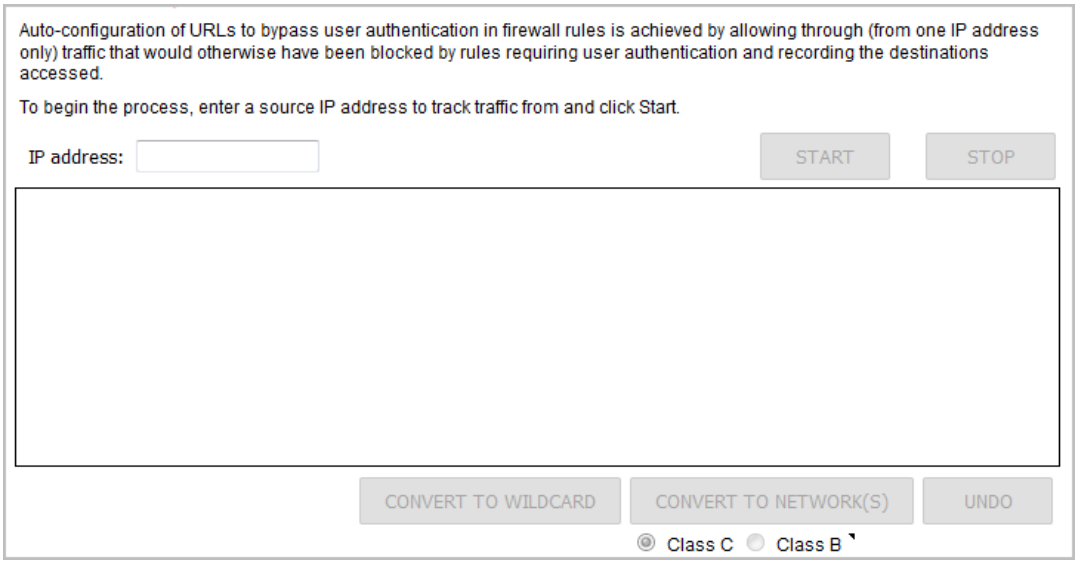
Auto-configuration of URLs to bypass user authentication in firewall rules is achieved by allowing through (from one IP address only) traffic that would otherwise have been blocked by rules requiring user authentication and recording the destinations accessed.

To configure Auto-configuration.

1 Navigate to **System Setup | Users > Settings > Authentication Bypass**.

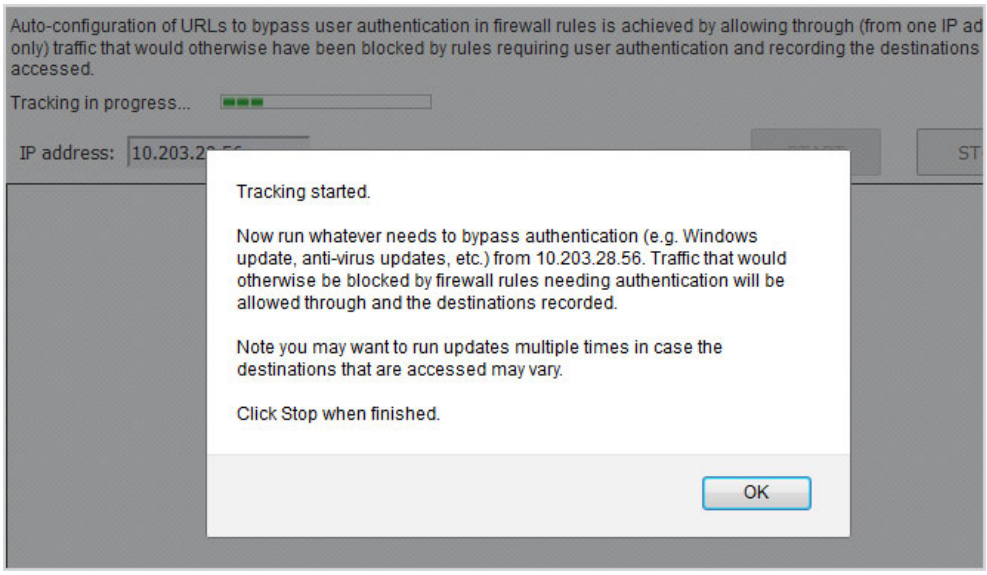


2 Click **AUTO-CONFIGURE**. The **Policy User Authentication Bypass Auto-Configuration** dialog displays.



3 Enter the source IP address in the **IP Address** field. **START** becomes available.

4 Click **START**. A **Tracking in progress** indicator and a **Tracking started** message displays.



5 Click **OK**.

Converting URLs for Wildcard Matching

Bypass authentication supports wildcard matching. This allows one or more tracked URLs to be converted to a single wildcard that matches against all currently selected URLs.

NOTE: The selected URLs must be in the same domain.

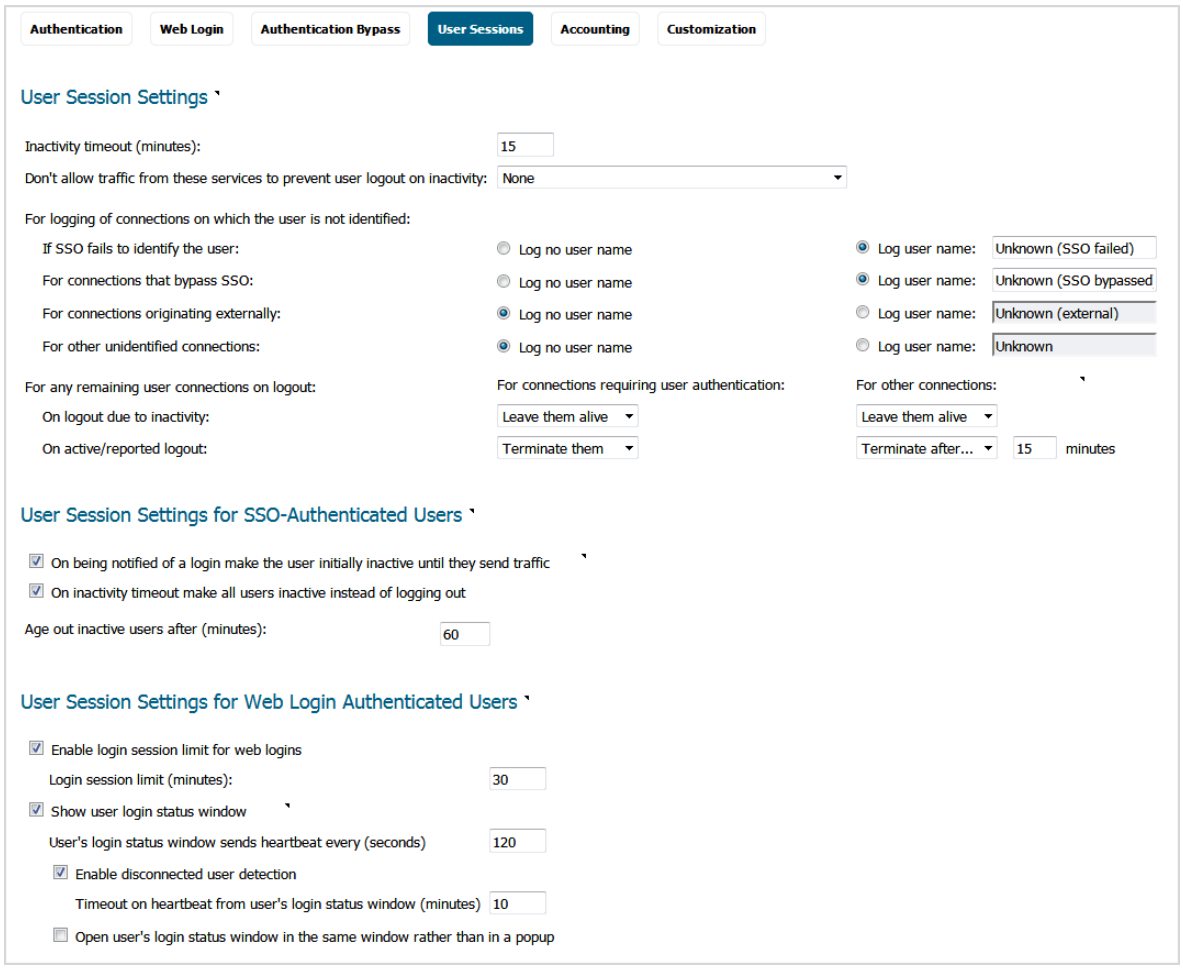
Converting to Network(s)

Windows Update accesses some destinations through HTTPS, and those destinations can be tracked only by IP address. The actual IP addresses accessed each time could vary, however, so rather than trying to set up a bypass for each such IP address, you can allow bypass for HTTPS to all IP addresses in that network.

Converting to network bypass allows tracked HTTPS destination IP address to be converted to either a:

- Class B (16 bit) network (default)
- Class C (24 bit) network

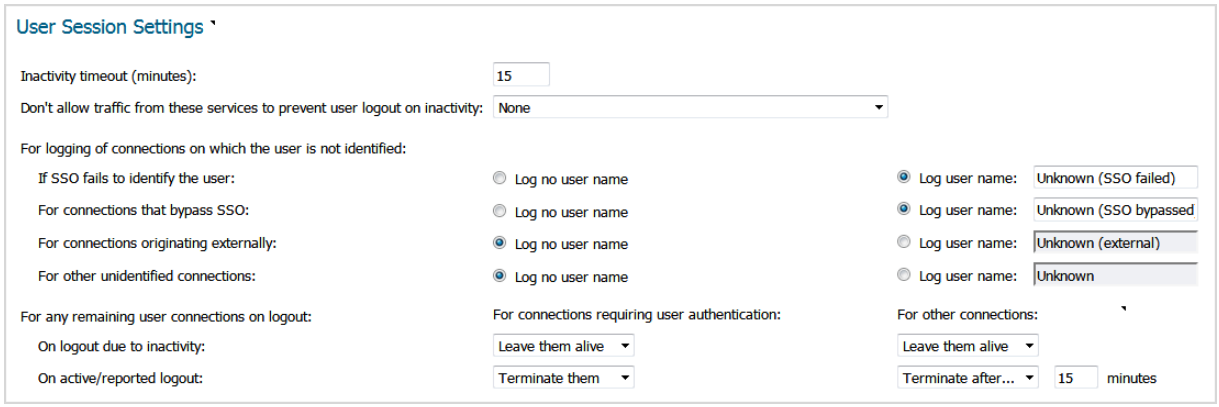
Configuring User Sessions



Topics:

- User Session Settings on page 129
- User Session Settings for SSO Authenticated Users on page 130
- User Session Settings for Web Login on page 131

User Session Settings



To configure settings that apply to all users who are authenticated through the Security Appliance:

1. Specify the length of time for inactivity after which users are logged out of the Security Appliance in the **Inactivity timeout (minutes)** field. The default is **15** minutes.
2. From the **Don't allow traffic from these services to prevent user logout on inactivity** drop-down menu, select the service or service group option to be prevented from logging out inactive users. This option saves system overhead and possible delays re-identifying aged-out authenticated users by making them inactive instead of logging them out. Inactive users do not use up system resources and can be displayed on the **Users > Status** page. The default is **None**.
3. For the following **For logging of connections on which the user is not identified** options, select the type of logging, **Log no user name** or **Log user name**, to be done, and optionally, the log user name:
   - **If SSO fails to identify the user: Log user name Unknown SSO failed** (default)
   - **For connections that bypass SSO: Log user name SSO Bypass** (default)

     NOTE: This option also can be set in the **SSO Bypass section of the Enforcement** tab of the **SSO Authentication Configuration** dialog.

   - **For connections originating externally: Log no user name** (default); if **Log user name** is selected, the default user name is **Unknown (external)**
   - **For other unidentified connects: Log no user name** (default); if **Log user name** is selected, the default user name is **Unknown**
4. Specify how to handle a user's connections that remain after the user logs out from the SonicWall appliance with the **Actions for remaining user connections on logout** options.

	Action	
Type of logout	**For connections requiring user authentication** [a]	**For other connections** [b]
On logout due to inactivity	Leave them alive (default) Terminate them Terminate after… minutes	Leave them alive (default) Terminate them Terminate after… minutes
On active/reported logout	Leave them alive Terminate them (default) Terminate after… minutes	Leave them alive Terminate them Terminate after… 15 minutes (default)

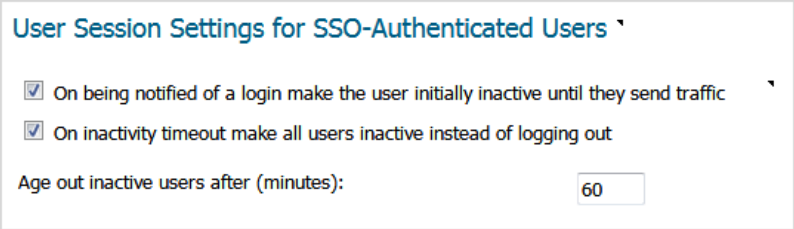
a. Applies for connections through access rules that allow only specific users.

b. Applies for other connections that do not have a specific user authentication requirement.

You can set different actions for:

- Inactivity logout, where the user might or might not still be logged into the domain/computer
- Users actively logging themselves out or being reported to the SonicWall appliance as being logged out (the latter normally means that the user has logged out from the domain/user)

User Session Settings for SSO Authenticated Users



To specify how inactive SSO-authenticated users are handled:

1 To put a user identified to the SonicWall appliance through an SSO mechanism, but no traffic has yet been received from the user, into an inactive state so they do not use resources, select the **On being notified of a login make the user initially inactive until they send traffic** checkbox. The users remain in an inactive state until traffic is received. This option is selected by default.

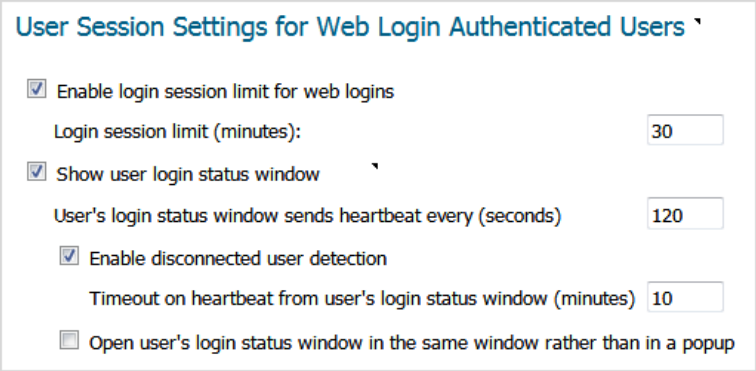
Some SSO mechanisms do not give any way for the SonicWall appliance to actively re-identify a user, and if users identified by such a mechanism do not send traffic, they remain in the inactive state until the appliance eventually receives a logout notification for the user. For other users who can be re-identified, if they stay inactive and do not send traffic, they are aged-out and removed after a period that can be set in Step 3.

2 If an SSO-identified user who has been actively logged in is timed out due to inactivity, then users who cannot be re-identified are returned to an inactive state. To have users who would otherwise be logged out on inactivity to be returned to an inactive state, select the **On inactivity timeout make all user inactive instead of logged out** checkbox. Doing this avoids overhead and possible delays re-identifying the users when they become active again. This setting is selected by default.

3 For inactive users who are subject to getting aged out, you can set the time, in minutes, after which they are aged-out and removed if they stay inactive and do not send traffic by selecting the **Age out inactive users after (minutes)** checkbox and specifying the timeout in the field. This setting is selected by default, and the minimum timeout value is 10 minutes, the maximum is 10000 minutes, and the default is **60** minutes.

NOTE: As the reason for keeping inactive user separate from active users is to minimize the resources used to manage them, the age-out timer runs once every 10 minutes. It might, therefore, take up to 10 minutes longer to remove inactive users from active status.

User Session Settings for Web Login



To configure user session settings for web login:

1 **Enable login session limit for web logins**: Limit the time a user is logged into the Security Appliance through web login by selecting the checkbox and typing the amount of time, in minutes, in the **Login session limit (minutes)** field. This setting is selected by default The default value is **30** minutes.

2 **Show user login status window** — For users logging in through web login, displays a status window with a **Log Out** button during the user's session. You can click **Log Out** to log out of your session.

> **NOTE:** The window must be kept open throughout the user's session as closing it logs the user out.

> **IMPORTANT:** If this option is not enabled, the status window is not displayed and users might not be able to log out. In this case, a login session limit must be set to ensure that they do eventually get logged out.

The **User Login Status** window displays the number of minutes the user has left in the login session. The user can set the remaining time to a smaller number of minutes by entering the number and clicking **Update**.

When this option is enabled, a mechanism that monitors heartbeats sent from that window also can be enabled to detect and log out users who disconnect without logging out.

If the user is a member of the SonicWall Administrators or Limited Administrators user group, the **User Login Status** window has a **Manage** button the user can click to automatically log into the Security Appliance's management interface. See Disabling the User Login Status Popup on page 113 for information about disabling the **User Login Status** window for administrative users. See Configuring Local Users and Groups on page 221 for group configuration procedures.

- **User's login status window sends heartbeat every (seconds)** — Sets the frequency of the heartbeat signal used to detect whether the user still has a valid connection. The minimum heartbeat frequency is 10 seconds, the maximum is 65530 seconds, and the default is **120** seconds.

3 **Enable disconnected user detection** — Causes the Security Appliance to detect when a user's connection is no longer valid and ends the session. This setting is selected by default.

- **Timeout on heartbeat from user's login status window (minutes)** — Sets the time needed without a reply from the heartbeat before ending the user session. The minimum delay before ending the user session is 1 minute, the maximum is 65535 minutes, and the default is **10** minutes.

4 Optionally, select to have the user's login status window display in the same window rather than a popup window by selecting **Open user's login status window in the same window rather than in a popup** checkbox.

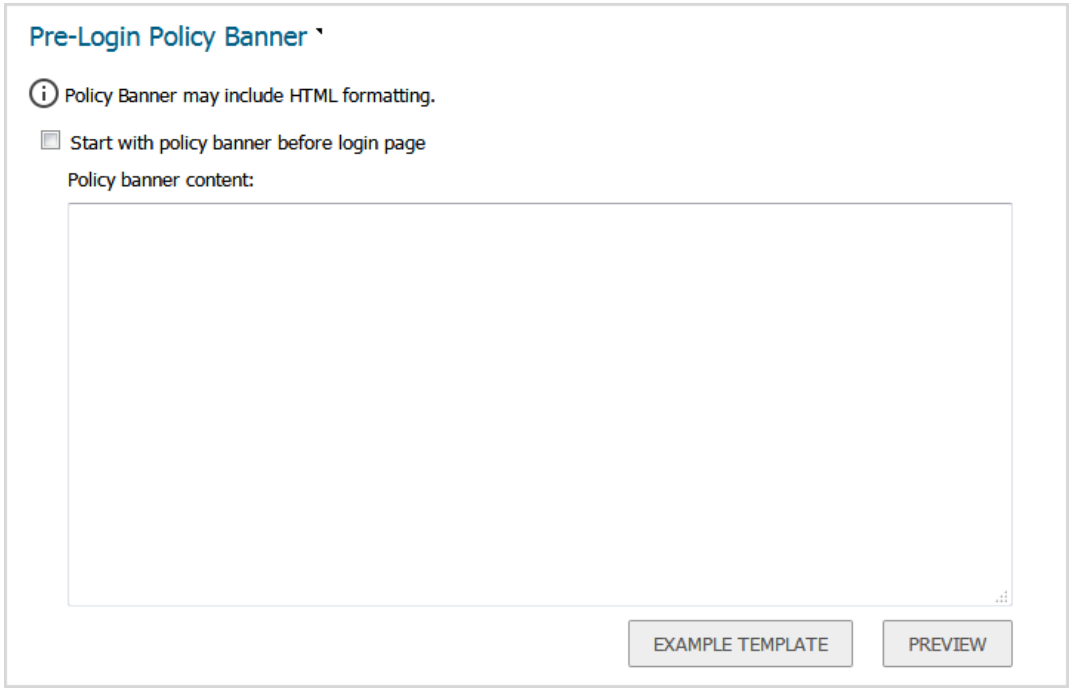
Customization

Topics:

- Pre-Login Policy Banner on page 132
- Post-Login Acceptable Use Policy on page 134
- Customized Login Pages on page 136

Pre-Login Policy Banner

In this section, you create a policy statement that is presented to all users as a banner in the window before web login. The policy banner might include HTML formatting.



To create a pre-login policy banner:

1. Navigate to **MANAGE | System Setup | Users > Settings**.
2. Click **Customization**.
3. Scroll to the **Pre-Login Policy Banner** section.
4. In the **Pre-Login Policy Banner** section, select **Start with policy banner before login page**. This option is not selected by default.
5. In the **Policy banner content** field, enter your policy text. You can include HTML formatting. The page that is displayed to the user includes **I Accept** and **Cancel** for user confirmation.

   TIP: Clicking **EXAMPLE TEMPLATE** creates a preformatted HTML template for your policy banner window; see Example Template on page 133.

6. Click **ACCEPT**.

Topics:

- Example Template on page 135
- Preview Message on page 133

Example Template

Click **EXAMPLE TEMPLATE** to populate the content with the default AUP template, which you can modify:

```
<font face=arial size=3>
<center><b><i>Welcome</center></b></i>
<font size=2>

<table width="100%" border="1">
<tr><td>
<font size=2>
<br><br><br>
<center>Enter your usage policy terms here.
<br><br><br>
</td></tr>
</table>

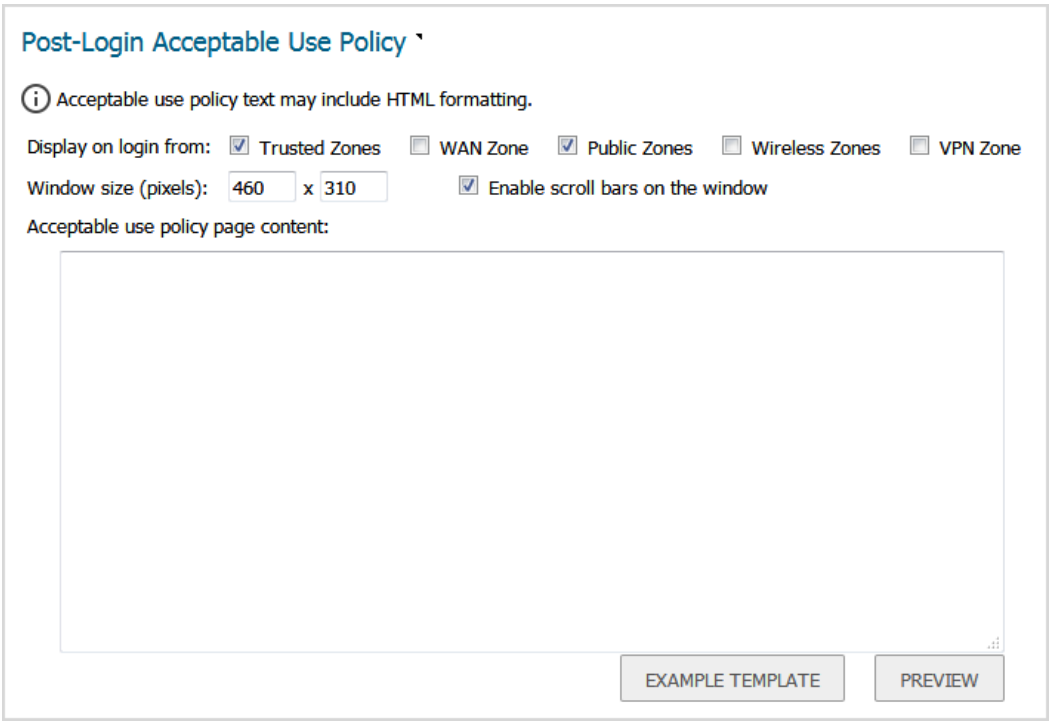
Click "I Accept" only if you wish to accept these terms and continue, or otherwise select
"Cancel".
```

Preview Message

Click **PREVIEW** to display your AUP message as it appears to the user.

Post-Login Acceptable Use Policy

An acceptable use policy (AUP) is a policy that users must agree to follow to access a network or the Internet. It is common practice for many businesses and educational facilities to require that employees or students agree to an acceptable use policy before accessing the network or Internet through the Security Appliance.



The **Post-Login Acceptable Use Policy** section allows you to create the AUP message window for users. You can use HTML formatting in the body of your message. Clicking **EXAMPLE TEMPLATE** creates a preformatted HTML template for your AUP window; see Example Template on page 135.

To create a post-login AUP message window:

1. Navigate to **MANAGE | System Setup | Users > Settings**.
2. Click **Customization**.
3. Scroll to the **Post-Login Acceptable Use Policy** section.
4. Specify these settings:
   - **Display on login from** - Select the network interface(s) you want to display the Acceptable Use Policy page when users login. You can choose **Trusted Zones** (default), **WAN Zone, Public Zones** (default), **Wireless Zones**, and **VPN Zone** in any combination.
   - **Window size (pixels)** - Allows you to specify the size of the AUP window, in pixels:
     - Width: Minimum size is 400 pixels, maximum size is 1280 pixels, and the default is **460** pixels.
     - Height: Minimum size is 200 pixels, maximum size is 1024 pixels, and the default is **310** pixels.
   - **Enable scroll bars on window** - Turns on the scroll bars if your content exceeds the display size of the window. This option is selected by default.
   - **Acceptable use policy page content** - Enter your Acceptable Use Policy text in this field. You can include HTML formatting. The page that is displayed to the user includes **I Accept** and **Cancel** for user confirmation.

5 Click **ACCEPT**.

Topics:

- Example Template on page 135
- Preview Message on page 135

Example Template

Click **EXAMPLE TEMPLATE** to populate the content with the default AUP template, which you can modify:

```
<font face=arial size=3>
<center><b><i>Welcome to the SonicWall</center></b></i>
<font size=2>

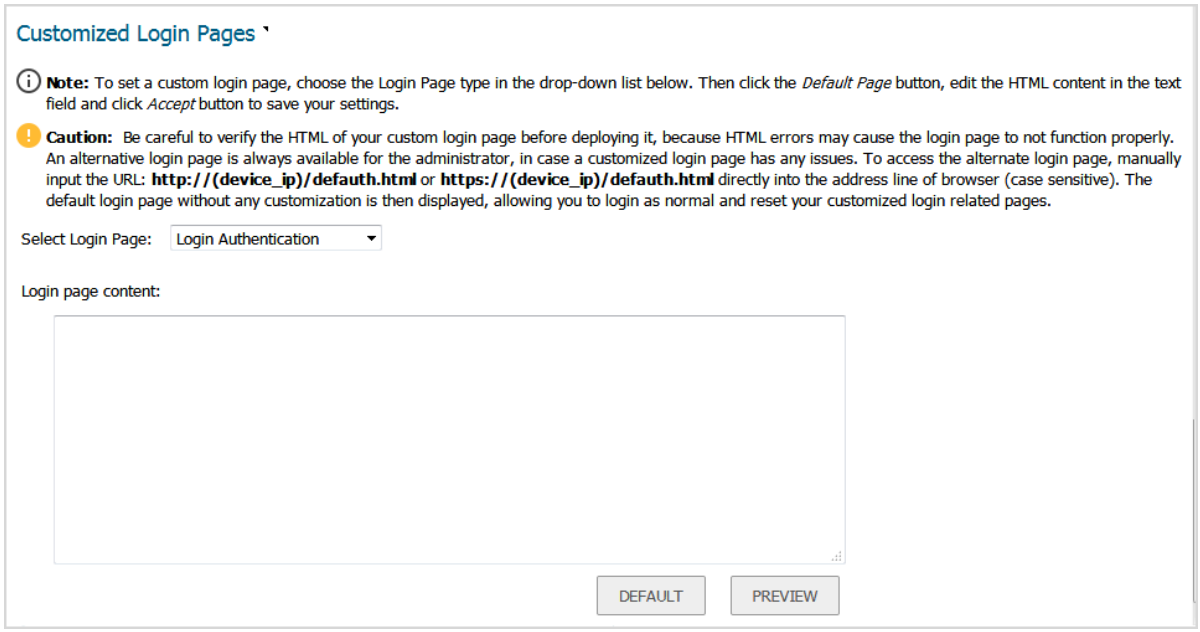
<table width="100%" border="1">
<tr><td>
<font size=2>
<br><br><br>
<center>Enter your usage policy terms here.
<br><br><br>
</td></tr>
</table>

Click "I Accept" only if you wish to accept these terms and continue,
or otherwise select "Cancel".
```

Preview Message

Click **PREVIEW** to display your AUP message as it appears to the user.

Customized Login Pages



SonicOS provides the ability to customize the text of the login authentication pages that are presented to users. You can translate the login-related pages with their own wording and apply the changes so that they take effect without rebooting.

Although the entire SonicOS Management Interface is available in different languages, sometimes you do not want to change the entire UI language to a specific local language.

However, if the Security Appliance requires authentication before users can access other networks, or enables external access services (for example, VPN, SSL-VPN), those login-related pages usually should be localized to make them more usable for typical users.

The **Customize Login Page** feature provides the following functionality:

- Keeps the style of original login by default
- Customizes login related pages
- Uses the default login related pages as templates
- Saves customized pages into system preferences
- Allows preview of changes before saving to preferences
- Presents customized login-related pages to typical users

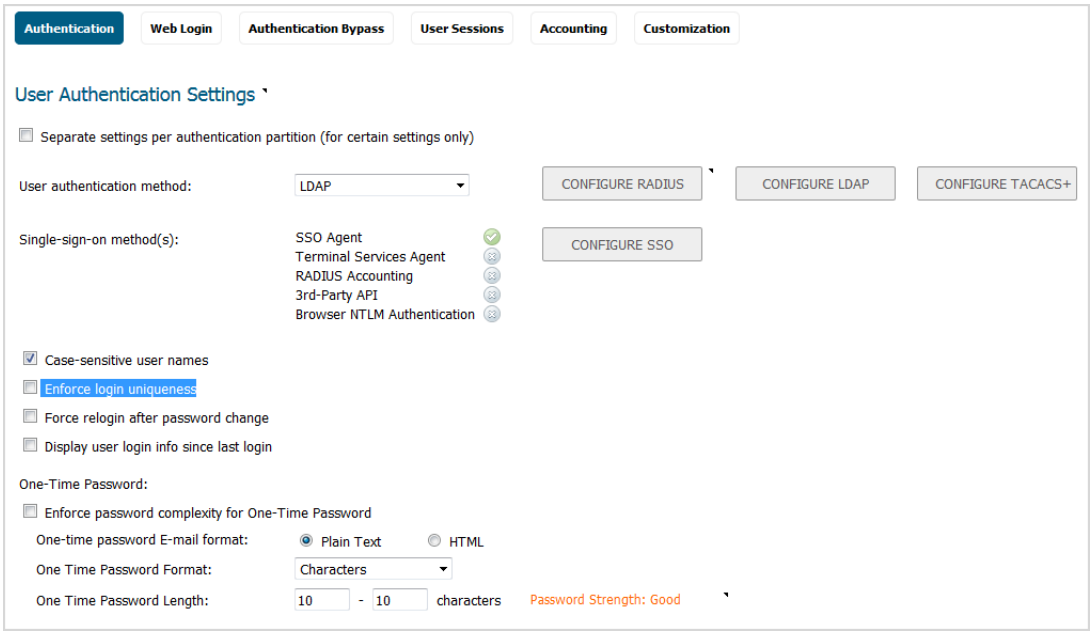
The following login-related pages can be customized:

- Admin Preempt
- Login Authentication
- Logged Out
- Login Full
- Login Disallowed
- Login Lockout
- Login Status
- Guest Login Status

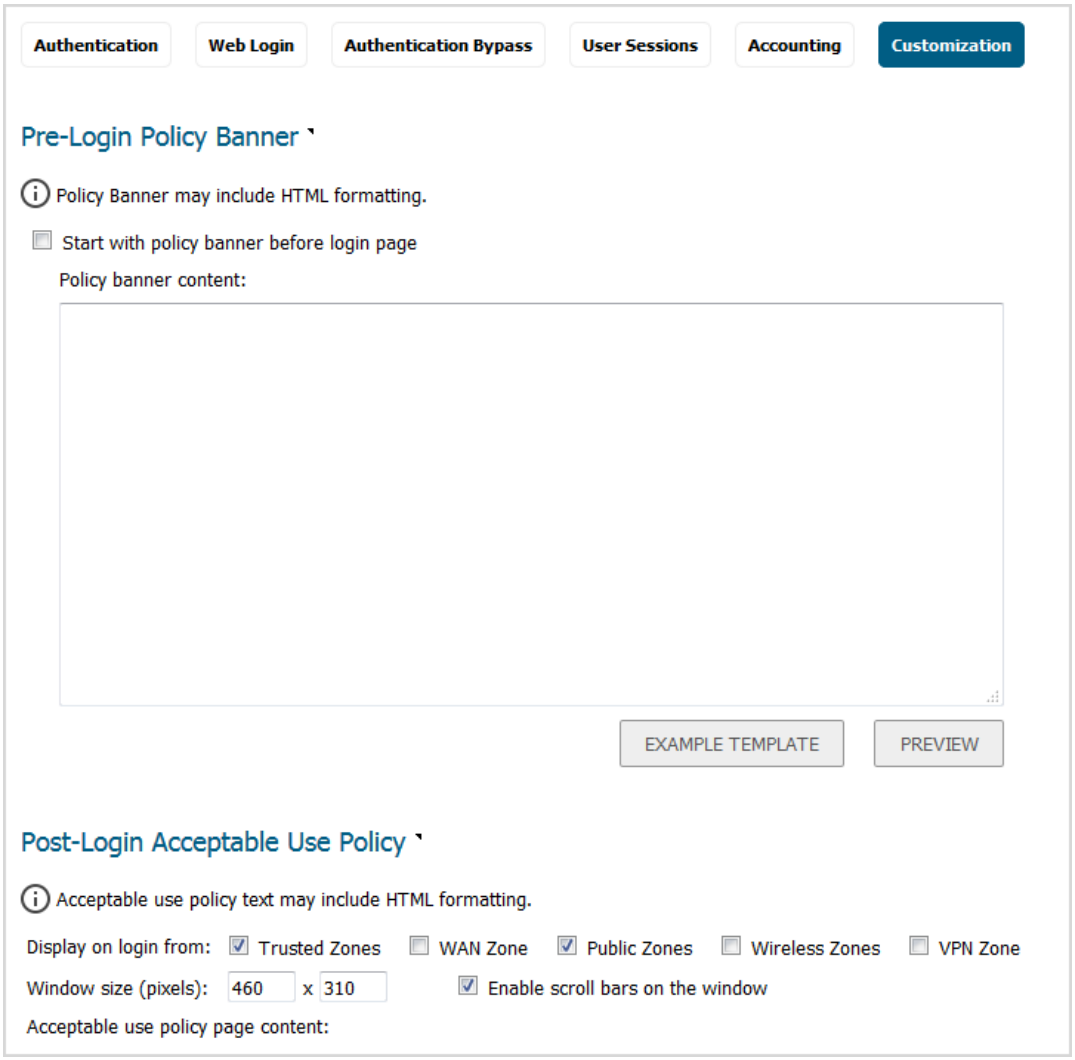
- Policy Access Barred
- Policy Access Down
- Policy Access Unavailable
- Policy Login Redirect
- Policy SSO Probe Failure
- User Password Update
- User Login Message

To customize one of these pages:

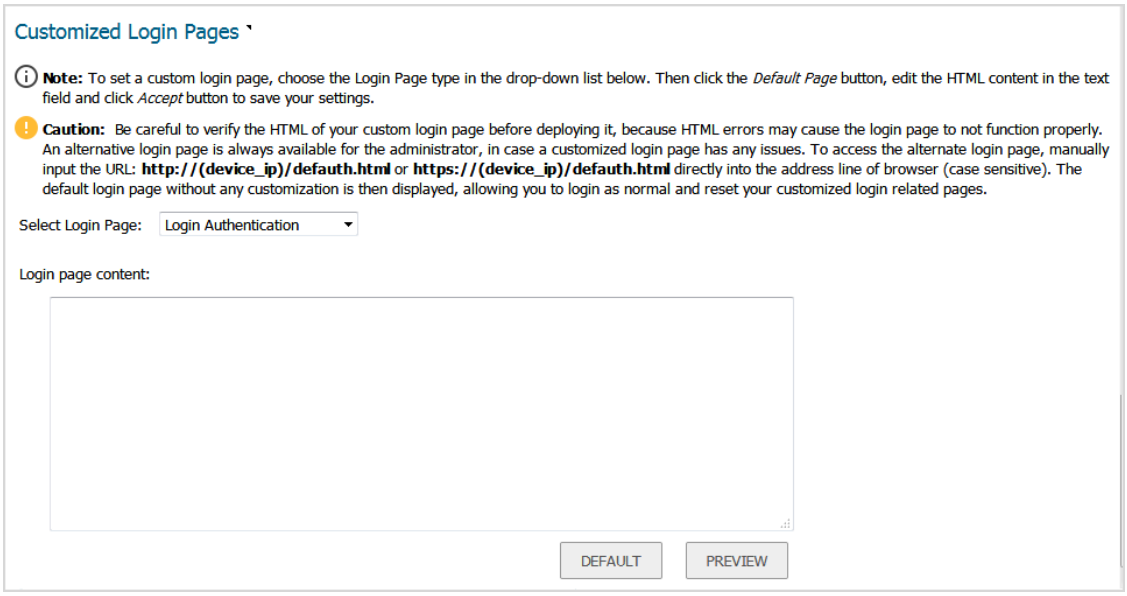
1 Navigate to **MANAGE | System Setup | Users > Settings**.



2 Click **Customization**.



3 Scroll to the **Customize Login Pages** section.



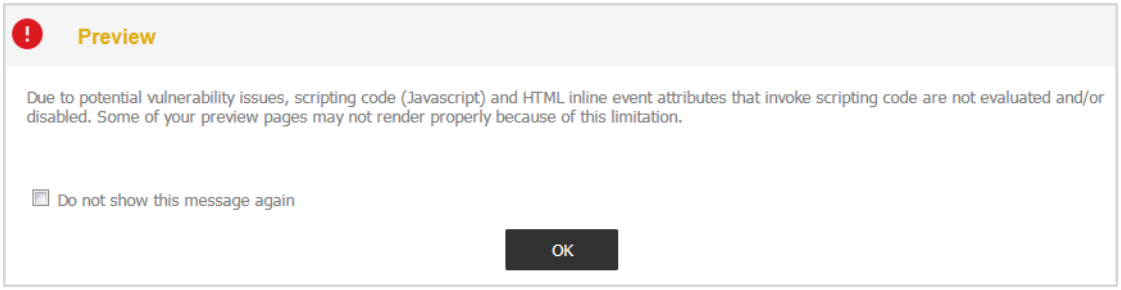
4 Select the page to be customized from **Select Login Page**.

5 Click **DEFAULT** to load the default content for the page.

6 Edit the content of the page.

NOTE: The `var strXXX =` lines in the template pages are customized JavaScript Strings. You can change them into your preferred wording. Modifications should follow the JavaScript syntax. You can also edit the wording in the HTML section.

7 Click **PREVIEW** to preview how the customized page looks. A message displays.



8 Click **OK**. Your customized page displays.

9 Close the window.

10 Make any changes.

11 When you are finished editing the page, click **ACCEPT**.

CAUTION: Be careful to verify the HTML of your custom login page before deploying it, because HTML errors might cause the login page to not function properly. An alternative login page is always available for the administrator, in case a customized login page has any issues. To access the alternate login page, manually input the URL: *https://(device_ip)/defauth.html* directly into the address line of browser (case sensitive). The default login page without any customization is then displayed, allowing you to login as normal and reset your customized login related pages.

TIP: Leave the **Login page content** field blank and apply the change to revert to the default page to users.

Accounting

SonicOS supports both RADIUS accounting and TACACS+ accounting. If both a RADIUS server and a TACACS+ server are configured, a user's accounting messages are sent to both servers.

Topics:

- Configuring TACACS+ Accounting on page 139

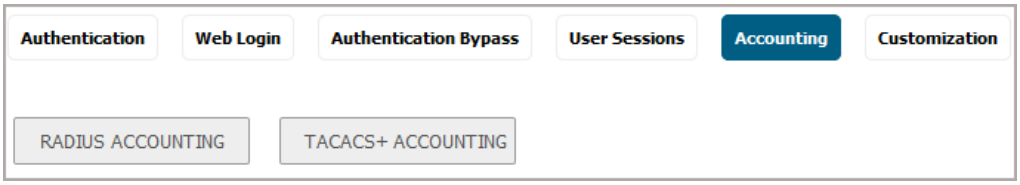
Configuring TACACS+ Accounting

SonicOS 6.5 supports TACACS+ accounting Start, Watchdog and Stop messages, but not the TACACS+ accounting proxy, that is, SonicOS does not forward the accounting request to the accounting server.

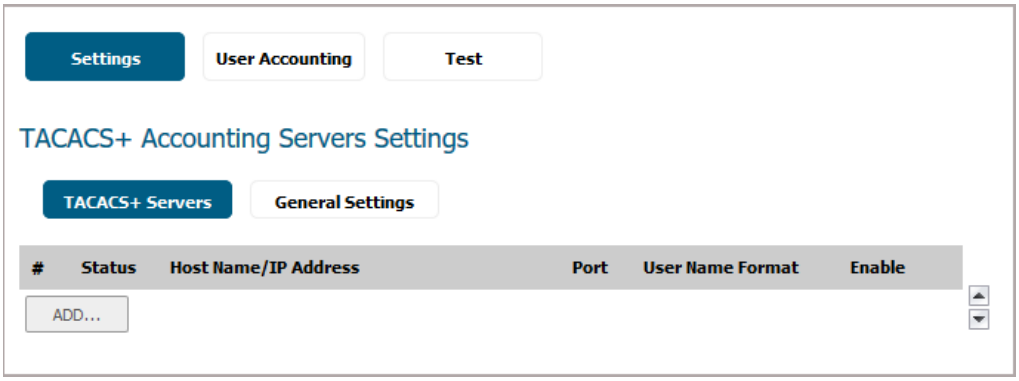
To configure TACACS+ accounting:

1 Navigate to **MANAGE | System Setup | Users > Settings**.

2 Click **Accounting**.

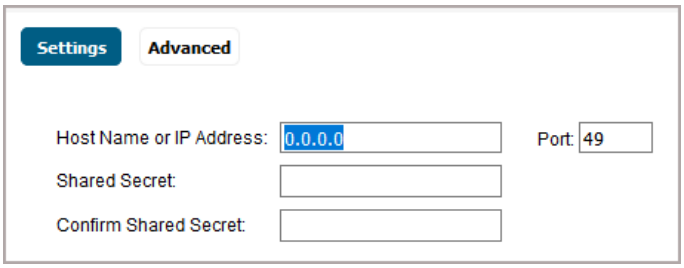


3 Click **TACACS+ ACCOUNTING**. The **TACACS+ Accounting Configuration** dialog displays.



4 To add a TACACS+ server, click **TACACS+ Servers**.

5 Click **Add**. The **Add TACACS+ accounting server** dialog displays.

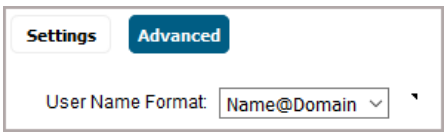


6 Enter the host name or IP address of the TACACS+ server in the **Host Name or IP Address** field.

7 Enter the port number of the server in the **Port** field. The default is **49**.

8 Enter the shared secret in the **Shared Secret** and **Confirm Shared Secret** fields.

9 Click **Advanced**.

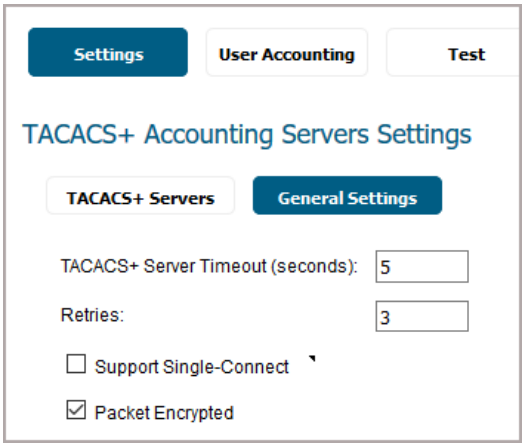


10 Select the format of the user name from **User Name Format**:

- **Simple-Name**
- **Name@Domain** (default)
- **Domain\Name**
- **Name.Domain**

11 Click **SAVE**.

12 Click **Settings**.



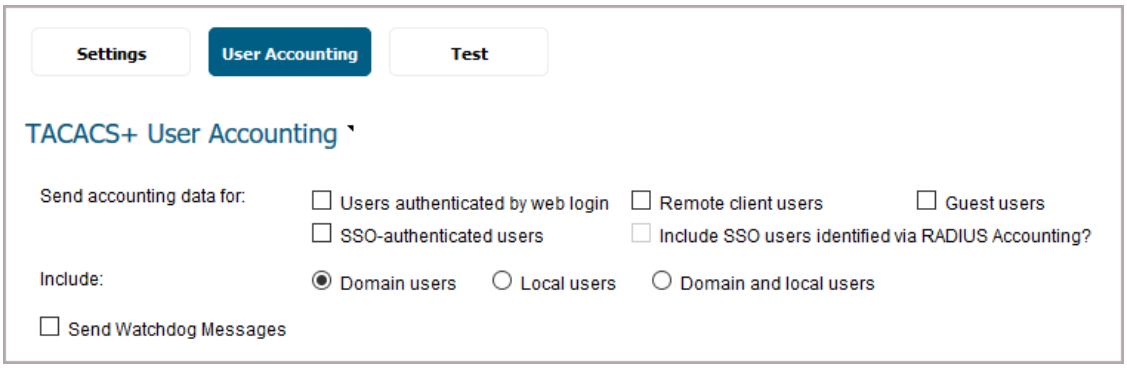
13 Enter the server timeout in the **TACACS+ Server Timeout (seconds)** field. The default is **5** seconds.

14 Enter the maximum number of retries in the **Retries** field. The default is **3**.

15 To support single connect, select **Support Single Connect**. This option is not selected by default.

16 To allow encrypted packets, select **Packet Encrypted**. This option is selected by default.

17 Click **User Accounting**.

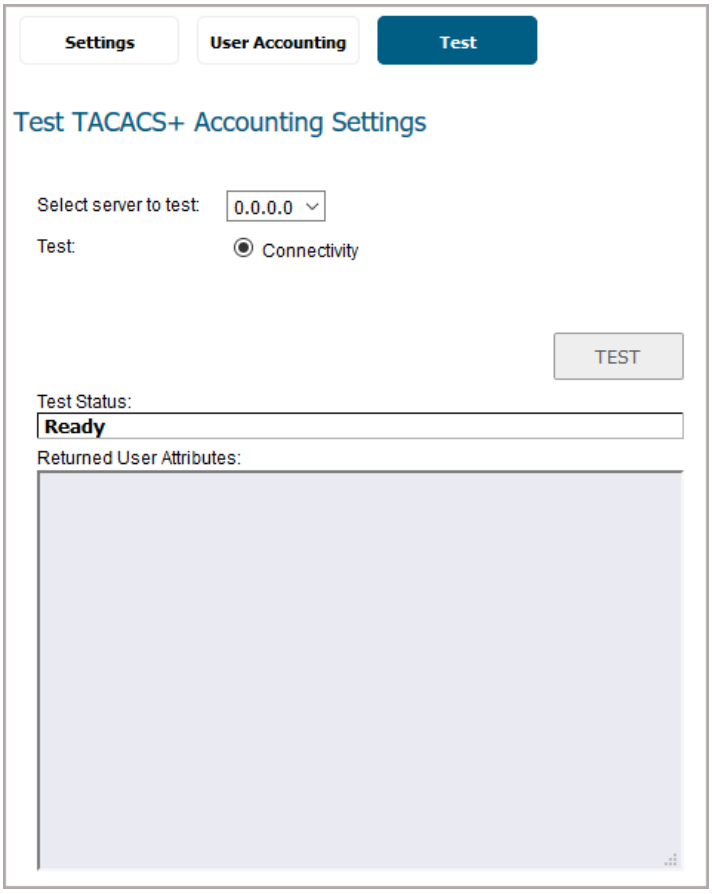


18 From **Send accounting data for**, select one or more types of users. **Include SSO users identified via RADIUS Accounting?** is not available by default. To make it selectable, first select the **SSO-authenticated users** field.

19 Choose whether to track domain and/or local users from **Include**. **Domain users** is selected by default.

20 To receive watchdog messages, select **Send Watchdog Messages**. This option is not selected by default. After selecting this option, the **Every:....minutes** option appears. Indicate how often you would like to receive **Watchdog Messages**.

21 Click **Test**.



22 From **Select server to test**, select the IP address of the TACACS+ server.

23 Choose the type of test from **Test**. **Connectivity** is selected by default.

24 Click **Test**. The results of the test display in **Returned User Attributes**.

25 Click **APPLY**.

26 Repeat the above steps for each server.

27 Click **OK**.

Configuring RADIUS Authentication

NOTE: For configuring RADIUS forSonicPoints or SonicWaves, see *SonicOS 6.5 Connectivity*.

For an introduction to RADIUS authentication in SonicOS, see Using RADIUS for Authentication on page 87. If you selected **RADIUS** or **RADIUS + Local Users** from the **Authentication method for login** drop-down menu on the **MANAGE | System Setup | Users > Settings** page, **CONFIGURE RADIUS** becomes available.

Topics:

- Configuring RADIUS Settings on page 143
- RADIUS Users Page on page 146

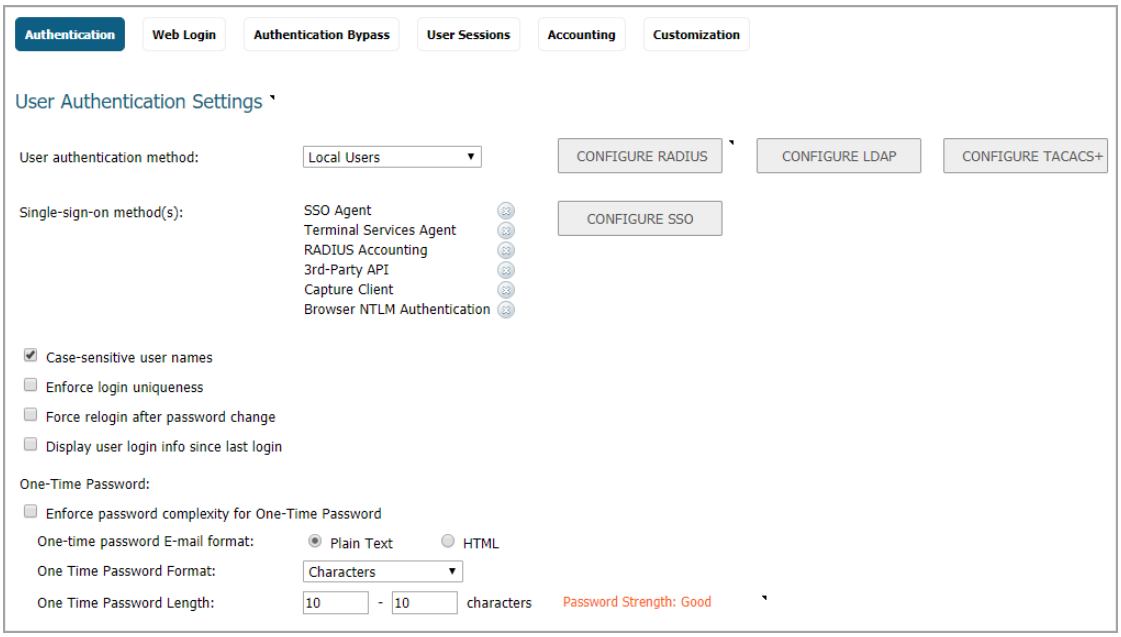
- RADIUS with LDAP for User Groups on page 147
- RADIUS Client Test on page 148

Configuring RADIUS Settings

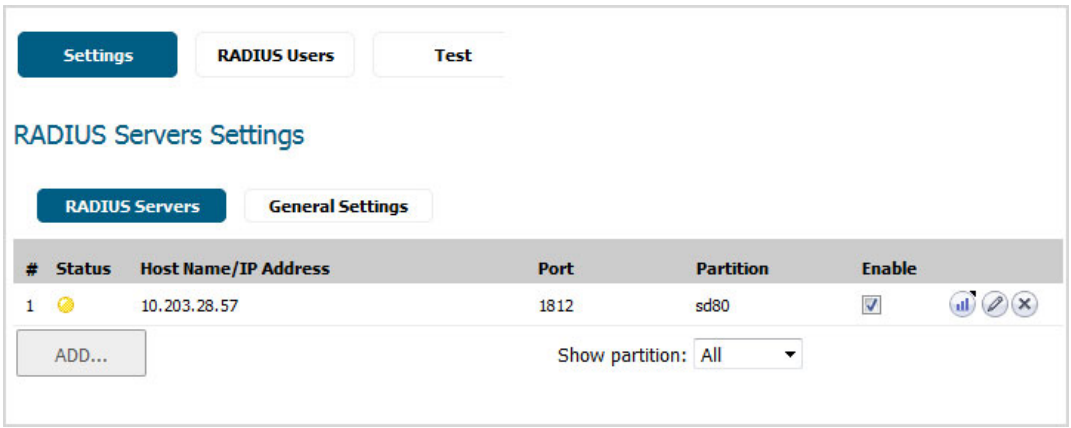
To configure RADIUS settings:

1 Navigate to **MANAGE | System Setup | Users > Settings**.

2 Click **Authentication**.



3 Click **CONFIGURE RADIUS**. The **RADIUS Configuration** dialog displays.



4 Under **Settings > RADIUS Servers Settings**, click **ADD**. The **Add server** dialog displays.



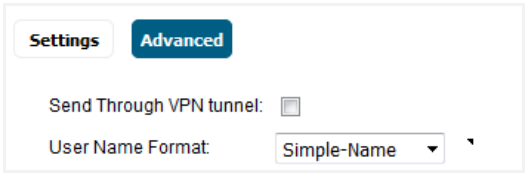
5 If partitioning is enabled, the **Select partition** popup displays over the **Add server** dialog. Select a partition and click **OK**.

6 Under **Settings**, enter an IP address or host name in the **Host Name or IP Address** field. The default is **0.0.0.0**.

7 In the **Port** field, enter the port for the RADIUS server to use for communication with SonicOS. The default is **1812**.

8 Enter the RADIUS server administrative password or shared secret in the **Shared Secret** and **Confirm Shared Secret** fields. The case-sensitive, alphanumeric **Shared Secret** can range from 1 to 31 characters in length.

9 Click **Advanced**.



10 Optionally, select **Send Through VPN tunnel**. This option is not selected by default.

11 Select the format for the user name from **User Name Format**:

- **Simple-Name** (default)
- **Name@Domain**
- **Domain\Name**
- **Name.Domain**

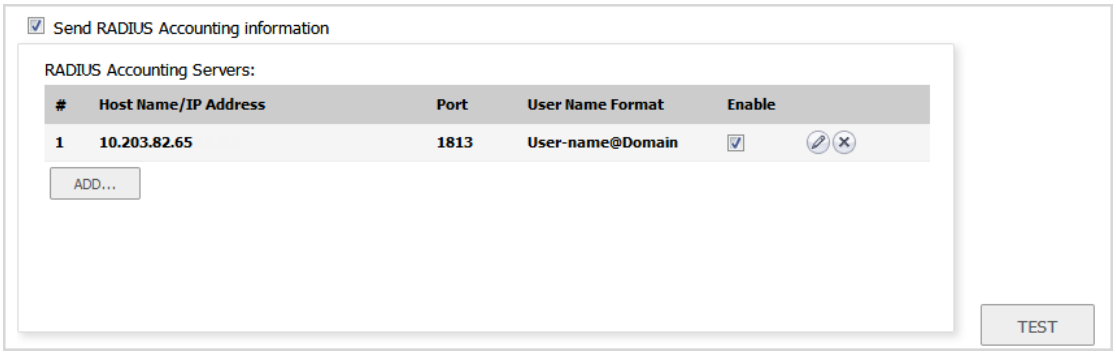
If the RADIUS server requires user names be sent with the domain component included, then select the format for that here.

NOTE: If the server accepts either the simple name without any domain component or a qualified name with the domain, then you can leave the selection as the default simple name unless you specifically want to force including the domain in the name sent to the server.

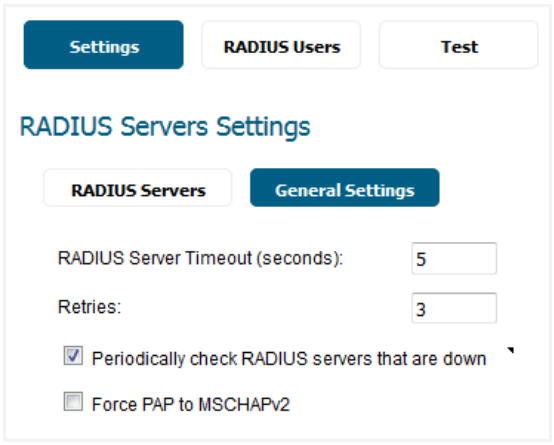
NOTE: In a Windows domain, if users are to be allowed to log in with a qualified user-name format that differs from what is set here (for example, to allow login to the firewall with `domain\name` if `name@domain` is selected or *vice versa*), then LDAP must be enabled for looking up the domain name mappings; otherwise, the user must enter a correctly formatted name acceptable by the RADIUS server.

12 Click **SAVE**.

13 Click **OK**. The server is added to the **RADIUS Accounting Servers** table.



14 Click **General Settings**.



15 Enter a timeout value in the **RADIUS Server Timeout (seconds)** field. The allowable range is 1-60 seconds with a default value of **5**.

16 In the **Retries** field, enter the number of times SonicOS attempts to contact the RADIUS server. If the RADIUS server does not respond within the specified number of retries, the connection is dropped. This field can range between 0 and 10, with a default setting of **3** RADIUS server retries.

17 To periodically check the status of RADIUS servers, select **Periodically check Radius server that are down**. This option is selected by default.

If the primary RADIUS server fails to respond to a request, then its status is changed to down (showing red on the **RADIUS servers** table of the **RADIUS Configuration** dialog and further authentication requests are sent to the secondary server until the primary comes back up. If this setting is checked, then while a a server is down, dummy authentication requests are periodically sent to check it. When the server responds to one, its status is restore to up. Your RADIUS server might log an occasional authentication request failure with user name, `status check`.

Disabling this option generally does not adversely affect user authentication. However, if it is disabled when the primary server goes down temporarily, then the firewall does not know when it becomes up, and so continues to show the server as down and sends authentication requests to the secondary server. This continues until secondary fails to respond to a request or the primary's status is checked manually, which can be done through the RADIUS test under **Test** on the **Configure RADIUS** dialog.

NOTE: If the secondary server goes down while the primary server is down, then the firewall reverts to sending requests to the primary server, so the firewall detects when/if the primary server is responsive regardless of this setting.

18 To Optionally, to enforce MS-CHAPv2 RADIUS authentication, select **Force PAP to MSCHAPv2.** This option is not selected by default.

19 Click **APPLY**.

20 Click **OK**.

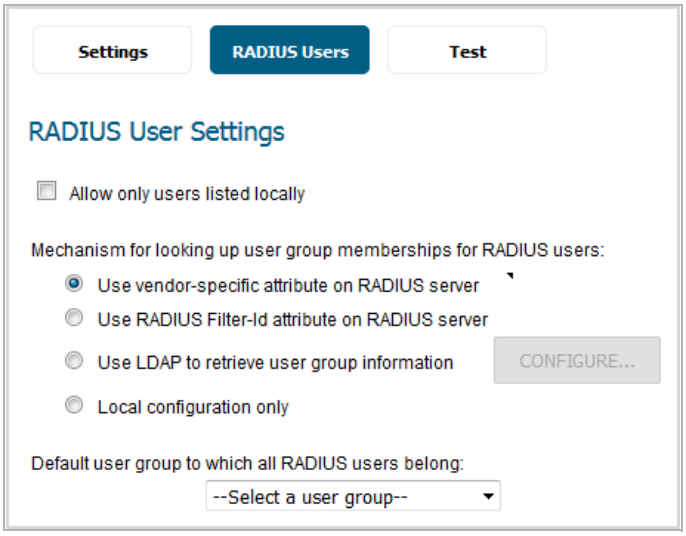
21 To configure RADIUS users, go to RADIUS Users Page on page 146

RADIUS Users Page

On the **RADIUS Users** page of the **RADIUS Configuration** dialog, you can specify what types of local or LDAP information to use in combination with RADIUS authentication. You can also define the default user group for RADIUS users.

To configure the RADIUS user settings:

1 Click **RADIUS Users**.



2 Select **Allow only users listed locally** if only the users listed in the SonicOS database are authenticated using RADIUS.

3 Select the **Mechanism used for setting user group memberships for RADIUS users** option:

> **NOTE:** If the **Use SonicWall vendor-specific attribute on Radius server** or **Use RADIUS Filter-ID attribute on RADIUS server** options are selected, the RADIUS server must be properly configured to return these attributes to the SonicWall appliance when a user is authenticated. The RADIUS server should return aero (0) or more instances of the selected attribute, each giving the name of a user group to which the user belongs.
>
> For details of the vendor-specific attribute settings, see the tech note, SonicOS Enhanced: Using User Level Authentication, and the SonicOS Enhanced RADIUS Dictionary file, `SonicWall.dct`. Both are available at https://support.SonicWall.com/.

- **Use SonicWall vendor-specific attribute on RADIUS server** – To apply a configured vendor-specific attribute from the RADIUS server. The attribute must provide the user group to which the user belongs. The preferred vendor-specific RADIUS attribute is `SonicWall-User-Group`. `SonicWall-User-Privilege` also works for certain user groups, but it is supported primarily for backwards compatibility and is not governed by the **Mechanism for setting user group memberships for RADIUS users** setting; that is, it is still effective even if something other than the **Use SonicWall vendor-specific attribute on RADIUS server** is selected.
- **Use RADIUS Filter-ID attribute on RADIUS server** – To apply a configured Filter-ID attribute from the RADIUS server. The attribute must provide the user group to which the user belongs.

- **Use LDAP to retrieve user group information** (default) – To obtain the user group from the LDAP server. You can click **Configure** to set up LDAP if you have not already configured it or if you need to make a change. For information about configuring LDAP, see Configuring the SonicWall for LDAP on page 149.
- **Local configuration only** – If you do not plan to retrieve user group information from RADIUS or LDAP.
- **Memberships can be set locally by duplicating RADIUS user names** – For a shortcut for managing RADIUS user groups. When you create users with the same name locally on the Security Appliance and manage their group memberships, the memberships in the RADIUS database automatically change to mirror your local changes.

4 If you have previously configured User Groups in SonicOS, select the group from the **Default user group to which all RADIUS users belong** drop-down menu. To create a new user group, see Creating a New User Group for RADIUS Users on page 147.

5 Either click

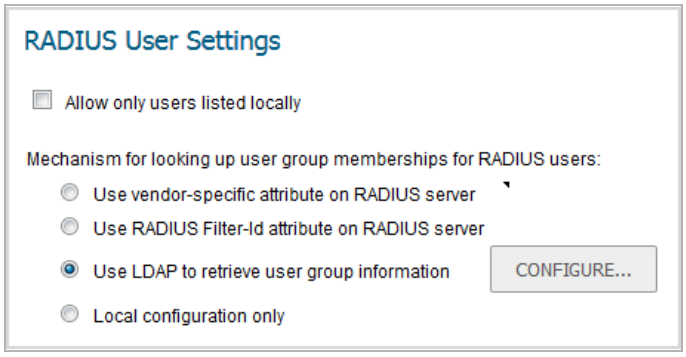
- **OK** if you have finished configuring the RADIUS server.
- **Apply**, to continue configuring RADIUS users and/or testing the settings.

Creating a New User Group for RADIUS Users

In the **RADIUS User Settings** dialog, you can create a new group by choosing **Create a new user group...** from the **Default user group to which all RADIUS users belong** drop-down menu. The **Add Group** dialog displays. For the procedure for creating a new user group, see Creating or Editing a Local Group on page 235.

RADIUS with LDAP for User Groups

When RADIUS is used for user authentication, there is an option on the **RADIUS Users** page in the **RADIUS Configuration** dialog to allow LDAP to be selected as the mechanism for setting user group memberships for RADIUS users:



When **Use LDAP to retrieve user group information** is selected, after authenticating a user through RADIUS, his/her user group membership information is looked up through LDAP in the directory on the LDAP/AD server.

NOTE: If this mechanism is **not** selected, and one-time password is enabled, a RADIUS user receives a one-time password fail message when attempting to login through SSL VPN.

Clicking **CONFIGURE** launches the **LDAP Configuration** dialog. For more information on configuring LDAP settings, see Preparing Your LDAP Server for Integration on page 90.

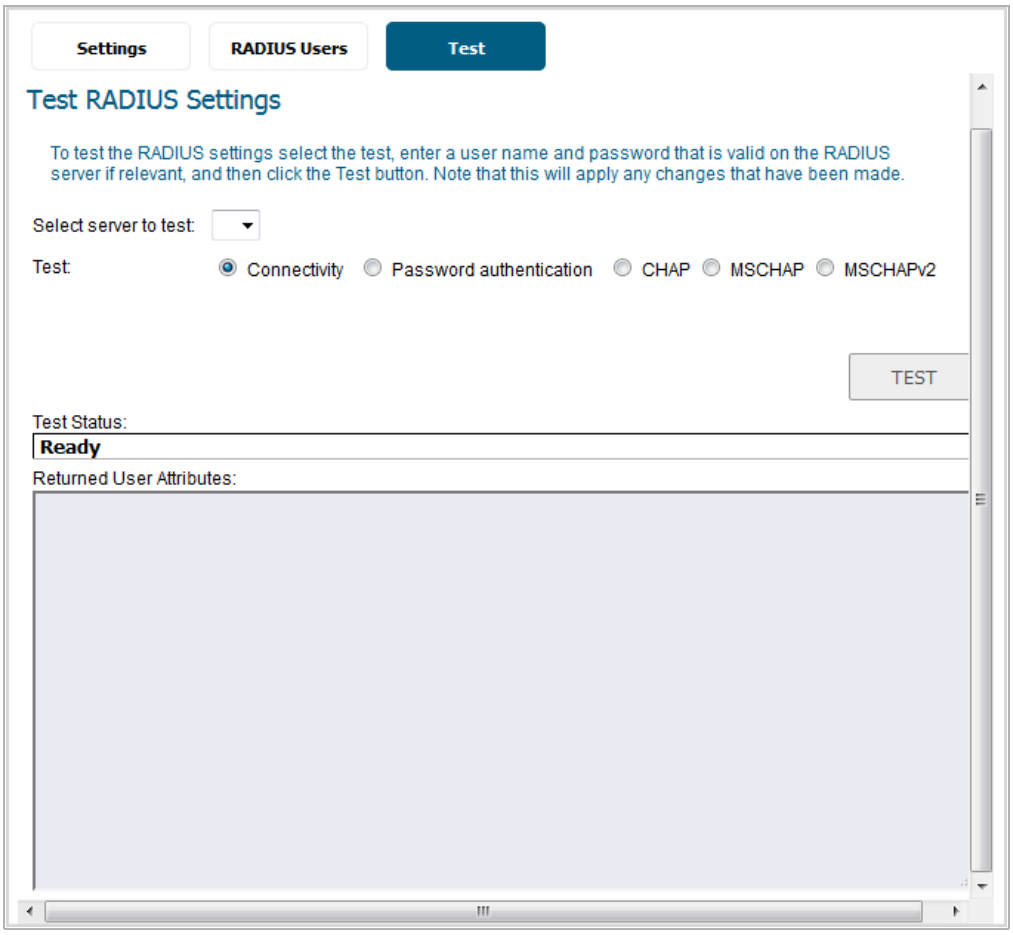
NOTE: In this case LDAP is not dealing with user passwords and the information that it reads from the directory is normally unrestricted, so operation without TLS could be selected, ignoring the warnings, if TLS is not available (for example, if certificate services are not installed with Active Directory). However, it must be ensured that security is not compromised by SonicOS doing a clear-text login to the LDAP server – for example, create a user account with read-only access to the directory dedicated for SonicOS use. Do not use the administrator account in this case.

RADIUS Client Test

In the **RADIUS Configuration** dialog, you can test your RADIUS Client user name, password and other settings by typing in a valid user name and password and selecting one of the authentication choices for **Test**. Performing the test applies any changes you have made.

To test your RADIUS settings:

1 Click **Test**.



2 In the **User** field, type a valid RADIUS login name.

3 In the **Password** field, type the password.

4 For **Test**, select one of the following:

- **Connectivity**: Select this to test RADIUS connectivity.
- **Password authentication**: Select this to use the password for authentication.

- **CHAP**: Select this to use the Challenge Handshake Authentication Protocol. After initial verification, CHAP periodically verifies the identity of the client by using a three-way handshake.
- **MSCHAP**: Select this to use the Microsoft implementation of CHAP. MSCHAP works for all Windows versions before Windows Vista.
- **MSCHAPv2**: Select this to use the Microsoft version 2 implementation of CHAP. MSCHAPv2 works for Windows 2000 and later versions of Windows.

5 Click **TEST**. If the validation is successful, the **Status** messages changes to **Success**. If the validation fails, the **Status** message changes to **Failure**.

6 To complete the RADIUS configuration, click **OK**.

After SonicOS has been configured, a VPN Security Association requiring RADIUS authentication prompts incoming VPN clients to enter a User Name and Password into a dialog.

Configuring the SonicWall for LDAP

Topics:

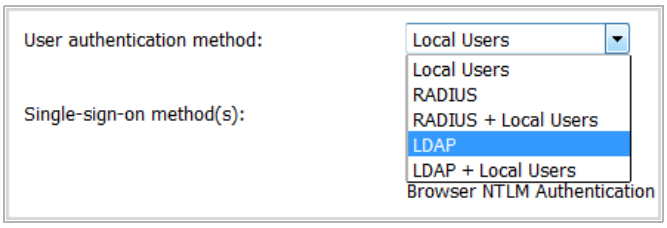
- Managing LDAP Integration on page 149
- About Extended Support for Multiple LDAP Servers on page 162
- About Importing and Mirroring from LDAP on page 163

Managing LDAP Integration

To manage your LDAP integration:

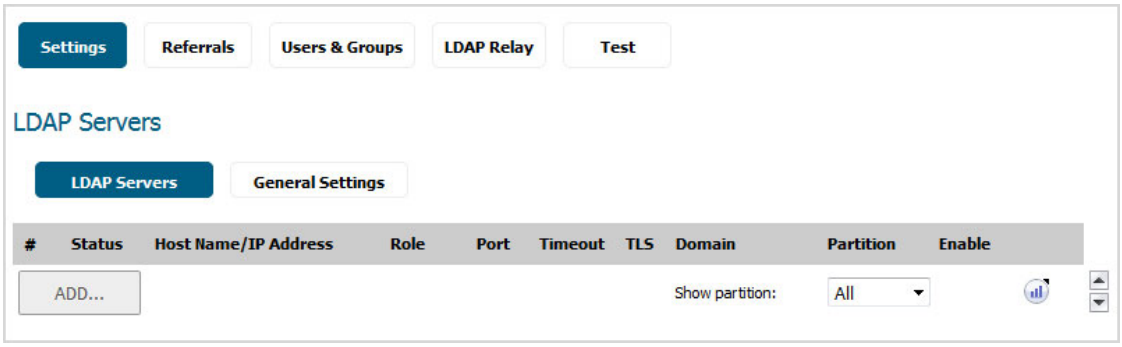
1 Navigate to **MANAGE | System Setup | Users > Settings**.

2 From **User Authentication method**, select either **LDAP** or **LDAP + Local Users**.



3 Click **CONFIGURE LDAP**.

4 If you are connected to your Security Appliance through HTTP rather than HTTPS, a message displays warning you of the sensitive nature of the information stored in directory services and offering to change your connection to HTTPS. If you have HTTPS management enabled for the interface to which you are connected (recommended), click **Yes**. The **LDAP Configuration** dialog displays.



NOTE: Dynamically learned secondary servers are shown in blue to distinguish them from configured servers.

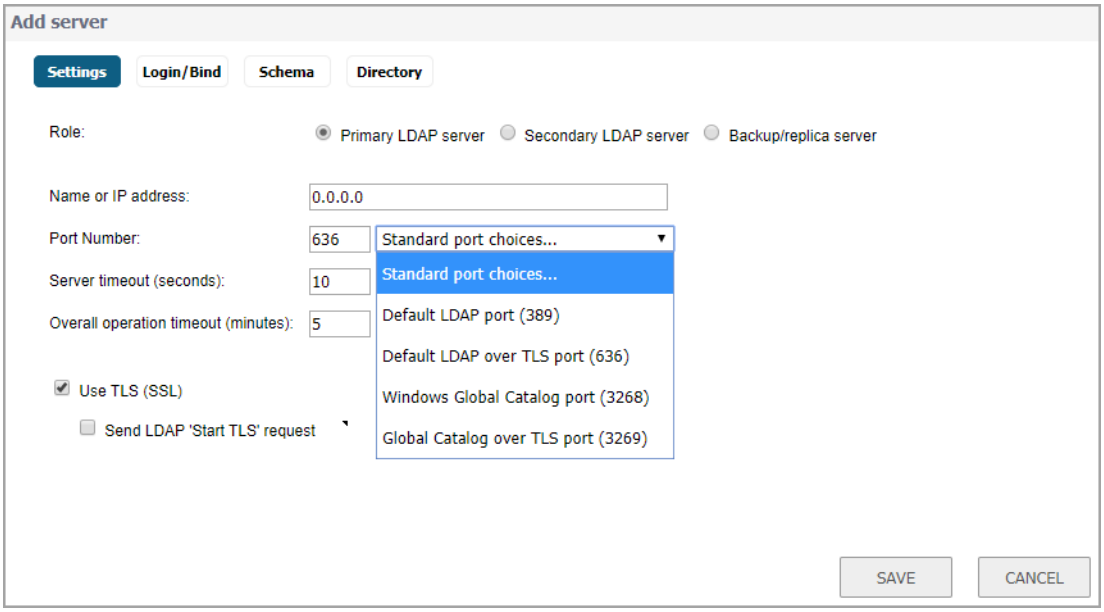
Topics:

- Settings on page 150
- Schema Page on page 152
- Directory Page on page 154
- Referrals Page on page 156
- Users & Groups Page on page 157
- LDAP Relay Page on page 160
- Test Page on page 161

Settings

To configure the LDAP server settings:

1 In the **CONFIGURE LDAP** interface, on the **Settings** page, click **ADD**. The **Add server** page displays.



2 Configure the following fields:

- **Name or IP Address** – The FQDN or the IP address of the LDAP server against which you wish to authenticate. If using a name, be certain that it can be resolved by your DNS server. Also, if using TLS with the 'Require valid certificate from server' option, the name provided here must match the name to which the server certificate was issued (such as the CN) or the TLS exchange fails.
- **Port Number** – The default LDAP over TLS port number is TCP 636. The default LDAP (unencrypted) port number is TCP 389. You can also select Windows Global Catalog port 3268 or Global Catalog over TLS port 3269. A Global Catalog is a Domain Controller that has data for every object in an Active Directory Forest rather than only the objects in one AD domain.

  If you are using a custom listening port on your LDAP server, specify it here.
- **Server timeout** – The amount of time, in seconds, that SonicOS waits for a response from the LDAP server before timing out. The range is 1 to 99999, with a default of **10** seconds.
- **Overall operation timeout** – The amount of time, in minutes, to spend on any automatic operation. Some operations, such as directory configuration or importing user groups, can take several minutes, especially when multiple LDAP servers are in use.
- Choose one of the following radio buttons:
  - **Anonymous Login** – Some LDAP servers allow for the tree to be accessed anonymously. If your server supports this (Active Directory generally does not), then you might select this option.
  - **Give login name/location in tree** – Select this option to build the distinguished name (dn) that is used to bind to the LDAP server from the `Login user name` and `User tree for login to server fields` according to the following rules:
    - The first name component begins `cn=`
    - The 'location in tree' components all use `ou=` (apart from certain Active Directory built-ins that begin with `cn=`)
    - The domain components all use `dc=`
    - If the User tree for login to server field is given as a `dn`, you can also select this option if the bind dn conforms to the first bullet above, but not to the second and/or the third bullet.
  - **Give bind distinguished name** – Select this option if the bind dn does not conform to the first bullet above (if the first name component does not begin with `cn=`). This option can always be selected if the dn is known. You must provide the bind dn explicitly if the bind dn does not conform to the first bullet above.
- **Login user name** – Specify a user name that has rights to log in to the LDAP directory. The login name is automatically presented to the LDAP server in full 'dn' notation. This can be any account with LDAP read privileges (essentially any user account); Administrative privileges are not required.

  NOTE: This is the user's name, not their login ID (for example, John Smith rather than jsmith).
- **Login password** – The password for the user account specified above.
- **Protocol version** – Select either LDAPv3 or LDAPv2. Most modern implementations of LDAP, including Active Directory, employ LDAPv3.
- **Use TLS (SSL)** – Use Transport Layer Security (SSL) to log in to the LDAP server. It is strongly recommended that TLS be used to protected the username and password information that is sent across the network. Most modern implementations of LDAP server, including Active Directory, support TLS. Deselecting this default setting displays an alert that you must accept to proceed.

- **Send LDAP 'Start TLS' Request** – Some LDAP server implementations support the Start TLS directive rather than using native LDAP over TLS. This allows the LDAP server to listen on one port (normally 389) for LDAP connections, and to switch to TLS as directed by the client. Active Directory does not use this option, and it should only be selected if required by your LDAP server.
- **Require valid certificate from server** – Validates the certificate presented by the server during the TLS exchange, matching the name specified above to the name on the certificate. Deselecting this default option presents an alert, but exchanges between SonicOS and the LDAP server still use TLS – only without issuance validation.
- **Local certificate for TLS** – Optional, to be used only if the LDAP server requires a client certificate for connections. Useful for LDAP server implementations that return passwords to ensure the identity of the LDAP client (Active Directory does not return passwords). This setting is not required for Active Directory.

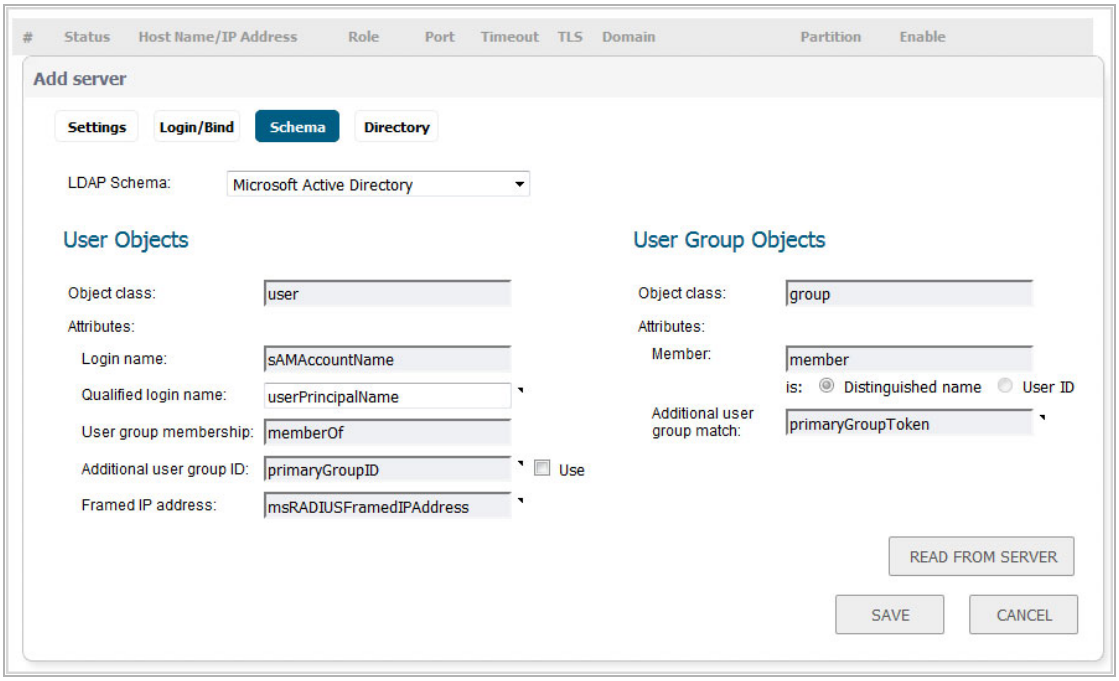
  If your network uses multiple LDAP/AD servers with referrals, then select one as the primary server (probably the one that holds the bulk of the users) and use the above settings for that server. It then refers SonicOS to the other servers for users in domains other than its own. For SonicOS to be able to log in to those other servers, each server must have a user configured with the same credentials (user name, password and location in the directory) as the login to the primary server. This might entail creating a special user in the directory for the SonicOS login. Note that only read access to the directory is required.
- **Force PAP to MSCHAPv2** – Optional, to enforce MS-CHAPv2 LDAP authentication select this option. If a RADIUS server is also configured, it provides authentication if LDAP authentication fails. This option is not selected by default.

3 Click **APPLY**.

Schema Page

To configure the LDAP server schema settings:

1 Click **Schema**.



2 **LDAP Schema** – Select one of the following from the **LDAP Schema** drop-down menu:

> **NOTE:** Selecting any of the predefined schemas automatically populates the fields used by that schema with their correct values. These values cannot be changed and their fields are dimmed.

- **Microsoft Active Directory** (default)
- **RFC2798 inetOrgPerson**
- **RFC2307 Network Information Service**
- **Samba SMB**
- **Novell eDirectory**
- **User defined** – Allows you to specify your own values; use this only if you have a specific or proprietary LDAP schema configuration.

3 **Object class** – Select the attribute that represents the individual user account to which the next two fields apply.

4 **Login name attribute** – Select one of the following to define the attribute that is used for login authentication:

- **sAMAccountName** for Microsoft Active Directory
- **inetOrgPerson** for RFC2798 inetOrgPerson
- **posixAccount** for RFC2307 Network Information Service
- **sambaSAMAccount** for Samba SMB
- **inetOrgPerson** for Novell eDirectory

5 **Qualified login name attribute** – Optionally, select an attribute of a user object that sets an alternative login name for the user in `name@domain` format. This might be needed with multiple domains in particular, where the simple login name might not be unique across domains.

> **NOTE:** For **Microsoft Active Directory**, this is normally set to **userPrinicpalName** for log in using `name@domain`, but could be set to **mail** to enable log in by email address. For **RFC2798 inetOrgPerson**, it is set to **mail**.

6 **User group membership attribute** – Select the attribute that contains information about the groups to which the user object belongs. This is **memberOf** in Microsoft Active Directory. The other predefined schemas store group membership information in the group object rather than the user object, and therefore do not use this field.

7 **Framed IP address attribute** – Select the attribute that can be used to retrieve a static IP address that is assigned to a user in the directory. Currently it is only used for a user connecting through L2TP with the SonicOS L2TP server. In the future this might also be supported for Global VPN Client. In Active Directory the static IP address is configured on the Dial-in tab of a user's properties.

8 **User Group Objects** – This section is auto-configured unless you select **User Defined** for the **LDAP Schema**.

- **Object class** – Specify the name associated with the group of attributes.
- **Member attribute** – Specify the attribute associated with a member.
  - Select whether this attribute is a **Distinguished name** or **User ID**.
- **Read from server** – Click to read the user group object information from the LDAP server.

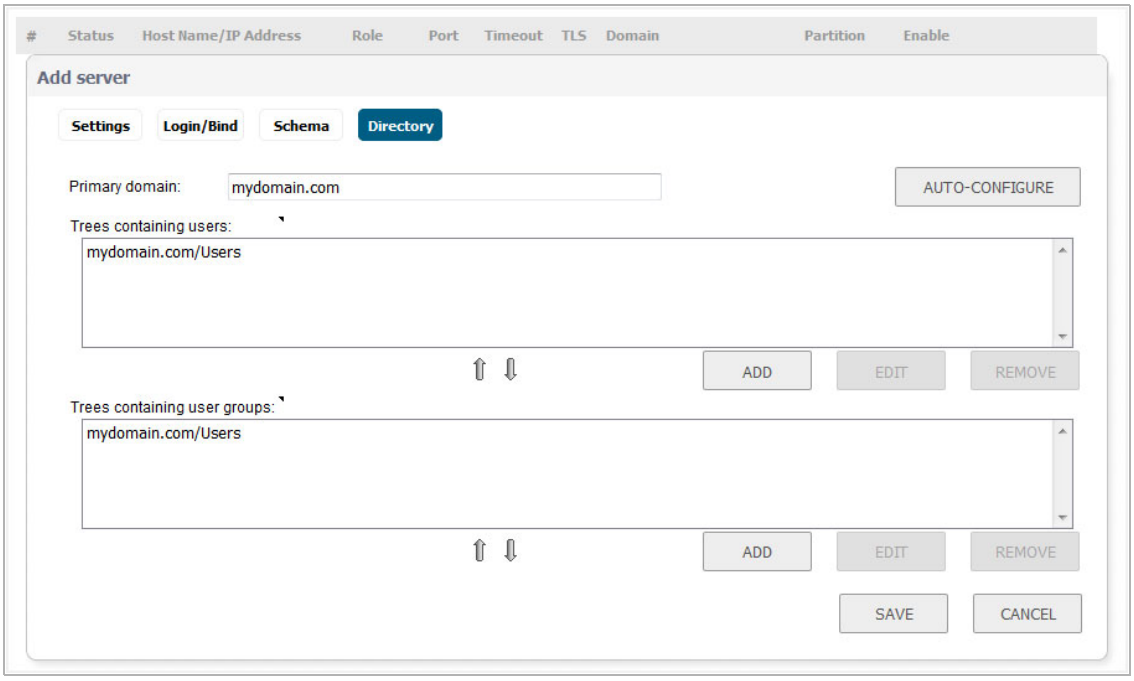
  > **NOTE:** You must enter the primary domain on the **Directory** tab first.

  - Select whether you want to **Automatically update the schema configuration** or **Export details of the schema**.

Directory Page

To configure the LDAP server directory settings:

1 On the **Directory** page, configure the following fields:



- **Primary Domain** – The user domain used by your LDAP implementation. For AD, this is the Active Directory domain name, for example, *yourADdomain.com*. Changes to this field, optionally, automatically update the tree information in the rest of the page. This is set to **mydomain.com** by default for all schemas except Novell eDirectory, for which it is set to **o=mydomain**.
- **Trees containing users** – The trees where users commonly reside in the LDAP directory. One default value is provided which can be edited, and up to a total of 64 DN values could be provided. SonicOS searches the directory using them all until a match is found, or the list is exhausted. If you have created other user containers within your LDAP or AD directory, you should specify them here.
- **Trees containing user groups** – Same as above, only with regard to user group containers, and a maximum of 32 DN values might be provided. These are only applicable when there is no user group membership attribute in the schema's user object, and are not used with AD.
- All the above trees are normally given in URL format but can alternatively be specified as distinguished names (for example, `myDom.com/Sales/Users` could alternatively be given as the DN `ou=Users,ou=Sales,dc=myDom,dc=com`). The latter form is necessary when the DN does not conform to the normal formatting rules as per that example. In Active Directory the URL corresponding to the distinguished name for a tree is displayed on the Object tab in the properties of the container at the top of the tree.

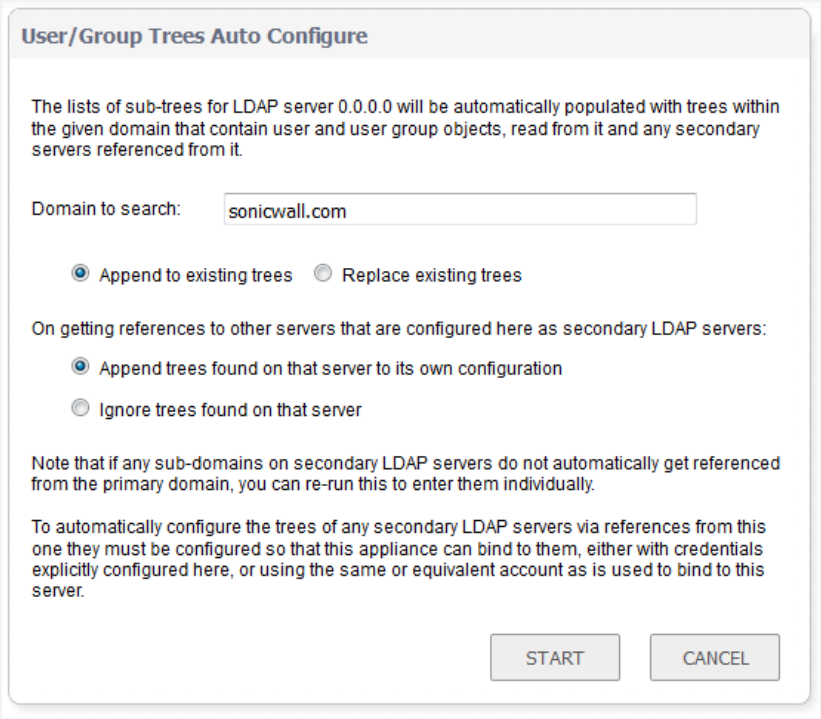
  NOTE: AD has some built-in containers that do not conform (for example, the DN for the top level Users container is formatted as `cn=Users,dc=…`, using `cn` rather than `ou`) but SonicOS knows about and deals with these, so they can be entered in the simpler URL format.

  Ordering is not critical, but because they are searched in the given order it is most efficient to place the most commonly used trees first in each list. If referrals between multiple LDAP servers

are to be used, then the trees are best ordered with those on the primary server first, and the rest in the same order that they are referred.

NOTE: When working with AD, to determine the location of a user in the directory for the **User tree for login to server** field, the directory can be searched manually from the Active Directory Users and Settings control panel applet on the server, or a directory search utility such as queryad.vbs in the Windows NT/2000/XP Resource Kit can be run from any PC in the domain.

- **AUTO-CONFIGURE** – This causes SonicOS to auto-configure the **Trees containing users** and **Trees containing user groups** fields by scanning through the directory/directories looking for all trees that contain user objects. To use auto-configure, first enter a value in the **User tree for login to server** field (unless anonymous login is set), and then click **AUTO-CONFIGURE** to display the **User/Group Trees Auto Configure** dialog:



  a) Enter the desired domain in the **Domain to search** field.

  b) Choose one of the following:

  - **Append to existing trees** – This selection appends newly located trees to the current configuration. This is the default.
  - **Replace existing trees** – This selection starts from scratch removing all currently configured trees first.

  c) Choose one of the following for **On getting references to other servers that are configured here as secondary LDAP servers**:

  - **Append trees found on that server to its own configuration** – This selection appends newly located trees found on the server to its own configuration. This is the default.
  - **Ignore trees found on that server** – This selection ignores trees on the specified server.

2 Click **OK**.

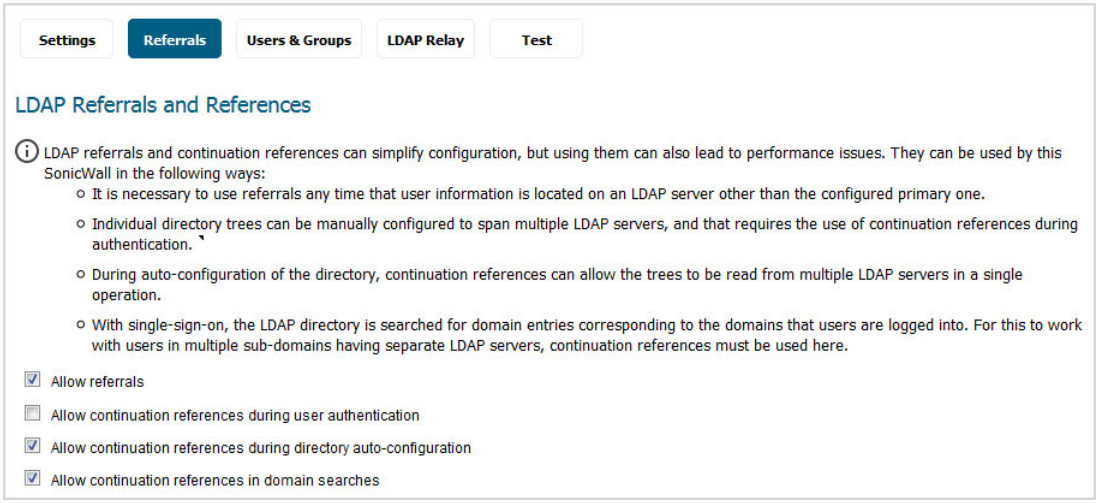
The auto-configuration process might also locate trees that are not needed for user login. You can manually remove these entries.

If using multiple LDAP/AD servers with referrals, this process can be repeated for each, replacing the **Domain to search** value accordingly and selecting **Append to existing trees** on each subsequent run.

Referrals Page

To configure the LDAP server referrals settings:

1 Click **Referrals**.



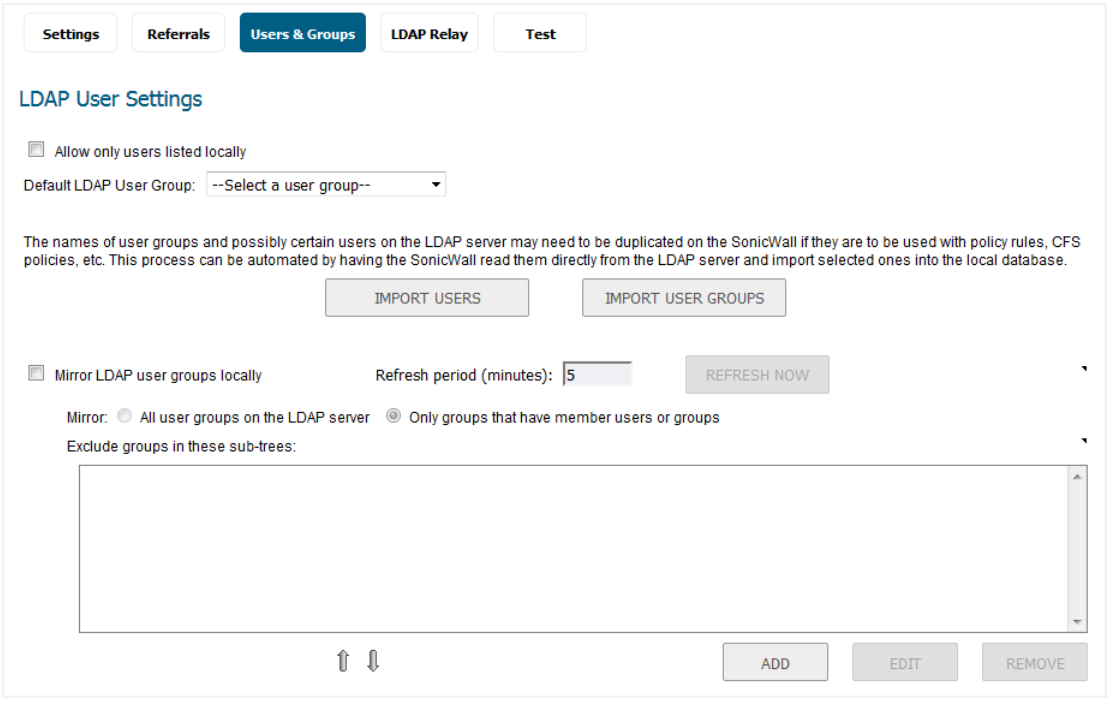
2 Configure the following fields:

- **Allow referrals** – Select this option any time that user information is located on an LDAP server other than the configured primary one. This option is selected by default.
- **Allow continuation references during user authentication** – Select this option any time that individual directory trees have been manually configured to span multiple LDAP servers. This option is not selected by default.
- **Allow continuation references during directory auto-configuration** – Select this option to allow the trees to be read from multiple LDAP servers in a single operation. This option is selected by default.
- **Allow continuation references in domain searches** – Select this option when using single-sign-on with users in multiple sub-domains having separate LDAP servers. This option is selected by default.

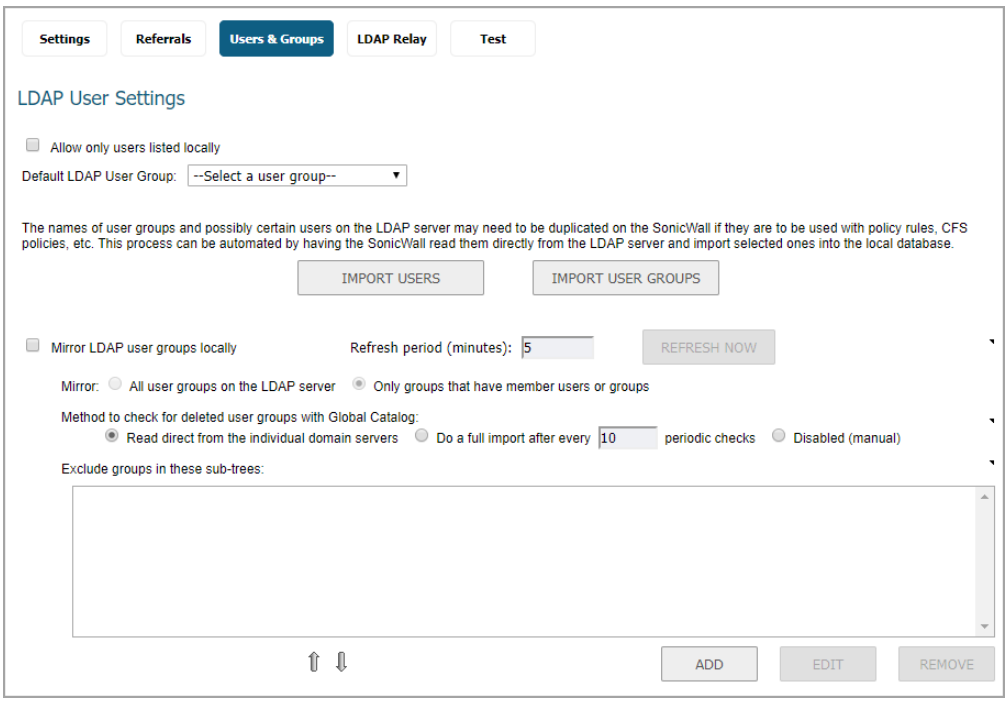
Users & Groups Page

To configure the LDAP users and groups settings:

1 Click **Users & Groups**.

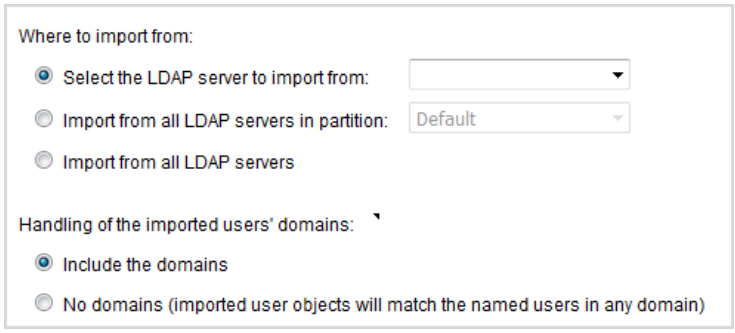


In SonicOS 6.5.4.5 and higher, if you have configured at least one LDAP server that uses Global Catalog, additional options are displayed on the Users & Groups page.



2 Configure the following fields:

- **Allow only users listed locally** – Requires that LDAP users also be present in the SonicOS local user database for logins to be allowed.
- **Default LDAP User Group** – A default group in SonicOS to which LDAP users belong in addition to group memberships configured on the LDAP server.
- **IMPORT USERS** – You can click this button to configure local users in SonicOS by retrieving the user names from your LDAP server. **IMPORT USERS** launches the **Select Server–User** dialog.



a) Choose the server from which to import users:
   - Select an LDAP server from **Select the LDAP server to import from**.
   - Select a partition containing the servers from **Import from all LDAP servers in partition**.
   - To import users from all LDAP servers, choose **Import from all LDAP servers**.

b) Choose whether to include domains in imported user objects:
   - **Include the domains**
   - **No domains (imported user objects will match the named users in any domain)**

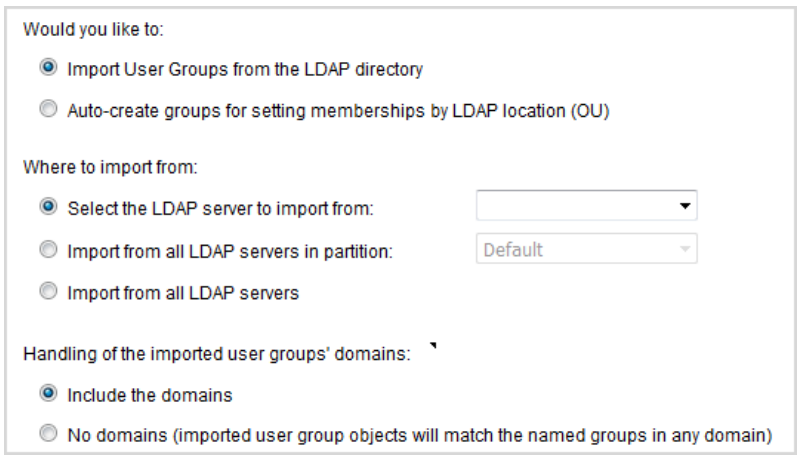
   Including domains when importing users allows for distinguishing users from different domains who might have the same user name. Including domains is generally recommended, especially if there are multiple domains. Settings, such as local user group memberships and account expiry time, made for an imported user object that includes the domain are applied only when the named user is authenticated from that domain.

   If domains are not included, then each imported user object matches any user with that name from any domain, and any settings made for such an imported user object is applied for all those matching users regardless of their domains.

c) Click **OK**.

Having users in SonicOS with the same name as existing LDAP users allows SonicWall user privileges to be granted upon successful LDAP authentication.

- **IMPORT USER GROUPS** – You can click this button to configure SonicOS user groups by retrieving the user group names from your LDAP server. **IMPORT USER GROUPS** launches the **Select Server–User Groups** dialog.



  a) Choose how to import users:

     - **Import User Groups from the LDAP directory**
     - **Auto-create groups for setting memberships by LDAP location (OU)**

     Having user groups in SonicOS with the same name as existing LDAP/AD user groups allows SonicWall group memberships and privileges to be granted upon successful LDAP authentication.

     Alternatively, you can manually create user groups on the LDAP/AD server with the same names as SonicWall built-in groups (such as 'Guest Services', 'Content Filtering Bypass', 'Limited Administrators') and assign users to these groups in the directory. This also allows SonicWall group memberships to be granted upon successful LDAP authentication.

  b) Choose the server from which to import users:

     - Select an LDAP server from **Select the LDAP server to import from**.
     - Select a partition containing the servers from **Import from all LDAP servers in partition**.
     - To import users from all LDAP servers, choose **Import from all LDAP servers**.

  c) Choose whether to include domains in imported user objects:

     - **Include the domains**
     - **No domains (imported user objects will match the named users in any domain)**

  The Security Appliance can retrieve group memberships efficiently in the case of Active Directory by taking advantage of its unique trait of returning a 'memberOf' attribute for a user.

- **Mirror LDAP user groups locally** – When enabled, the firewall periodically auto-imports user groups and nested user groups from the LDAP server(s) to create local user groups that mirror those in the LDAP directory.

  Configure the following related options:

  - **Refresh period (minutes)** – The number of minutes between periodic auto-import of changes in user groups from the LDAP server.

- **Mirror** – Select either **All user groups on the LDAP server** or **Only groups that have member users or groups**. Select the latter if the LDAP server has many user groups and you want to avoid exceeding the limit that can be imported.

- If you are mirroring user groups locally from an LDAP server using Global Catalog, deleted user groups on the server are not detected during the periodic auto-imports since Global Catalog does not support doing that. To handle deleted user groups, select one of the following options under **Method to check for deleted user groups with Global Catalog:**

  NOTE: These settings only appear if you have configured at least one LDAP server that uses Global Catalog.

  - **Read direct from the individual domain servers** – Search for deleted user groups on each domain server. This requires that the domain servers can be accessed using the same credentials as configured for the Global Catalog LDAP server. This is the default selection.
  - **Do a full import after every [X] periodic checks** – Occasionally perform a full read of all user groups on the Global Catalog server. Doing a full read too often is inefficient. The default is to do this after every 10 periodic checks.
  - **Disabled (manual)** – Select this option for no automatic handling of deleted user groups. You can trigger a full import by doing a manual refresh from the web management interface or CLI to clean up any deleted user groups that are locally mirrored.

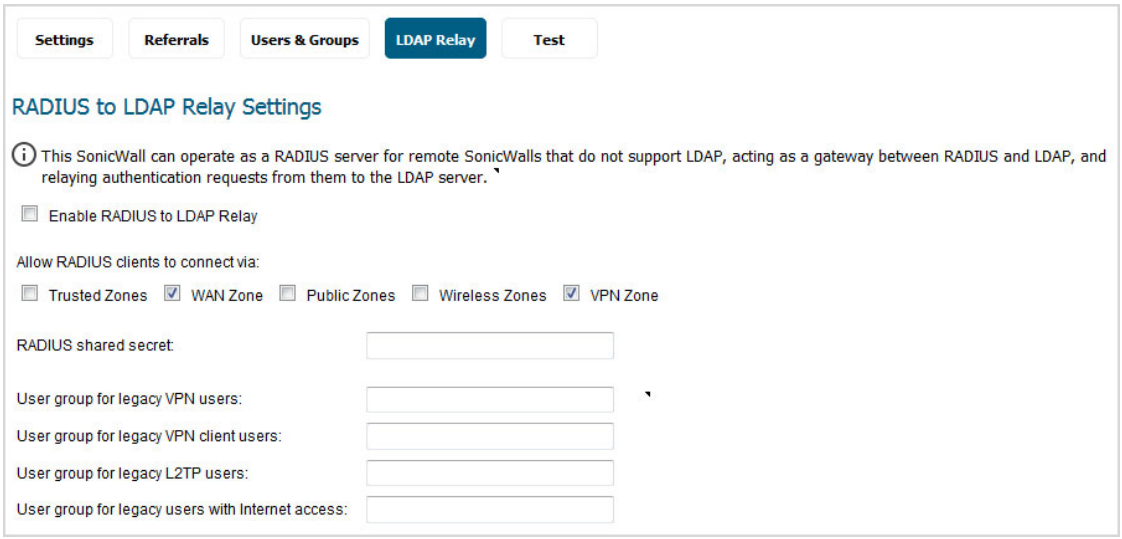
  These options provide different methods of detecting when mirrored user groups have been deleted from the LDAP server, and adjusting the local mirrored groups on the firewall.

- **Exclude groups in these sub-trees** – You can configure user group mirroring exclusions in this text box. Up to 64 sub-trees in the LDAP directory can be excluded from mirroring. Click **ADD** to enter a tree to be excluded, either in path format or as a distinguished name. Limited wildcard characters can be used. To change an existing tree in the exclusion list, select it and then click **EDIT**. To remove a tree from the exclusion list, select it and then click **REMOVE**.

LDAP Relay Page

To configure the LDAP server relay settings:

1 Click **LDAP Relay**.



The RADIUS to LDAP Relay feature is designed for use in a topology where there is a central site with an LDAP/AD server and a central SonicWall with remote satellite sites connected into it through low-end

Security Appliances that might not support LDAP. In that case the central SonicWall can operate as a RADIUS server for the remote SonicWalls, acting as a gateway between RADIUS and LDAP, and relaying authentication requests from them to the LDAP server.

2 Configure the following fields:

- **Enable RADIUS to LDAP Relay** – Enables this feature. This option is not selected by default.
- **Allow RADIUS clients to connect via** – Check the relevant checkboxes and policy rules are added to allow incoming RADIUS requests accordingly. **WAN Zone** and **VPN Zone** are selected by default.
- **RADIUS shared secret** – This is a shared secret common to all remote SonicWalls.
- **User groups for legacy VPN users** – Defines the user group that corresponds to the legacy 'Access to VPNs' privileges. When a user in this user group is authenticated, the remote SonicWall is notified to give the user the relevant privileges.
- **User groups for legacy VPN client users** – Defines the user group that corresponds to the legacy 'Access from VPN client with XAUTH' privileges. When a user in this user group is authenticated, the remote SonicWall is notified to give the user the relevant privileges.
- **User groups for legacy L2TP users** – Defines the user group that corresponds to the legacy 'Access from L2TP VPN client' privileges. When a user in this user group is authenticated, the remote SonicWall is notified to give the user the relevant privileges.
- **User groups for legacy users with Internet access** – Defines the user group that corresponds to the legacy 'Allow Internet access (when access is restricted)' privileges. When a user in this user group is authenticated, the remote SonicWall is notified to give the user the relevant privileges.

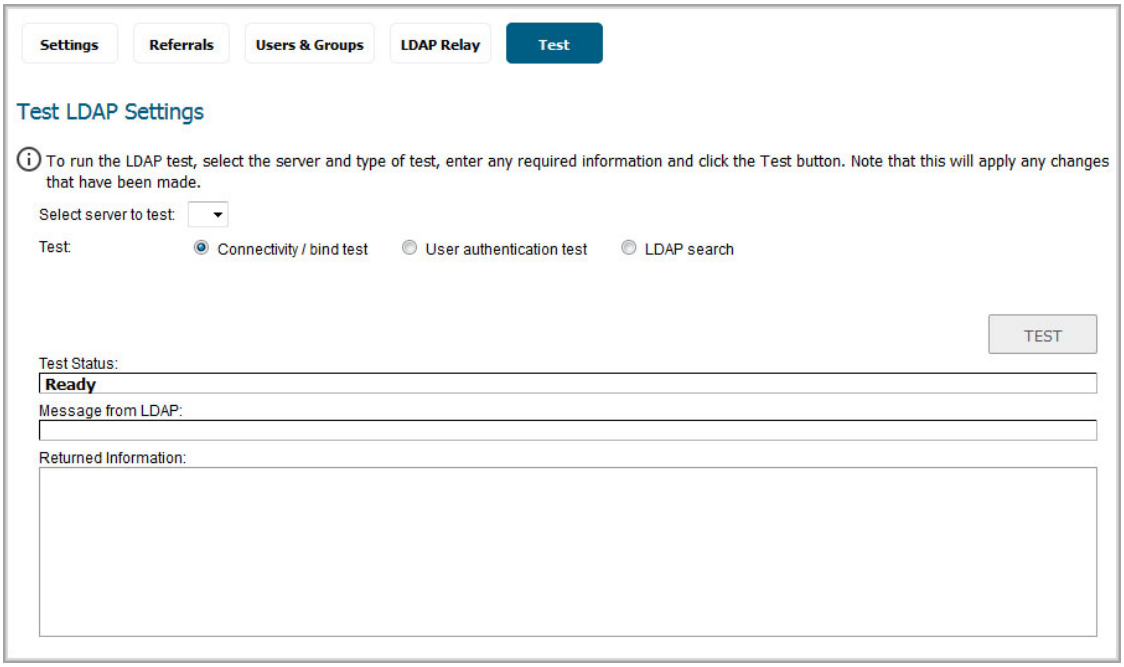
NOTE: The 'Bypass filters' and 'Limited management capabilities' privileges are returned based on membership to user groups named 'Content Filtering Bypass' and 'Limited Administrators' – these are not configurable.

Test Page

IMPORTANT: Testing the LDAP settings applies any changes you have made.

To configure the LDAP server test settings:

1 Select **Test** to test the configured LDAP settings:



The **Test LDAP Settings** page allows for the configured LDAP settings to be tested by attempting authentication with specified user and password credentials. Any user group memberships and/or framed IP address configured on the LDAP/AD server for the user are displayed.

About Extended Support for Multiple LDAP Servers

Multiple primary LDAP servers can be configured, one for each authentication partition, plus a list of additional servers for each. Each primary LDAP server is configured as per the current LDAP server. For the additional servers, configuration is minimal (common configuration from the primary server applies), but includes the login (bind) credentials and the sub-domain that the server controls .

NOTE: Active Directory has a 1:1 mapping of LDAP server to domain, which might not be the case with other LDAP servers. When there is a 1:1 mapping, configuring a domain for each LDAP server makes selection of the server efficient, but if that is not the case, selection is just less efficient.

The settings that are configurable separately per-server are those currently in the **System Setup | Users > Settings > Configure LDAP** dialog in the Management Interface. For more information about configuring LDAP, see Configuring the SonicWall for LDAP on page 149.

IMPORTANT: For correct operation, all the LDAP servers within a partition must be set to the same schema. If this is not the case, a warning is issued.

The **Referrals** settings are configured globally and are common across all the LDAP servers in all the authentication partitions.

NOTE: Explicitly configuring the secondary servers is optional. Each primary and secondary server can be configured separately, or a primary can be configured with all the user/group trees that can be accessed through it through referrals.

Topics:

- About Configuring Secondary Servers on page 163
- About Dynamically Learned Secondary Servers on page 163
- About Backup Servers on page 163

About Configuring Secondary Servers

Creating/configuring permanent secondary servers is the same as for the primary server, except for the primary/secondary setting. The only functional difference between them is that when a search is to be made and the location is not known from the configured user/group trees, then the search is sent to the primary server, with the primary server sending references/referrals when passing the search on to the secondary server(s) if needed.

About Dynamically Learned Secondary Servers

When a secondary server is accessed through a referral or reference for the first time, the Security Appliance binds to the secondary server after possibly trying multiple bind domain names [DNs) based on the various configured user trees. The Security Appliance internally creates a record for the secondary server in which the Security Appliance saves the bind information for future attempts. This process includes secondary servers that are not configured, thus creating a dynamic server object that is kept internally along with the server objects for the configured servers.

These dynamically learned server objects allow storing additional information as well as the current bind information, as per configured servers. The information includes the user/group trees that are learned by the server, plus statistics for the object.

NOTE: This information is not persistent over reboot and is re-learned as necessary. The configuration of the user/group trees of dynamic secondary servers, however, are saved with the primary server.

About Backup Servers

Support for backup servers is provided with Active Directory, where backup is achieved through the DNS name system. An Active Directory domain controller is accessed through the DNS name of either the machine or the domain; in the latter case, the domain name resolves to a list of the IP addresses of all of the domain's controller replicas. When the LDAP server DNS name resolves to a list of IP addresses, the SonicWall Security Appliance tries each in turn until one responds. Hence, configuring the LDAP server DNS name as the primary domain name rather than the domain controller machine name gives redundancy, with a backup server being used if the primary does not respond.

This mechanism also works in Active Directory for referrals and references because it returns the secondary domain's DNS name in a referral to the domain.

NOTE: In Active Directory, backup servers are generally referred to as replica servers.

One or more backups can be configured for each configured primary or secondary LDAP server. This configuration enables recording status and statistics for each individual server and provides support for redundancy with backup servers in non-Active Directory installations where the above DNS name mechanism does not provide such support.

A backup server has only a subset of the configuration that is set for other servers, as most of the configuration is identical to the server for which it is a backup. By default, only the host name or IP address of the backup server is needed.

About Importing and Mirroring from LDAP

To create local user groups that mirror those in the LDAP directory when LDAP User Group Mirroring is enabled, the SonicWall Security Appliance periodically auto imports user groups and user group nestings (memberships where groups are members of other groups) from the LDAP server(s)

You can select mirror user groups anywhere you can select regular user groups, such as in access rules and CFS policies. Mirror user groups do have a few restrictions, however, such as they cannot have other user groups added as members locally on the SonicWall Security Appliance, although mirror user groups can be made

members of other local user groups and local users can be made members of them. Users who are members of a user group on the LDAP server receive any access privileges set through its local mirror group automatically.

Topics:

- User Importation on page 164
- User Group Importation and Mirroring on page 164

User Importation

When user importation from LDAP is launched from the **LDAP Configuration** dialog or the **System Setup | Users > Local Users & Groups** page, there is an option to specify the LDAP server(s) from which to import:

- One specific LDAP server
- All the servers in an authentication partition (when the latter is enabled)
- All LDAP servers

To be able to distinguish users imported from different domains on different LDAP servers who might have the same usernames, there is also an option to create the local user object with one of a number of qualified username formats that include the domain. This option is in addition to using the simple username.

If a user account is imported with one of the qualified username formats, then:

- For web or client login using that account, the fully qualified username must be entered exactly as imported.
- When a user is identified through SSO, because the name formats can vary depending on the SSO source, the username and domain components are matched separately against those of the user object. For example, if a user is imported from LDAP as `jdoe@mydomain.com` and an SSO agent reports `MYDOMAIN/jdoe`, those match, and that user account is used to set additional group memberships for the user. Hence, for SSO, which qualified name format is selected does not really matter, and the choice comes down mainly to display preference.

  NOTE: This applies only if the **Use LDAP to retrieve user group information** or **Allow only users listed locally** option is set in **System Setup | Users > Settings**. For more information, see Configuring the SonicWall for LDAP on page 149 and Configuring SonicOS to Use the SonicWall SSO Agent on page 167.

User Group Importation and Mirroring

When using authentication partitioning, the users in a partition must get access permissions set for user groups imported from that partition, but not the access permissions for user groups of the same name imported from other partitions.

For example, imported/mirrored user groups are used in policies to select the applicable user groups by matching the group names in the policy against the group names read from LDAP at login time for the users. Imported and mirrored user groups work a bit differently (mainly for historical reasons):

- When a user group is imported manually, the local user group object is created with the simple group name with no domain component. Then, when users' group memberships are matched against the local group name, just the simple group names are compared, and any domain component is ignored. So, where user groups exist with the same names in different domains, users from any domain would get memberships set for the local group.
- When LDAP user group mirroring mirrors a group, the local user group object is created with the name, `group-name@domain.com`, to distinguish groups mirrored from different domains. Then, when users' group memberships are read from LDAP, they are put into the same format, and the full group name, including the domain component, is compared. Where user groups exist with the same names in

different domains, users from a domain only get memberships set for the groups mirrored from their own domain.

Manually imported user groups also have the option to import a qualified group name so they can be used as per mirrored groups above to set memberships separately for each domain's users. When importation of the groups is launched from the **System Setup | Users > Settings > Configure LDAP** dialog or the **Users > Local Users & Groups** page, the dialog has the same options as for users, except that the only choices for the format are **Simple name** or **name@domain** (the default).

NOTE: No explicit authentication partition recording/checking is needed for imported/mirrored user groups because matching the domain component implicitly ensures that only the groups from the domains in a user's partition are selected.

For backwards compatibility and ease of setting common access for members of standard groups across different partitions, if a user group is imported from LDAP (or manually created) with a simple name, then the domain is ignored when matching against that; hence, a simple name can be used to set access rights for users in any domain/partition.

For example, if you have:

- Partition A: `domain dom_a.com`
- Partition B: `domain dom_b.com`

and you then import the Administrators groups from both, selecting to import as **name@domain.com**, you import local user groups `Administrators@dom_a.com` and `Administrators@dom_b.com`. Users in each partition only receive access rights set for the relevant group; that is:

- When an administrative user from partition A logs in and an LDAP lookup finds that they are a member of the Administrators group in `dom_a.com`, they are given membership in `Administrators@dom_a.com`.
- Similarly, when an administrative user from partition B logs in, they receive membership in `Administrators@dom_b.com`.

If, however, you imported the Administrators group from either domain as a simple name, you would get a local user group named `Administrators` and administrative users in either partition would get any access rights set for that group.

Mirroring is enabled globally. When it is enabled, the user groups are mirrored from all configured and learned LDAP servers.

NOTE: It is possible to use the exclude feature with wild cards to exclude all groups on a server.

About LDAP Test

In LDAP test, you can select the LDAP server to test, add connectivity and search tests in addition to doing a user authentication test. See the LDAP Tests table.

LDAP Tests

Test	Function
Connectivity/bind	Simply tries to bind to the LDAP server with the configured bind credentials.
User authentication	Tests that a given username and password can be sent to and authenticated by the LDAP server.

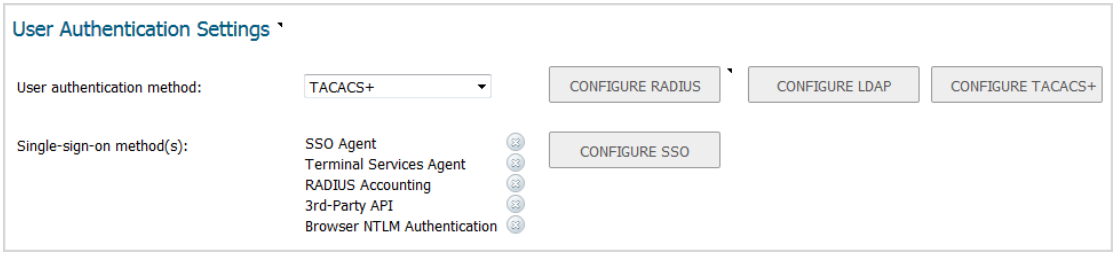
LDAP Tests

Test	Function
LDAP search	Has basic and advanced modes: **Basic mode** searches for a: • User with a given login name, qualified login name, or common name • User group with a given name or member **Advanced mode** allows: • An explicit search filter • Optionally, changing the search base and scope (the default is to search from the top of the domain sub-tree, with scope to search that entire sub-tree) • Searching for multiple objects • Limiting the information returned **TIP:** If you set up a user or user group search in basic mode and then switch to advanced mode, the fields will be pre-filled with the search parameters for that user/group search which you can then use as a starting point to set up a more complex search.

Configuring TACACS + for Authentication

To configure TACACS+:

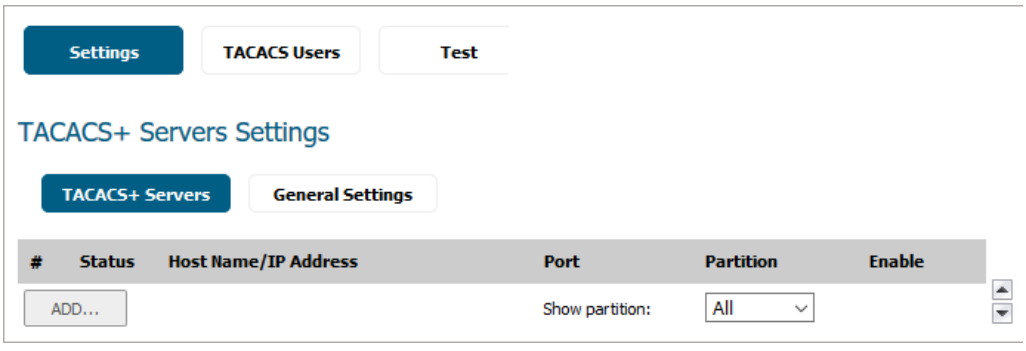
1 Navigate to **MANAGE | System Setup | Users > Settings**.



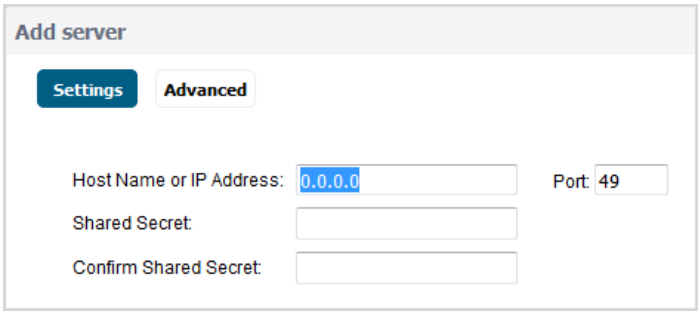
2 From **User authentication method**, select either:

- **TACACS+**
- **TACACS+ + Local User**

3 Click **CONFIGURE TACACS+**. The **TACACS+ Configuration** dialog displays.



4 Under the **TACACS+ Servers** table, click **ADD**. The **Add Server** dialog displays.



5 Enter the server's ID in the **Host Name or IP Address** field.

6 Enter the port number in the **Port** field. The default is **49**.

7 Enter the shared secret in the **Shared Secret** field.

8 Reenter the shared secret in the **Confirm Shared Secret** field.

9 Click **Save**. The server is added to the TACACS

10 Click **OK**.

Configuring SonicOS to Use the SonicWall SSO Agent

To configure your Security Appliance to use the SonicWall SSO Agent:

1 Go to **MANAGE | System Setup | Users > Settings**.

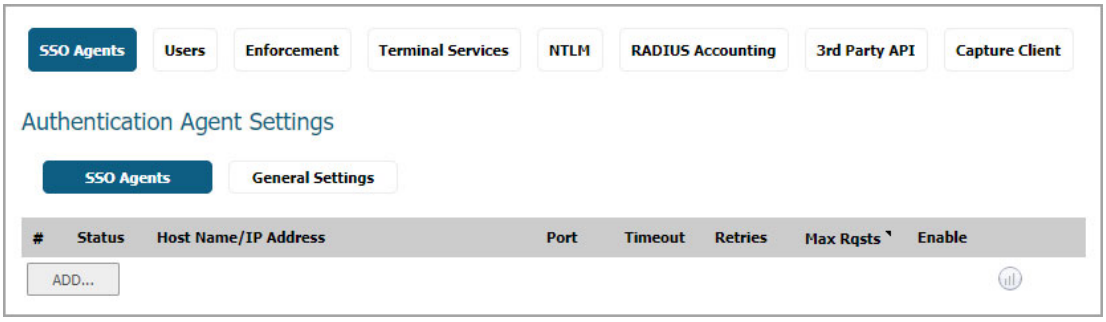
2 In the **User Authentication Settings** section, for **Single-sign-on method(s)**, select **SSO Agent**. Use this choice to add and configure a TSA as well as an SSO Agent for the SSO method.

3 Click **CONFIGURE SSO**. The **SSO Authentication Configuration** dialog displays.

Topics:

- SSO Agents Page on page 168
- Users Page on page 170
- Enforcement Page on page 174
- Terminal Services Page on page 176
- NTLM Page on page 178
- RADIUS Accounting Page on page 180
- 3rd Party API Page on page 187
- Capture Client Page on page 188
- Test Page on page 188

SSO Agents Page



On the **SSO Agents** page under **Authentication Agent Settings** you can view any SSO Agents already configured:

- The green LED next to the Agent's IP address indicates that the agent is currently up and running.
- A red LED would indicate that the agent is down.
- A grey LED shows that the agent is disabled.

The LEDs are dynamically updated using AJAX.

To configure an SSO agent:

TIP: As you type in values for the fields, the row at the top is updated in red to highlight the new information.

1 Click **ADD** to create an agent. The **Add Agent** dialog displays.

2 If you have partitions, choose which partition to add the agent.

3 Click **OK**.

4 Configure the SSO Settings options.

- For **Host Name or IP Address**, enter the name or IP address of the workstation on which SonicWall SSO Agent is installed. By default, **0.0.0.0** is entered.
- At **Port**, enter the port number that the SonicWall SSO Agent is using to communicate with the appliance. The default port is **2258**.

  NOTE: Agents at different IP addresses can have the same port number.

- At **Shared Key**, enter the shared key that you created or generated in the SonicWall SSO Agent. The shared key must match exactly. Re-enter the shared key in the **Confirm Shared Key** field.
- At **Timeout (seconds),** enter a number of seconds before the authentication attempt times out. This field is automatically populated with the default of **10** seconds.
- At **Retries**, enter the number of authentication attempts. The default is **6**.

5 Click **Advanced**.

6 At **Maximum requests to send at a time**, enter the maximum number of simultaneous requests to send from the appliance to the agent at one time. The default is **32**.

The agent processes multiple requests concurrently, spawning a separate thread in the agent PC to handle each. The number of simultaneous requests that the authentication agent can handle depends on the performance level of the PC that it runs on and of the network. Increasing this setting could make SSO user authentication more efficient, but setting it too high could swamp the agent by sending too many requests at a time, thus overloading the PC and resulting in timeouts and authentication failures.

On the other hand, if the number of simultaneous requests sent from the appliance is too low, some requests have to wait, possibly causing ring buffer overflows. Too many requests waiting could lead to

slow response times in Single Sign On authentication. If this setting cannot be increased high enough to avoid ring buffer warnings without getting a significant numbers of timeouts, then consider moving the agent to a higher-performance, dedicated machine, or possibly adding additional agents. For more information about checking for ring buffer overflows and related statistics in the SonicOS TSR, see Single Sign-On Advanced Features on page 104.

TIP: Look at the statistics in the **Single Sign On Authentication** section of the Tech Support Report. If significant numbers of timeouts are shown, then decreasing this value might help. If the **Maximum time spent on ring** approaches or exceeds the polling rate (configured on the **Users** tab) or if any ring buffer overflows are shown, then this value should probably be increased.

The page is updated to display a new row in the table at the top.

TIP: You can modify any of the entries by clicking on it. The entry turns into an editable field.

7 Click **General Settings** under **Authentication Agent Settings**.

8 Configure the following options:

- Select the **Enable SSO agent authentication** checkbox to use the SSO Agent for user authentication. This setting is selected by default.
- Select the **Try next agent on getting no name from NetAPI/WMI** checkbox to force a retry of the authentication through a different SSO agent if there is no response or error from the first agent. This setting is not selected by default.

  NOTE: This setting affects only agents using NetAPI/WMI, not any agents that use just the domain controller security log lookup mechanism.

  IMPORTANT: See also the **Poll the same agent that authenticated the user** setting on the **Users** tab, which needs to be set if this setting is enabled.

  The NetAPI/WMI protocols used by the SSO agent for user identification are provided by Windows, and what they actually do is outside the control of the agent or appliance. When using NetAPI or WMI, should Windows respond with no user name and no error to a request from an agent, then by default, the appliance assumes that other agents get the same and does not retry the request through another agent (as it would do should it receive an error response).

  If you see authentication failures logged as `SSO agent returned no user name` when you think the users should have been identified, try enabling this setting. If this setting is enabled when the appliance receives a no-user-name response from an agent, the appliance treats the response as an error and retries the request through a different agent.

  Typically, enabling this setting is needed in a situation where only some of the agents can reach certain users; for example, if it is necessary to place an agent at a remote site to identify the users there because they cannot be reached easily by the agents at the central site.

- Select the **Don't block user traffic while waiting for SSO** checkbox to use the default policy while the user is being identified. This prevents browsing delays. This setting is not selected by default.

  When a user is being identified through SSO, traffic from the user is normally blocked until identification is complete so proper policies can be applied where applicable. Sometimes an SSO agent takes a significant time to identify a user, however, and that delay can result in users experiencing browsing delays.

  This setting allows you to override that delay and instead allow users traffic through while waiting for SSO, with default policies applied until the user is identified.

You also can choose whether to allow through traffic when a user needs to be identified for an access rule that requires user authentication (that is, when a user would not otherwise be allowed any access if not identified).

CAUTION: Take care with doing this as it can temporarily allow through a user who would not be allowed when identified. If you choose to do this for selected access rules, then a setting for it appears in the advanced settings of those rules that require user authentication.

- Select the **Including for** checkbox and either the **All access rules** (default) or the **Selected access rules** radio button to allow traffic affected by access rules that require user authentication, while waiting for user identification.

CAUTION: This can temporarily allow access that would not be allowed when the user is identified.

- To have all the SSO agents synchronize their user databases, select either:
  - **Sync all agents** – To synchronize together no matter what identification mechanisms they use, thus giving a single, homogeneous user database duplicated on every agent.
  - **Synch those with the same user identification mechanisms** – To synchronize only those databases using the same identification mechanism; this is the default.

  Each SSO agent maintains its own database of the users that it has identified, and the agents can optionally be configured to synchronize those databases, thus giving a common user database duplicated on each agent. A common, synchronized user database makes user lookups more efficient and gives better redundancy. By specifying synchronicity here, the appliance can inform each agent of the other agents with which to synchronize, thereby avoiding the complexity of having to configure it in the agents.

  By default, the appliance has those agents configured to use the same user identification mechanisms synchronize together. For example, if some agents are reading domain controller logs while others use NetAPI, then two separate, external databases in the two groups of agents result, one database of those user found in the domain controller logs and a separate database of the users identified by NetAPI.

  NOTE: This setting can be overridden by explicitly configuring in each SSO agent the list of other agents with which to synchronize.

- Configure the list of Windows service user names in the **User names used by Windows services** table. You can list up to 64 user names that might be used by services on the end-users' PCs; any log ins with these names are assumed to be service log ins and are ignored by the SSO agent(s).

  a) Click **ADD**. the **Service User name** dialog displays.

  b) Enter the service user name.

  c) Click **OK**.

  d) Repeat Step a through Step c for each user account.

  Windows services log on to the machine or domain using user accounts just as real users do. Some of the Windows' APIs used by the SSO agent do not provide for distinguishing these service log ins from real user log ins, which can lead to the SSO agent incorrectly reporting the user name used by a service instead of that of the actual user.

Users Page

1 Click **Users** to specify the following the **User Settings** options:

- Select the **Allow only users listed locally** checkbox to allow only users listed locally on the appliance to be authenticated. This setting is disabled by default.

- Select the **Allow limited access for non-domain users** checkbox to allow limited access to users who are logged in to a computer but not into a domain. These users are not given membership in the Trusted Users user group, even when set locally, and so do not get any access set for Trusted Users. The are given access through policies, etc., that apply to Everyone or that specifically list them as allowed users. This setting is disabled by default.

  These users are identified in logs as *computer-name/user-name*. When using the local user database to authenticate users and the **Simple user names in local database** option is disabled, user names must be configured in the local database using the full *computer-name/user-name* identification.

  NOTE: This does not apply for users authenticated through NTLM. With NTLM, authentication non-domain users are given access only if the name/password matches a local user account created on the appliance.

- If your network includes non-Windows devices or Windows computers with personal Security Appliances running:

  a) Select the **Probe user for** checkbox.

  b) Choose one of the following, depending on which is configured for the SSO Agent:

  - **NetAPI over NetBIOS**
  - **NetAPI over TCP**
  - **WMI**

  TIP: Hovering the mouse over these options displays a small tooltip contain the TCP port number.

  When the SSO agent attempts to identify a user in a Windows domain, if the agent uses NetAPI or WMI, then when the agent tries to communicate directly with the user's computer from which the traffic is originating. This can cause problems:

  - When traffic is coming from non-Windows devices as such devices do not respond to, or might block, the Windows networking messages used by the SSO Agent to identify a user.
  - With Windows computers if personal Security Appliances on them are blocking them.

  The result can be that the agent might get overloaded with multiple threads waiting for requests that are not getting replies.

  To avoid these problems, enable this setting (it is disabled by default) and select the correct NetAPI/WMI protocol that the SSO agent is configured to use. Before sending a request to the agent to identify a user through NetAPI or WMI, the SonicWall appliance probes the machine from which the traffic originated to verify if it responds on the port used by the NetAPI or WMA protocol. If it does not, then the device fails SSO immediately without the agent getting involved.

  NOTE: This setting does not affect an agent that reads user login information from the domain controller(s).

- If the **Probe users for** setting is enabled, it causes the Security Appliance to probe for a response on the NetAPI/WMI port before requesting that the SSO Agent identify a user. The **Probe timeout (seconds)** is set to **5** seconds by default.
- Select the **Probe test mode** checkbox to test that SSO probes are functioning correctly during SSO without interfering with user authentications. Probes are sent after initiating user authentication through the SSO agent. This setting is disabled by default.

  If this setting is enabled, the probes are sent after initiating user authentication through the SSO agent (normally, the latter is done if the probe is successful). Statistics for the probing are updated as normal, and if a probe fails for a user who is successfully authenticated by the agent, then that is reported through a message on the console port.

- For the **Mechanism for setting user group memberships**, choose either:
  - **Use LDAP to retrieve user group information** radio button to use LDAP to retrieve user information. This option is selected by default.
    - To configure the LDAP settings click **CONFIGURE**. The **LDAP Configuration** dialog displays. For configuration information for this dialog, refer to Advanced LDAP Configuration on page 189.
  - **Local configuration** radio button to use locally configured user group settings.
- In the **Polling rate (minutes)** field, enter a polling interval, in minutes (the default is **5**). After a user has been identified and logged in, the SonicWall polls the authentication agent at this rate to verify the user is still logged on.

  If you are using NTLM authentication, then in the NTLM settings you can selectively choose to have the appliance poll users by forcing them to re-authenticate through NTLM rather than polling through the agent.
- Select the **Poll the same agent that authenticated the user** checkbox if the network topology requires that particular agents be used depending on the location of users, rather than polling any agent to determine if the user is still logged in. This setting is disabled by default.

  IMPORTANT: The **Try next agent on getting no name from NetAPI/WMI** setting on the **SSO Agents General Settings** tab also needs to be set if this is set.

  By default, the appliance assumes that any SSO agent can send NetAPI or WMI requests to any user, so when polling to check if users are still logged in, the appliance can choose any agent based on current loadings. If this is not the case, and the network topography requires particular agents be used depending on the location of the users, then enable this setting. When it is enabled, after a user is successfully identified by an agent, subsequent polling of the user is performed through that same agent.

  NOTE: This setting affects only agents using NetAPI/WMI, not any agents that use just the domain controller security log lookup mechanism.
- In the **Hold time after (minutes)** field, enter a time, in minutes, that the Security Appliance waits before trying again to identify traffic after an initial failure to do so. This feature rate limits requests to the agent to avoid possibly flooding it with requests if further traffic continues to be received from sources that repeatedly fail SSO. The default is **1** minute.

  NOTE: The times to hold off after getting errors from the SSO agent and after the agent reports that no user is logged in are set separately, so they are configured separately.
- In the **…after finding no user** field, enter the number of minutes that the appliance should wait before trying again if it gets errors from the SSO agent or when the agent reports that no user is logged in. The default is **1** minute.
- If **Ramp Up** is enabled, then the first retries after an SSO failure is quick, with the hold-off period growing exponentially on subsequent failures until it reaches configured rate. This can help avoid significant delays in identifying a user after a transient failure, which can happen, for example, if traffic emanates from their PC during boot-up. This option is selected by default.

  TIP: If you encounter a problem that ramp-up does not help, then that might be because it ramps up too quickly before the problem has cleared, so try slowing down the ramp-up rate.
- Select the ramp-up rate:
  - **4 (fast ramp-up)** (default)
  - **3**
  - **2**

- **1 (slow ramp-up)**

2 To give consistent naming for a domain in logging, select one of the following radio buttons for **When different SSO sources report different name variants for a user's domain**:

- **Use the domain name as received** (default)
- **Always use a consistent domain name**; go to Step a.

By default, a user identified through SSO is logged in on the SonicWall appliance with whatever domain name is reported to it by the external source that identified the user. A doman, however, typically has two or three different variants of its domain name (for example, a Windows domain has its DNS name, its NetBIOS name, and its Kerberos realm name), and different SSO sources might report different variants of these for a user in the same domain.

This difference can cause difficulty in tracking users by domain in logging, so you can instead select to make the names consistent by having the same domain name variant used for all the users in a domain, no matter which variant is reported to the SonicWall appliance.

a If you have selected **Always use a consistent domain name**, click **SELECT**. The **Select the name variant to use for each domain** pop-up dialog lists known domains from which you can select the names to use is displayed.

b Select the variant(s) to use. The initial default variant for each domain is **None**, which means that behavior of using whatever domain name is reported to the appliance through SSO does not change until **Always use a consistent domain name** is enabled and the domain name to use is selected here.

> **NOTE:** If a domain is not shown in this list, wait until the SSO has identified some users in the domain, then repeat this step.

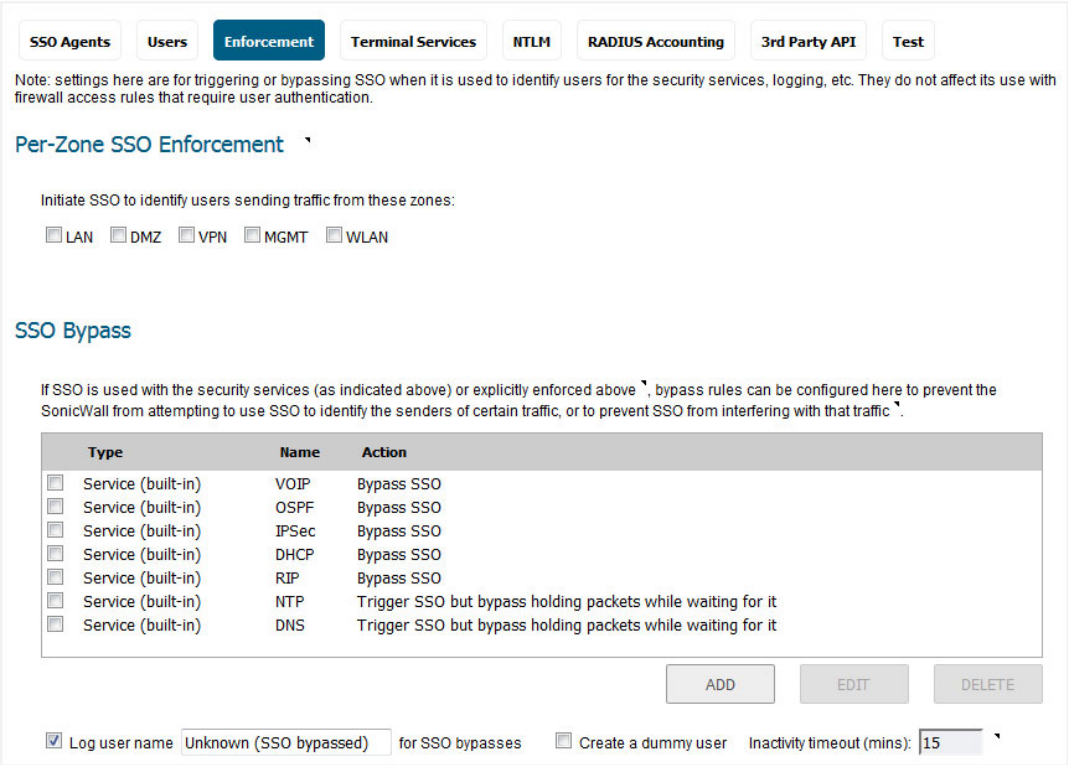
c Click **OK**.

If, when using Single Sign On, you see unexpected user names shown on the **MANAGE | System Setup | Users > Status** page, or logs of user login or user login failure with unexpected user names, those might be due to Windows service logins and those user names should be configured here so that the SSO agent knows to ignore them.

In cases where there are multiple Security Appliances communicating with an SSO agent, the list of service account names should be configured on only one of them. The effect of configuring multiple lists on different appliances is undefined.

Enforcement Page

1 Click **Enforcement** if you want to either trigger SSO on traffic from a particular zone, or bypass SSO for traffic from non-user devices such as internal proxy web servers or IP phones.



2 Under **Per-Zone SSO Enforcement**, select the checkboxes for any zones on which you want to trigger SSO to identify users when traffic is sent:

- **LAN**
- **DMZ**
- **VPN**
- **MGMT**
- **WLAN**

If SSO is already required on a zone by Application Control or other policies, those checkboxes are pre-selected and cannot be cleared. If Guest Services is enabled on a zone, SSO cannot be enforced and you cannot select the checkbox. On zones where it is not otherwise initiated, SSO enforcement can be enabled by this option.

NOTE: On zones where security services policies or access rules are set to require user authentication, SSO is always initiated for the affected traffic, and it is not necessary to also enable SSO enforcement here.

These per-zone SSO enforcement settings are useful for identifying and tracking users in event logging and AppFlow Monitor visualizations, even when SSO is not otherwise triggered by content filtering, IPS, or Application Control policies, or by access rules requiring user authentication.

3 To bypass SSO for traffic from certain services or locations and apply the default content filtering policy to the traffic, select the appropriate service or location from the list in the **SSO Bypass** table or add a new

service or location to the table. The table displays the built-in services that bypass SSO; these services cannot be delete.

TIP: You could create SSO bypass address and/or service group objects for this and reference those same ones both here and in those access rules.

NOTE: SSO bypass settings do not apply when SSO is triggered by access rules requiring user authentication. To configure this type of SSO bypass, add separate access rules that do not require user authentication for the affected traffic. For more information on configuring access rules, see *SonicOS 6.5 Policies*.

By default, Linux and Mac users who are not authenticated by SSO through Samba are assigned the default content filtering policy. To redirect all such users who are not authenticated by SSO to manually enter their credentials, create an access rule from the **WAN** zone to the **LAN** zone for the **HTTP** service with **Users Allowed** set to **All**. Then configure the appropriate CFS policy for the users or user groups. For more information on configuring access rules, see *SonicOS 6.5 Policies.*

SSO bypass might be necessary, for example, for:

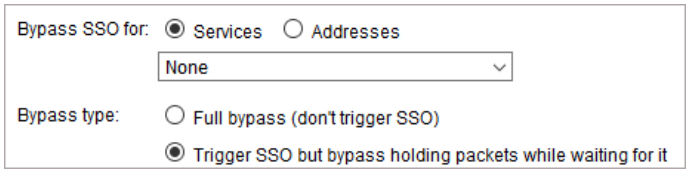
- Traffic emanating from a non-user device, such as an internal mail server or an IP phone.
- User traffic that does not need to be authenticated and might be adversely affected by delays waiting for SSO.

For traffic that bypasses SSO, the default content filtering policy is applied. If any APP rules or IPS/Anti-Spyware policies are set to include/exclude users, then that traffic is no included/excluded respectively with those.

The second setting is appropriate for user traffic that does not need to be authenticated, and triggering SSO might cause an unacceptable delay for the service.

4 Optionally, to add a service or location:

a Click **ADD**. The **Add an SSO bypass rule** dialog displays.



b For **Bypass SSO for**, choose either **Services** or **Addresses**.

c Select a service or address from the drop-down menu.

d Choose the **Bypass type**:

- **Full bypass (don’t trigger SSO**
- **Trigger SSO but bypass holding packets while waiting for it** (default)

e Click **ADD**. The entry is added to the table

5 To select a SSO bypass user name for logging:

a Select **Log user name *<bypass name>* for SSO bypasses**.

b Specify a name for the SSO bypassed user. The default is **Unknown (SSO bypassed)**.

This setting is selected by default, and a default name of **SSO Bypass** is specified. When this setting is enabled, and when traffic bypasses SSO (as configured here), the traffic is shown in logs and AppFlow Monitor with the given user name rather than as from an unknown user, thus allowing it to be differentiated from traffic sent by users whom SSO could not identify.

TIP: You also can configure logging on the **Users > Settings** page under **User Session Settings**.

6 Optionally, **Create a dummy user**. This setting is not selected by default.

If this setting is enabled along with **Log user name <*bypass name*> for SSO bypasses**, on receiving SSO bypass traffic, a dummy user entry is created with the given user name for the originating IP address. Then, in addition to the name appearing in logs and the AppFlow Monitor, the dummy user entry displays on the **MANAGE | System Setup | Users > Status** page. The dummy name remains in existence until traffic from the IP address stops for the given inactivity time or, in the case of bypass services, until non-bypass traffic is received from it.

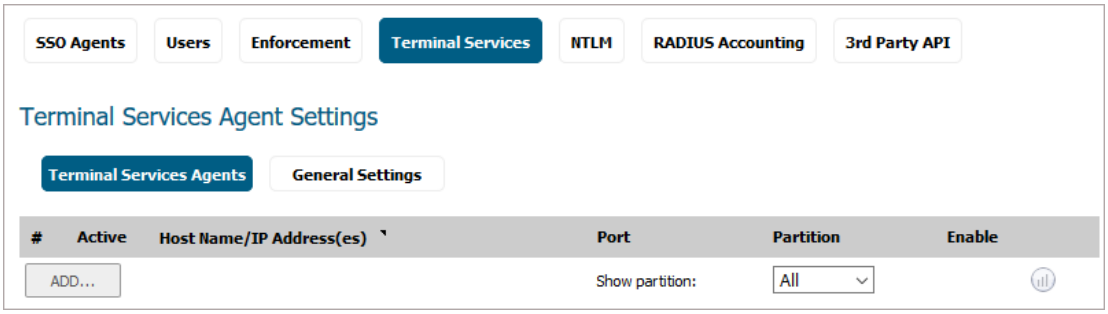
NOTE: This dummy user name applies only for bypass rules set for full SSO bypass. Any set to trigger SSO, but bypass holding packets while waiting for it results in the user being set according to the result of the triggered SSO identification.

NOTE: The logging part of this option also can be configured by the **For logging of connections on which the user is not identified** option in the **User Session Settings** section of the **Users > Settings** page.

a Optionally, specify an inactivity timeout, in minutes, in the **Inactivity timeout (mins)** field. The default is **15** minutes.

Terminal Services Page

1 Click **Terminal Services** to specify the following **Terminal Services Agent Settings** options.



2 To add agents, click **ADD**. The **Add Terminal Services Agent** dialog displays.



TIP: If partitioning is enabled, the **Select partition** dialog displays over the **Add Terminal Services Agent** dialog.

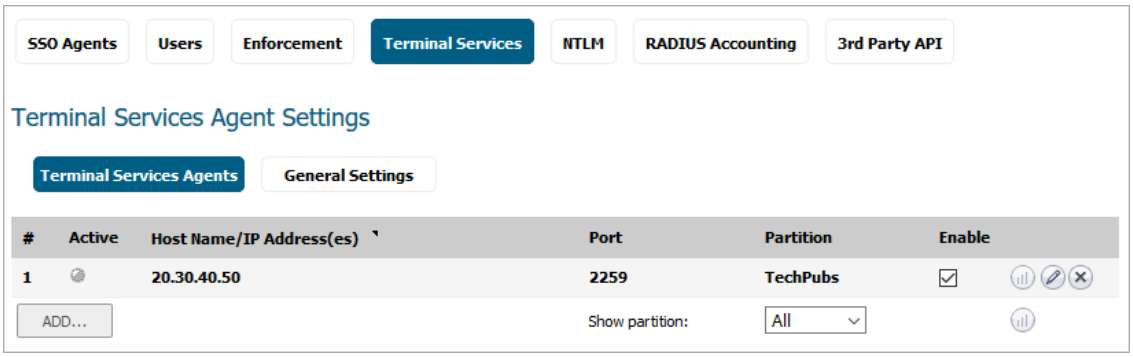
1 Choose the partition to which to add the new agent.
2 Click **OK**.

- In the **Host Name or IP Address(es)** field, enter the name or IP address of the terminal server on which SonicWall TSA is installed. If the terminal server is multi-homed (has multiple IP addresses) and you are identifying the host by IP address rather than DNS name, enter all the IP addresses as a comma-separated list.
- At **Port**, enter the port number that the SonicWall TSA is using to communicate with the appliance. The default port is **2259**.

  NOTE: Agents at different IP addresses can have the same port number.

- In the **Shared Key** field, enter the shared key that you created or generated in the SonicWall TSA. The shared key must match exactly. Re-enter the shared key in the **Confirm Shared Key** field.

  TIP: The shared key must be an even number of characters (

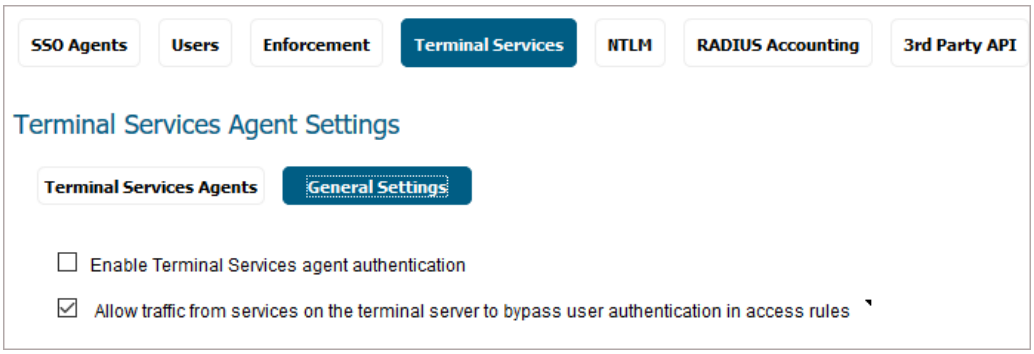


The page is updated to display a new row in the table at the top and new input fields in the lower half of the page. For existing agents:

- Green LED-style icon next to an agent indicates that the agent is up and running.
- Red LED icon indicates that the agent is down.
- Yellow LED icon means that the TSA is idle and the firewall has not heard anything from it for 5 minutes or more.

Because TSA sends notifications to the appliance rather than the appliance sending requests to the agent, a lack of notifications could mean that there is a problem, but more likely means simply that no user on the terminal server is currently active.

3 Click **General Settings** under **Terminal Services Agent Settings** to configure the following options:



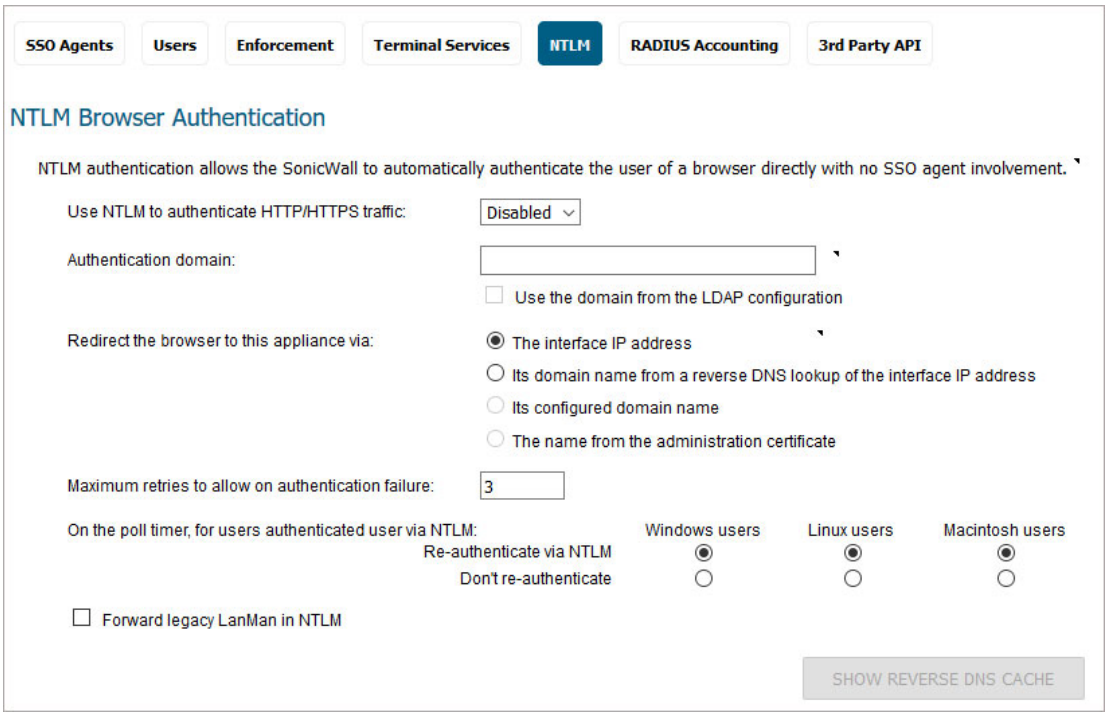
- Select **Enable Terminal Services agent authentication** to use the TSA for user authentication. This option is not selected by default.
- **Allow traffic from services on the terminal server to bypass user authentication in access rules** is selected by default. This allows service traffic, such as Windows updates or anti-virus updates not associated with any user login session, to pass without authentication. That traffic normally would be blocked if the applicable access rules are set to require user authentication.

  If you deselect this option, traffic from services can be blocked if access rules require user authentication. In this case, you can add rules to allow access for **All** to the services traffic destinations, or configure the destinations as HTTP URLs that can bypass user authentication in access rules.

NTLM Page

IMPORTANT: RADIUS must be configured.

1 Click **NTLM**.



NTLM authentication is supported by Mozilla-based browsers and can be used as a supplement to identifying users through an SSO agent or, with some limitations, on its own without the agent. The Security Appliance interacts directly with the browser to authenticate the user. Users logged in with domain credentials are authenticated transparently; in other cases the user might need to enter credentials to login to the appliance, but should only need to do so after the credentials are saved.

For more information about NTLM, see How Does Browser NTLM Authentication Work? on page 99.

2 Configure these settings;

- Select one of the following choices from **Use NTLM to authenticate HTTP traffic**:
  - **Disabled** – Never uses NTML authentication
  - **Enabled** – Try to authenticate users with NTLM before using the SonicWall SSO agent
- For **Authentication domain**, do one of the following:
  - Enter either:
    - The full DNS name of the domain in the form `www.somedomain.com`.
    - If using LDAP, the same domain that is entered in the LDAP configuration.

    In the NTLM protocol, the server (this firewall) might need to communicate its domain to the browser, and authentication might only be fully transparent to the user if the browser sees that as the local domain.

TIP: In the NTLM protocol, the domain is something that might be requested/required by the browser in addition to the actual authentication of the name/password. Whether the domain affects the outcome of the authentication depends on the browser.

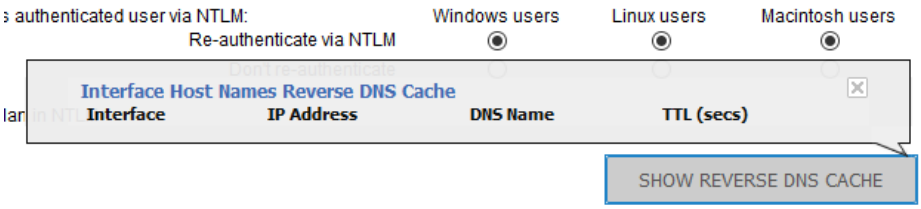
- Select **Use the domain from the LDAP configuration** to use the same domain that is used in the LDAP configuration.

Fully transparent authentication can only occur if the browser sees the appliance domain as the local domain.

- For **Redirect the browser to this appliance via**, select one of the following options to determine how a user's browser is initially redirected to the Security Appliance's own Web server:
  - **The interface IP address** – Select this to redirect the browser to the IP address of the appliance Web server interface. This option is selected by default.
  - **Its domain name from a reverse DNS lookup of the interface IP address** – Enables **Show Reverse DNS Cache** at the bottom of the window; when clicked, a popup displays the appliance Web server's Interface, IP Address, DNS Name, and TTL in seconds. Click the button to verify the domain name (DNS name) being used for redirecting the user's browser.
  - **Its configured domain name** – Use the Security Appliance's domain name as configured on the **System > Administration** page.
  - **The name from the administration certificate** – Use the imported certificate that is selected for HTTPS Web Management on the **Appliance > Base Settings > Appliance > Base Settings** page.
- Enter a number of retries in the **Maximum retries to allow on authentication failure**. The default is **3**, and the maximum is 99.
- To detect when users log out, select the polling method to be used by the appliance for Windows, Linux, and Macintosh users in the **On the poll timer, for users authenticated user via NTLM** options. Choose the option for users on each type of computer (**Windows users**, **Linux users**, **Macintosh users**):
  - **Re-authenticate via NTLM** – This method is transparent to the user if the browser is configured to store the domain credentials, or the user instructed the browser to save the credentials.
  - **Don't re-authenticate** – If you select this option, logout is not be detected other than through the inactivity timeout.

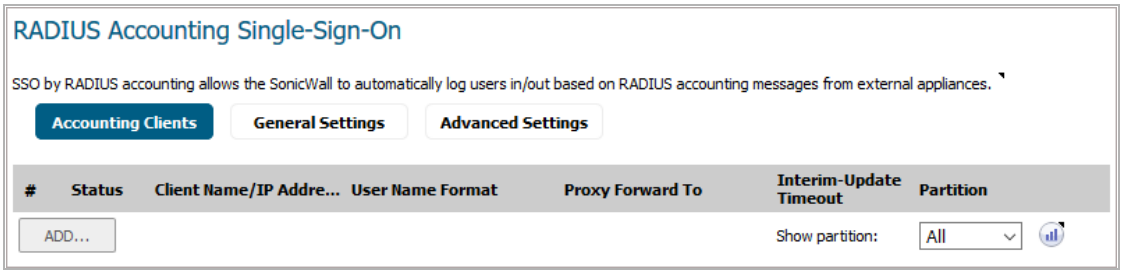
NOTE: When multiple Content Filter policies are configured and NTLM is enabled for Single Sign-On enforcement, an HTTP/HTTPS access rule with Trusted Users as Users Allowed must be added to the LAN to WAN rules in the **MANAGE | Policies > Rules > Access Rules** page. This rule triggers an NTLM authentication request to the user. Without the access rule, restrictive CFS policies might block the user from Internet access and prevent authentication.

- If you are using older legacy servers that require legacy LAN Manager components to be included in NTLM messages, select the **Forward legacy LanMan in NTLM** checkbox. This might cause authentication to fail in newer Windows servers that do not allow LanMan in NTLM by default because it is not secure.
- If you chose Its domain name from a reverse DNS lookup of the interface IP address, click **SHOW REVERSE DNS CACHE**. A popup displays:



RADIUS Accounting Page

1 Click **RADIUS Accounting** to display **RADIUS Accounting Single-Sign-On**.



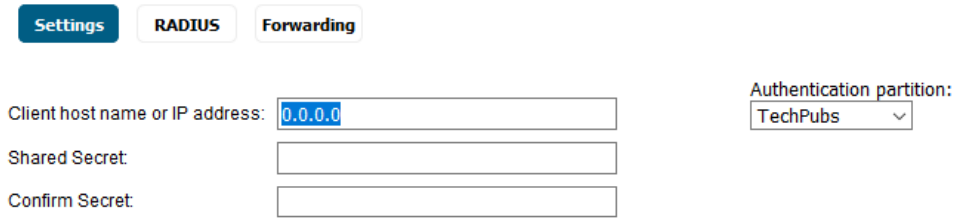
Single Sign-On by RADIUS accounting allows the appliance to act as a RADIUS accounting server for external third-party appliances, and to log users in or out based on the accounting messages from those devices. For third-party appliances that use RADIUS accounting for other purposes, SonicOS can also forward the RADIUS accounting messages to another RADIUS accounting server.

The **Status** column shows the current status for each RADIUS accounting client listed in the panel:

- Green—the client is active
- Yellow—the client is idle
- Grey—the client is not detected

2 To add a new RADIUS client, click **ADD**. The **Add RADIUS Accounting Client** dialog displays..

NOTE: Changes made in the dialog are instated directly into the highlighted entry in the **Accounting Clients** table as they are made. On completion, click anywhere outside of the pane to close it. Individual fields in the **Accounting Clients** table also can be updated by clicking on them directly in the table.



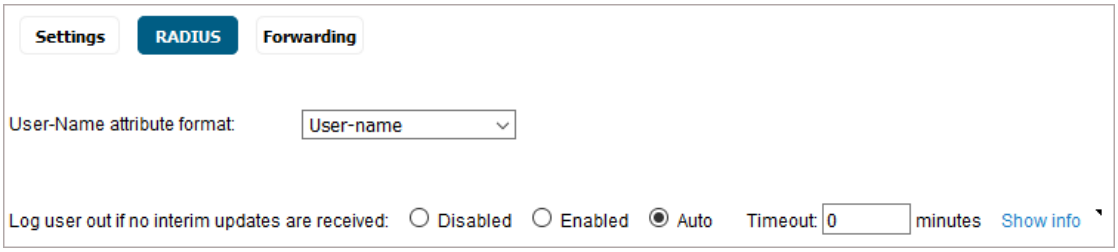
TIP: If partitioning is enabled, the **Select partition** dialog displays over the **Add RADIUS Accounting Client** dialog.

1 Choose the partition to which to add the new agent.
2 Click **OK**.

3 In the **Client host name or IP address** field, enter the name or the IP address for the RADIUS client host.

4 In the **Shared Secret** field and the **Confirm Secret** field, enter your shared secret for the client.

5 Click **RADIUS**.



6 From **User-Name attribute format**, select the format for the user name login.

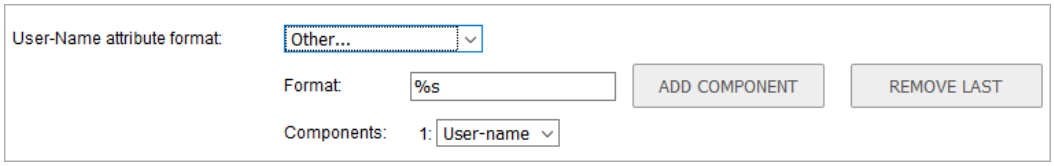
RADIUS Accounting does not specify the format of the content of the User-Name attribute passed in RADIUS Accounting messages. You need to enter, therefore, the format that is sent by the client. You can select from some common formats:

- **User-name**
- **Domain\User-name**
- **Domain/User-name**
- **User-name@Domain**
- **SonicWall Aventail** (**SonicWall SMA**)
- **Other** – Non-standard format

> **IMPORTANT:** The predefined formats are for common cases. If those do not match what your network access server sends, then you must select **Other** as the **User-Name** attribute format and then enter a customized format.

7 If you selected:

- A standard format, go to Step 8.
- If you select **Other**, more settings appear so you can configure the components to be found in the attribute:



  - **Format**
  - **Components**
  - **ADD COMPONENT**
  - **REMOVE LAST**

a In the **Format** field, enter a limited scanf-style string, with either a `%s` or `%[...]` directive for each component. This directive tells the appliance what the network access device (NAS) sends in the **User-Name** attribute. This format is not specified by the RADIUS Accounting RFC. Devices are not constrained as to what they can send in this attribute, so, its content can be very variable. What you set here specifies how the appliance must decode the **User-Name** attribute to extract the user name, domain, and/or DN.

> **TIP:** When you select **Other**, these fields are set to the format string and components of the previously selected format. So, first select the pre-defined format that most closely matches what your network access server sends. This gives you a good starting point for entering your customized format. Then, change to **Other**.

b From **Components**, select one of these

- **Not used**
- **User-Name** (default)
- **Domain**
- **DN**

The components that you enter as a limited scanf-style string in the **Format** field consist of one or more of the following items:

- User-Name
- Domain
- Fully qualified distinguished name (DN)

NOTE: You can double click **Components** to display a tooltip with instructions on how to enter the scanf-style format.

c Click **ADD COMPONENT**. The **Add a component to the User-Name attribute format** dialog displays.

NOTE: If you understand the scanf-style format, you can edit the **Format** field directly instead of using **Add component**.

TIP: Use %s for a component that is followed by white space or is at the end. For a component followed by some other character, use %[^x]x. For example, the **Format** string for the `name@domain` format would be `%[^@]@%s`, with the three components set to **User-Name**, **Domain**, and **Not used**.



d Select the type of component from **Component to add**:

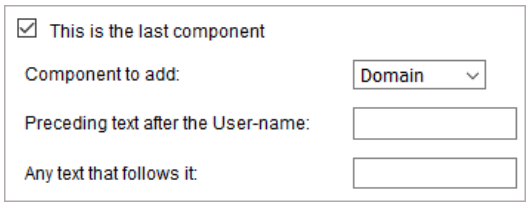
- **User-name**
- **Domain** (default)
- **DN**

e Enter text to separate entries in the **Preceding text after the User-name** field.

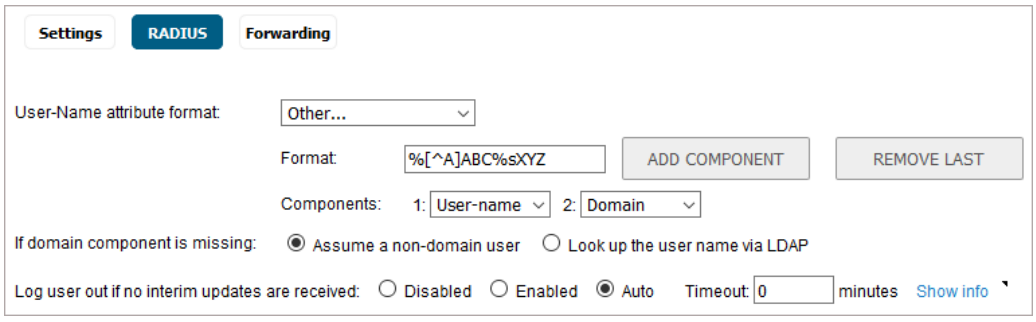
f Repeat Step c through Step e for each component.

To delete the last component you added, click **REMOVE LAST**.

g If this is the last user-name component, select **This is the last component**. Another option displays.



h Click **ADD**. More options appear on the **Add a component to the User-Name attribute format** dialog.



i From **If domain component is missing**, choose:

- **Assume a non-domain user** – This option is selected by default.
- **Look up the user name via LDAP**.

8 A RADIUS Accounting client can optionally send Interim Update messages periodically while a user is logged in. If the client does sends the messages at a reasonably constant interval, then the SonicWall appliance can monitor them and assume that the user has been logged out should the messages stop being sent. This process gives a fallback mechanism to guard against missing RADIUS Accounting Stop messages, which are sent on user log out.

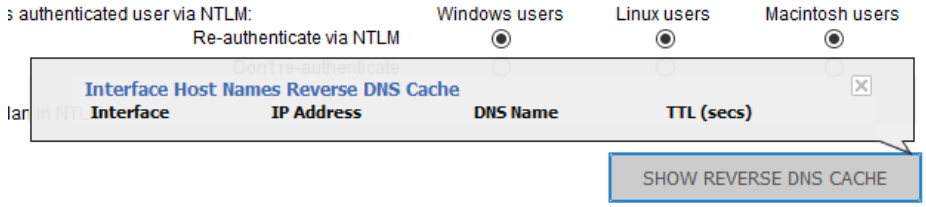
Select a **Log user out if no accounting interim updates are received** option:

- **Disabled** – to not have messages sent.
- **Enabled** – to manually specify the **Timeout** interval. Set the timeout value greater than the period at which the RADIUS Accounting client sends the Interim-Update messages, and for dropped/missed Interim-Update messages, set the **Timeout** value at least 2 to 3 times greater than the period.
- **Auto** (This option is selected by default.) – to have the firewall detect automatically whether Interim-Update messages are being sent periodically and, if they are, to use them as specified under Enabled and setting automatically the timeout accordingly.

NOTE: If, after some time, the timeout stays at 0 (zero) when the page is reloaded, then it has not detected them being sent and is not timing them out.

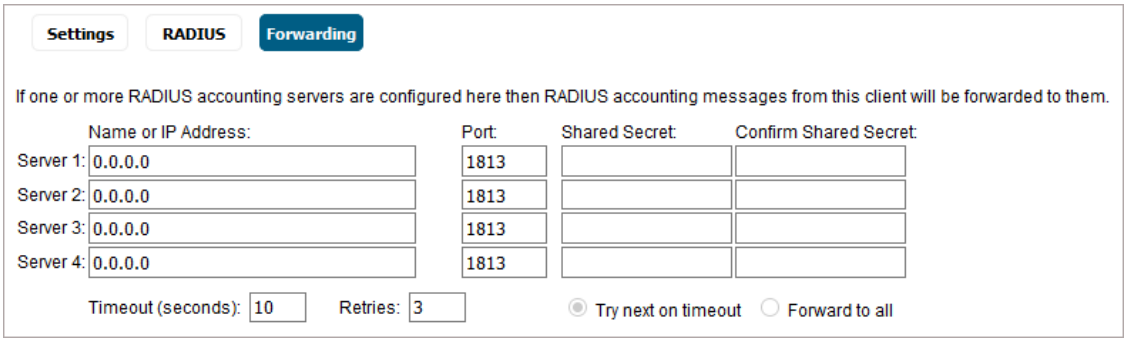
It could take quite a considerable time to complete auto-detection, depending on how frequently the client sends them. For example, if the client sends them every 10 minutes, then it could take over 30 minutes before the measure timeout is shown here.

TIP: You can click the **Show info** link to monitor progress in a popup dialog.



TIP: To rerun auto-detection, change the setting to **Disabled** and then back to **Auto**, clicking **SAVE** after each change.

9 Click **Forwarding**.

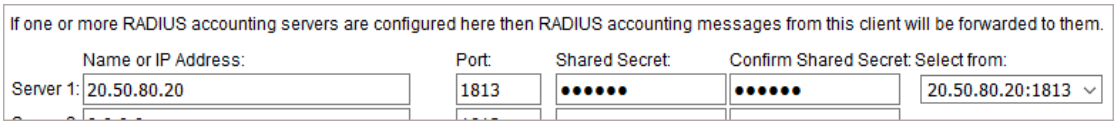


10 You can enter up to four RADIUS accounting servers in these fields:

- **Name or IP address**
- **Port** (default **1813**)

- **Shared Secret** for the RADIUS accounting servers to which you want the client to forward message
- **Confirm Shared Secret**

When you enter this information for a server, the **Select from** drop-down menu displays for each server.



11 For each server, from **Select from**, select either:

- **No forwarding**
- IP address of the accounting server

If requests from more than one client are to be forwarded to the same accounting server, then after it has been configured for any one client, it can be selected from **Select from** for the others. All the information for the selected accounting server, including its shared secret, is copied and instated for this client.

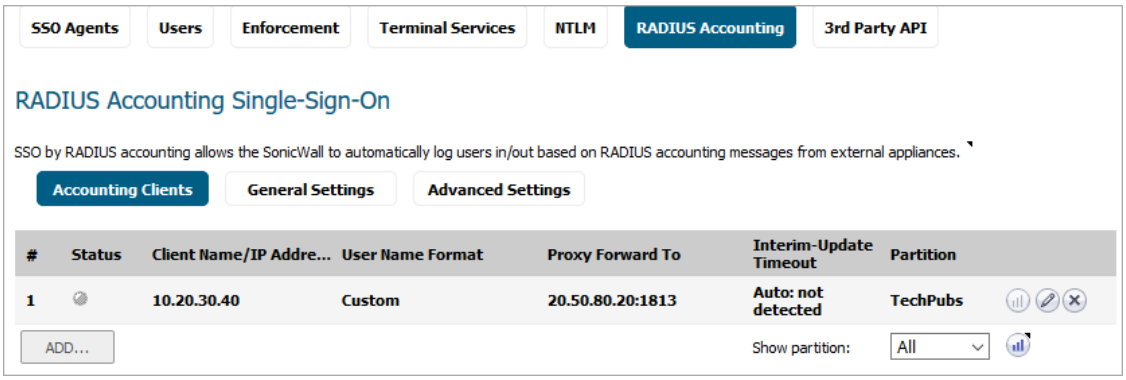
12 In the **Timeout (seconds)** field and **Retries** field, enter the timeout period in seconds and the number of retries. The default for **Timeout (seconds)** is **10** seconds, and the default for **Retries** is **3**.

To determine which users have logged out, the SonicWall Security Appliance polls the SSO Agent by sending requests to multiple logged-in users in a single request message to the SSO Agent. To configure the number of user requests the Security Appliance can send in a single request message to the Test tab

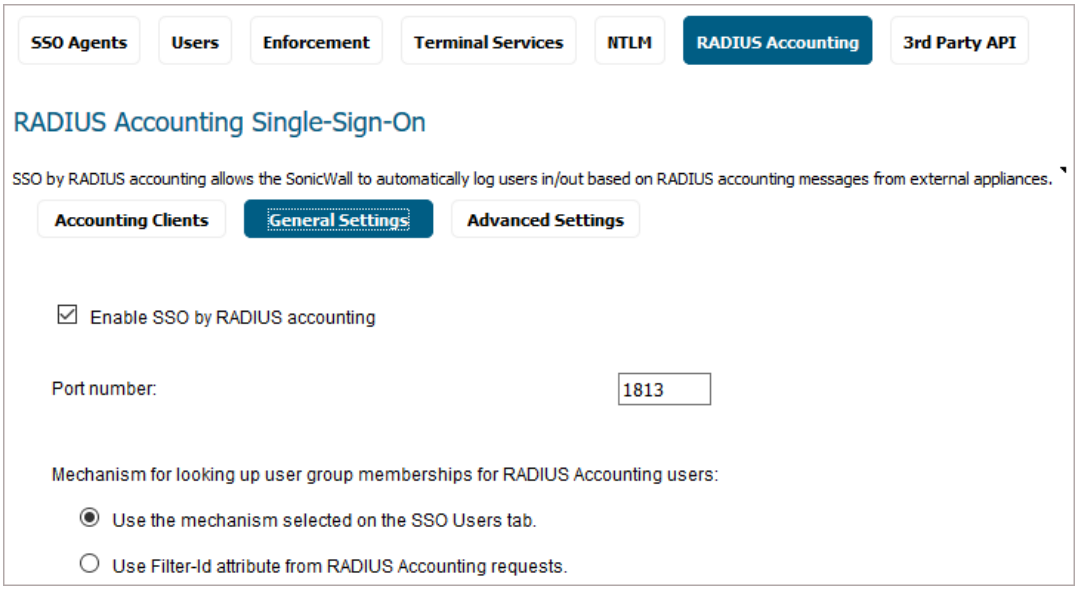
13 Select how the RADIUS accounting messages are forwarded from this client, either:

- **Try next on timeout**
- **Forward to all**

14 Click **SAVE**. The **Accounting Clients** table is updated.



15 Click **General Settings**.



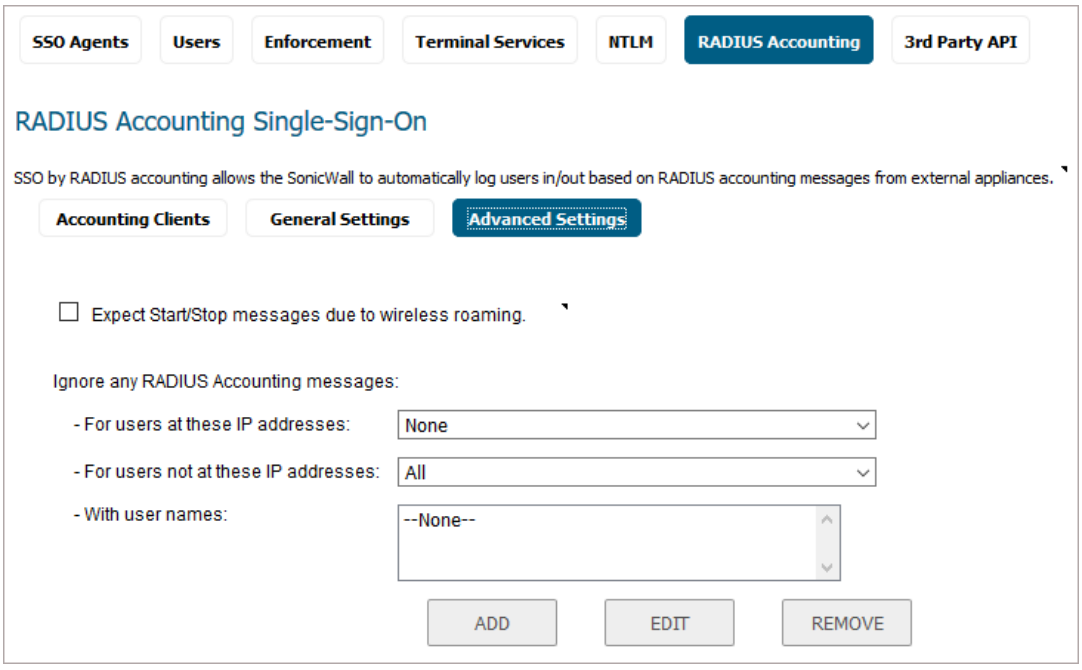
16 Enable SSO or RADIUS accounting by selecting the **Enable SSO or RADIUS accounting** checkbox. This option is selected by default.

17 Specify the port in the **Port number** field. The default port is **1813**.

18 For Mechanism for looking up user group memberships for RADIUS Accounting users, choose:

- **Use the mechanism selected on the SSO Users tab.** (default)
- **Use Filter-Id attribute from RADIUS Accounting requests.**

19 Click **Advanced Settings**.



20 To have the appliance track RADIUS Accounting messages for Start/Stop messages, select the **Expect Start/Stop messages due to wireless roaming** checkbox. This option is not selected by default.

RADIUS Accounting clients send Start/Stop messages to notify the Security Appliance of users connecting/disconnecting. If those clients are or use wireless access points, then the wireless users could roam between access points, which might cause them to generate spurious Start/Stop messages as the user connects to a new access point and disconnects form the old one. These roaming Start/Stop messages could interfere with the SSO authentication process, which normally processes Stop messages as notifications of user logout.

If this option is enabled, then the Security Appliance tracks the RADIUS Accounting messages to look for this Start/Stop sequence. If the sequence is found, then the Security Appliance considers the Stop messages as indications of roaming rather than as notifications of user logout.

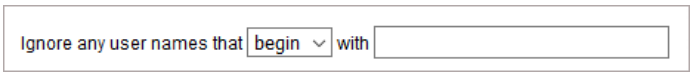
That is, the Security Appliance assumes Start/Stop messages are due to roaming switch-over between access points if those messages:

- Are received (in any order): a Start message for a currently connected user that indicates the same user is at a different access point, along with a Stop message from the previous location.
- Occur together within the specified time.

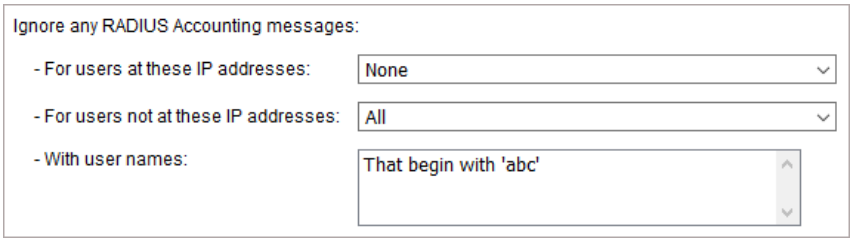
NOTE: The maximum switch-over time should allow for the RADIUS Accounting message possibly getting dropped and retransmitted. The recommended time is the same as the timeout multiplied by the maximum retries for the RADIUS Accounting clients.

21 To have the Security Appliance ignore any RADIUS Accounting messages for users, under **Ignore any RADIUS Accounting messages**:

- At specific IP addresses, select an address object or address group from **For users at these IP addresses** or create a new address object or address group. The default is **None**.
- Not at specific IP addresses, select an address object or address group from **For users not at these IP addresses** or create a new address object or address group. The default is **All**.
- With specific user names:
  1) Click **ADD**. The **Add RADIUS Accounting User Name Exclusion** dialog displays.



  2) From the Ignore any user names that drop-down menu select
     - **begin**
     - **end**
  3) Enter the user name in the **with** field.
  4) Click **APPLY**. The entry is added to the list.



  To edit an entry, select it and then click **EDIT**.

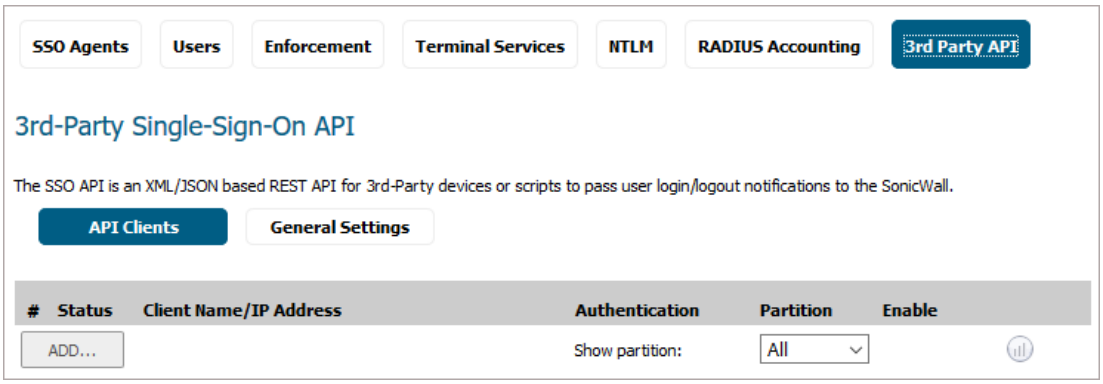
  To remove an entry, select it and then click **REMOVE**.

3rd Party API Page

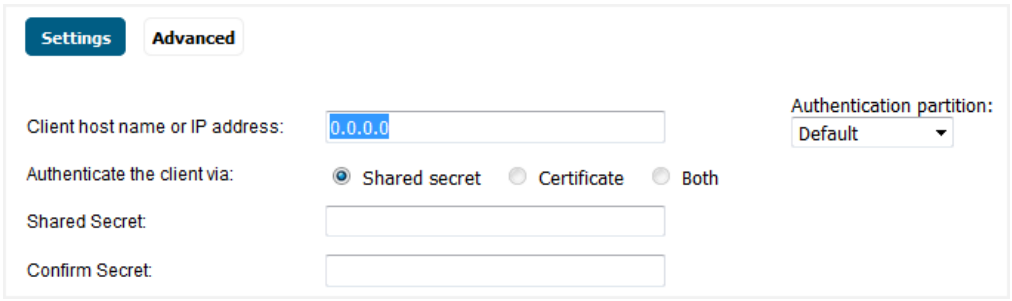
The SSO API is an XML-/JSON-based REST API for third-party devices or scripts to pass user login/logout notifications to the firewall.

To configure third-party API:

1 Click **3rd Party API**.



2 Under the **API Clients** table, click **ADD**. The **Add API Client** dialog displays.



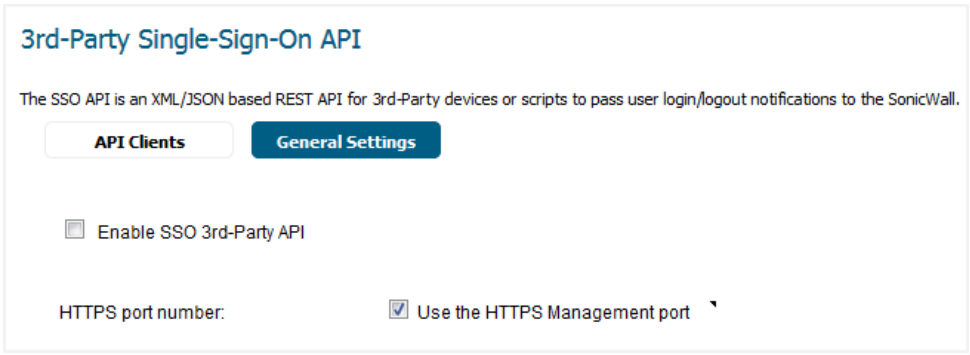
TIP: If partitioning is enabled, the **Select partition** dialog displays over the **Add RADIUS Accounting Client** dialog.

1 Choose the partition to which to add the new agent.
2 Click **OK**.

3 Configure the options.

4 Click **SAVE**. The client is added to the **API Clients** table.

5 Click **General Settings**.



6 Select **Enable SSO 3rd-Party API**. This option is not selected by default.

7 To have user login/logout notifications on the API sent to:

- The same TCP port number used for HTTPS web management, select **Use the HTTPS Management port**. This option is selected by default.
- A separate, dedicated TCP port number, clear Use the HTTPS Management port; a field in which to enter the port number displays.



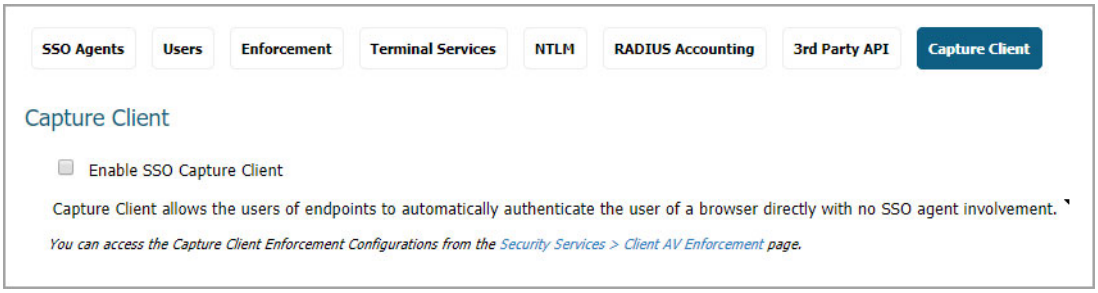
IMPORTANT: You must use a separate port number if client certificates are used to authenticate API clients.

TIP: A separate, dedicated TCP port number is recommended for larger installations to prevent a possible high rate of API requests from potentially adversely affecting web management.

8 Click **APPLY**.

Capture Client Page

1 Select the **Enable SSO Capture Client** checkbox to enable user authentication via SonicWall Capture Client by client PCs configured in the Client AV Enforcement lists, with any zone.



This allows the user to access the internet via their browser with no SSO agent involvement.

It is not necessary to enable this checkbox if you just want to use Client AV Enforcement with Capture Client.

For information about Client AV Enforcement with Capture Client, refer to the *SonicOS 6.5 Security Configuration* administration documentation, available at https://www.sonicwall.com/support/technical-documentation.

Test Page

1 To test the agent settings you configured, click **TEST**.

IMPORTANT: Performing tests on this page applies any changes that have been made.

You can test the connectivity between the appliance and an SSO agent or TSA. You can also test whether the SSO agent is properly configured to identify a user logged into a workstation.

2 If you have multiple agents configured, select the SSO agent or TSA to test from the **Select agent to test** drop-down menu. The drop-down menu includes SSO agents at the top, and TSA's at the end under the heading **--Terminal Server Agents--**.

3 Select the type of test to perform:

- **Check agent connectivity** radio button – Tests communication with the authentication agent. If the Security Appliance can connect to the SSO agent, the message **Agent is ready** displays. If

testing a TSA, the **Test Status** field displays the message, and the version and server IP address are displayed in the **Information returned from the agent** field.

- For SSO agents only, select **Check user**, enter the IP address of a workstation in the **Workstation IP address** field. This tests if the SSO agent is properly configured to identify the user logged into a workstation.

  TIP: If the messages **Agent is not responding** or **Configuration error** display, check your settings and perform these tests again.

4 Click **TEST.**

5 When you are finished with all Authentication Agent configuration, click **OK**.

Configuring RADIUS Accounting for SSO

RADIUS accounting for Single Sign-On is configured on the **Users > Settings** page.

To configure RADIUS accounting for SSO:

1 Display the **MANAGE | System Setup | Users > Settings** page.

2 Click **CONFIGURE SSO**. The **SSO Authentication Configuration** dialog appears.

3 Click the **RADIUS Accounting** tab. For the procedure to configure RADIUS Accounting, see RADIUS Accounting Page on page 180.

4 Click **APPLY**.

Advanced LDAP Configuration

If you selected **Use LDAP** to retrieve user group information on the **Users** page as described in Configuring SonicOS to Use the SonicWall SSO Agent on page 167, you must configure your LDAP settings.

To configure LDAP to retrieve user group information:

1 Navigate to **MANAGE | System Setup | Users > Settings**.

2 Click **CONFIGURE SSO**. The **SSO Configuration Authentication** dialog displays.

3 Click **Users**.

4 Click **CONFIGURE** next to the **Use LDAP to retrieve user group information** option. The **LDAP Configuration** dialog displays.

5 For configuring LDAP, see Configuring the SonicWall for LDAP on page 149.

9

Managing Authentication Partitions

Topics:

- About Authentication Partitioning on page 190
  - About User Authentication Partitioning on page 191
  - About Subpartitions on page 192
  - About Inter-Partition User Roaming on page 194
  - About Authentication Partition Selection on page 195
  - About Extended Support for Multiple LDAP Servers on page 198
  - Per-Partition DNS Servers and Split DNS on page 198
  - About RADIUS Authentication on page 198
  - Upgrading from a Non-Partitioned Configuration on page 198
- Configuring Authentication Partitions and Policies on page 199
  - Displaying and Filtering Users/Partitions on page 199
  - Configuring and Managing Partitions on page 201
  - Configuring Partition Selection Policies on page 215
  - Configuring Servers, Agents, and Clients for Authentication Partitioning on page 219

About Authentication Partitioning

Topics:

- About User Authentication Partitioning on page 191
- About Subpartitions on page 192
- About Inter-Partition User Roaming on page 194
- About Authentication Partition Selection on page 195
- About Extended Support for Multiple LDAP Servers on page 198
- Per-Partition DNS Servers and Split DNS on page 198
- Upgrading from a Non-Partitioned Configuration on page 198

About User Authentication Partitioning

NOTE: For definitions of terms used in this section, see Terms and acronyms.

SonicWall Security Appliances provide a mechanism for LDAP, RADIUS, and/or Single-Sign On (SSO) authentication in an environment where you manage multiple non-interconnected domains. Such an environment needs users in a particular domain to be authenticated through the specific:

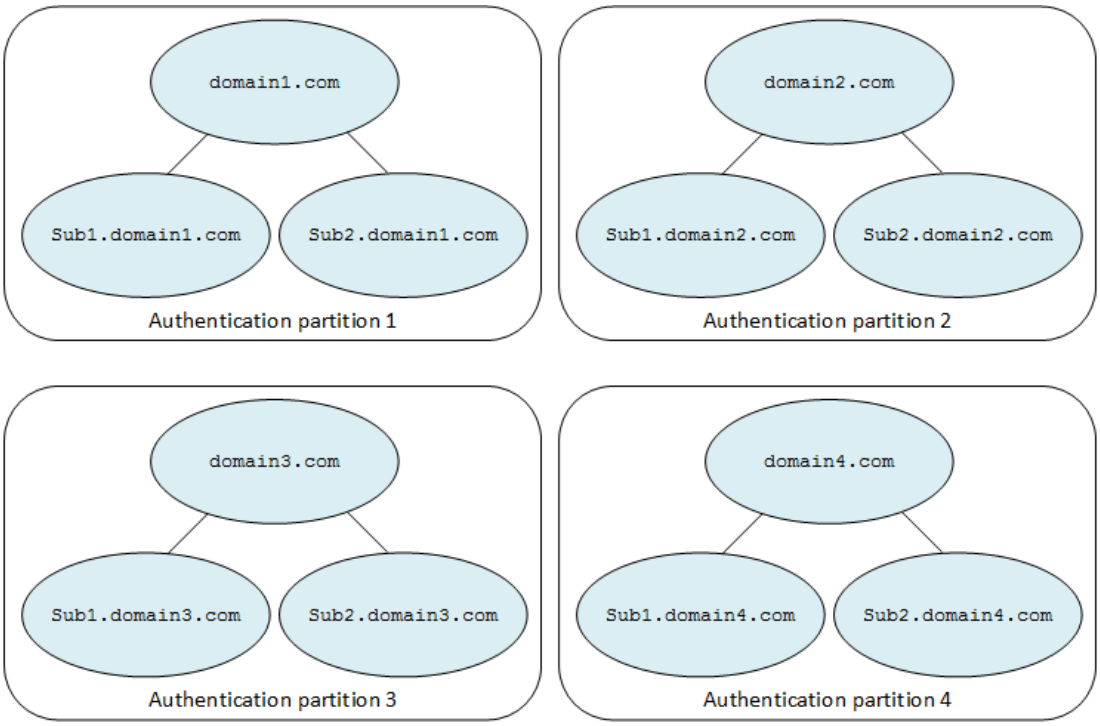
- LDAP/RADIUS server for that domain
- SSO agent(s) located in that domain

The mechanism for such environments is User Authentication Partitioning, which means:

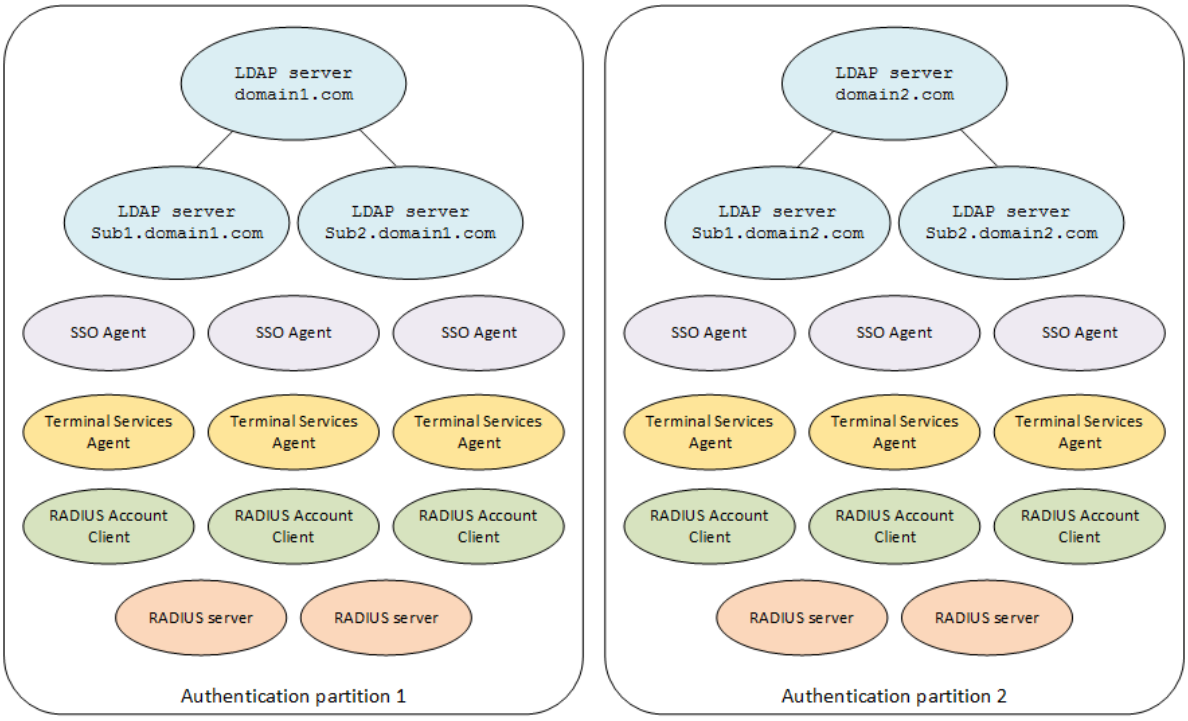
- First, partitioning your network(s) into separate partitions, each with its own authentication servers/agents/clients.
- Then, authenticating each user against the relevant authentication devices (servers/agents/clients) according to the authentication partition in which the user is located. The user's partition is selected by either:
    - Matching the user's domain names against those configured in the domains.
    - If the user's domain names are not available, basing their physical location as set by the partition-selection policies.

An authentication partition typically corresponds to one or more domains; for example, in a Windows domain, a partition usually corresponds to an Active Directory forest. Each partition has separate LDAP servers, RADIUS servers, SSO agents, and/or Terminal Service agents (TSAs). See Authentication partitions and Installation with central and remote sites.

Authentication partitions



Partition contents



Terms and acronyms

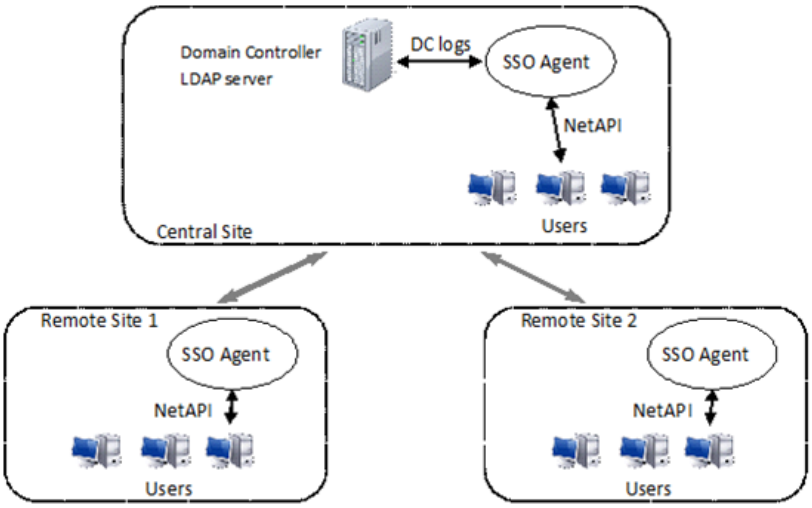
Authentication Partition	A part of a network with its own authentication servers/agents/clients, separate from those in other parts of the network
DC	Domain controller
LDAP	Lightweight Directory Access Protocol
RADIUS	Remote Authentication Dial In User Service
SSO	Single Sign On
TSA	Terminal Services Agent

About Subpartitions

Authentication partitions select the LDAP servers, RADIUS servers, SSO agents, and TSAs used to authenticate particular users. In addition to assigning the servers and agents to a partition, there can be instances where it is necessary to then further assign certain of them to different subsets of the users in the partition. Subpartitions allow assigning particular agents for certain subsets of a partition's users if specific ones need to be used for them. If an authentication partition is set as a subpartition of another one, then agents specific to the top-level, or parent, authentication partition's users can be assigned to the subpartition. The subpartition's agents are used when relevant, but the servers and agents of the parent partition can be used as appropriate.

For example, an installation with central and remote sites has the domain controllers (DCs)/LDAP server located at the central site. Access policies, however, prevent an SSO agent located at the central site from accessing the remote user computers. For SSO with NetAPI or WMI to work in this topology, one or more SSO agents must be placed at each remote site in addition to those located at the central site; see Installation with central and remote sites. For NetAPI/WMI, the SSO agent talks directly to the users' computers while user identification in the DC security log uses the SSO agent of the domain controller(s).

Installation with central and remote sites



Subpartitioning of top-level partitions solves the problems of:

- Telling the appliance that the SSO agents at each individual site are used for NetAPI or WMI identification of the users located there.
- Using the SSO agents and LDAP servers at the central site for DC logs identification and user group lookups for all of the users across all the sites.

Each of the remote sites can be configured as a subpartition of the central site, with the SSO agent(s) at a remote site assigned to the subpartition. One or more partitions can be configured as subpartitions of a parent partition, with different selection policies defining the locations of the user subsets for each. In Installation with central and remote sites, the entire installation is one partition, and the remote sites are each a subpartition within that partition. After selecting the relevant agent from the subpartition to identify a user at the remote site, the user's group memberships are subsequently looked up through the LDAP servers of the parent partition.

Some special cases for subpartitions are:

- LDAP servers cannot be assigned to a subpartition. If a subpartition corresponds to a sub-domain that has its own LDAP servers, those servers should be assigned to the parent partition. The LDAP server manages referring requests to the sub-domain.
- For RADIUS servers, a subpartition uses either servers assigned to it or those of the parent partition, but not both. If RADIUS servers are assigned to a subpartition, they are used for users located in it; otherwise, those of the parent partition are used.
- With SSO agents using NetAPI or WMI versus reading from domain controller logs, if there are agents in both a subpartition and its parent, then only the subpartition's own SSO agents are used for NetAPI/WMI identification of users located in the subpartition; the subpartition's SSO agents involve direct access to the user's computer. Domain controller log reading is done by SSO agents in the parent and/or subpartitions, if the subpartition's SSO agent's are configured for that.

Operation of Servers, Agents, and Clients with Subpartitions

Generally, any servers, agents, and/or clients assigned to a subpartition are used for the users located in it, but in addition, certain servers, agents, and/or clients of the parent partition are used; see Use of servers, agents, and clients with subpartitions.

Use of servers, agents, and clients with subpartitions

Server, agent, client	Use
LDAP servers	Only assignable to a top-level partition, not to a subpartition. If a subpartition corresponds to a sub-domain with its own LDAP server(s), those should be assigned to the parent partition, and the LDAP's referral/reference mechanism refers requests to the sub-domain's servers. Where a sub-partition corresponds to a sub-domain with its own LDAP server(s), however, you may find it more logical to assign those servers to the sub-partition, and that is allowed. Servers assigned to a sub-partition are linked internally into the parent partition.
RADIUS servers	A subpartition uses either the RADIUS servers assigned to it or those of the parent partition, but not both. If RADIUS servers are assigned to a subpartition, then they are used for users located in it; otherwise, those of the parent partition are used.
SSO agents	When using NetAPI or WMI, the agents need to be located where they can access the user PCs directly, and when reading from DC (domain controller) logs, they need to access the DCs. An SSO agent can be configured to do both activities. When using both DC logs and NetAPI/WMI, the SonicWall Security Appliance controls which are used and in what order. The Security Appliance: 1 Has the agents look for the user in the DC logs read from each DC. 2 If the user is not found in the logs, makes a separate follow-up request to try NetAPI/WMI. When using subpartitions, this mechanism operates as follows to identify a user located in a subpartition: 1 If any SSO agents assigned to the subpartition have DC logs enabled, then requests are sent to those SSO agents to look for the user in their DC logs. 2 If the user was not identified in Step 1, then if any SSO agents assigned to the parent partition have DC logs enabled, requests are sent to those SSO agents to look for the user in their DC logs. 3 If the user was not identified in Step 2, then if any SSO agents assigned to the subpartition have NetAPI or WMI enabled, requests are sent to one of those to identify the user. **NOTE:** NetAPI/WMI is not attempted through SSO agents in the parent partition for users located in a subpartition. If there are no agents assigned to the subpartition with NetAPI or WMI enabled, then authentication is not attempted.
TSAs and RADIUS accounting clients	The partition to which these agents/clients send user are assigned affects only the choice of LDAP server to use for user group lookups. As the LDAP servers of the parent partition are also used for all of its sub-partitions, the TSAs and RADIUS accounting clients can be assigned to either. The only difference would be which partition is displayed for their users, and the users are assigned to them based on their physical location. **NOTE:** This is applies only when a domain is not supplied.

About Inter-Partition User Roaming

Users who log into a domain in one partition are able to roam and connect from the physical network of a different partition if the network topology has been set up to allow them to access their own partition's domain server from the login partition. If an SSO agent is used in such a case, the appliance selects the SSO agent(s) of the local partition based on the user's physical location, not the agent(s) of their home partition.

The SSO agents of the local partition are not able to identify the roaming user from domain controller logs because the agents do not read from the correct domain controllers. The agents can identify the roaming user through NetAPI or WMI if they have the correct privileges, which requires Windows inter-domain trust. Thus, when the Security Appliance gets the username from an SSO agent, the Security Appliance checks the partition where the specified domain is located and allows that to override the partition initially selected based on the user's physical location.

The process to identify the roaming user and set their access permissions is:

1. A user who is logged in to domain-1 (in partition-1) connects from a subnet in partition-2; the user's partition is initially recorded as partition-2.
2. If the partition-2 agents are reading domain controller logs, then requests are first sent to check those. These requests fail to find the user, who is not logged into partition-2's domain.
3. A request is sent to an SSO agent in partition-2 to try NetAPI. The agent does this, and identifies the user as one from domain-1.
4. The Security Appliance sees that domain-1 is in partition-1, and switches the user's partition to partition-1. The Security Appliance then looks up the user's group memberships through the LDAP servers in partition-1.

About Authentication Partition Selection

Topics:

- Selection Policies on page 195
- Remote Users on page 196
- Appliance Notification of User Logins on page 196
- Web User Login on page 197

Selection Policies

Network topology can affect how SonicOS locates authentication partition users on the network. SonicOS provides several options for locating and selecting the user's partition.

Selection options

When selected by	Each authentication partition
IP address	Corresponds to a set of IP addresses selected through an address object (network, range, or group) in its configuration.
Network interface	Corresponds to the networks that are accessed through one or more interfaces that are selected in its configuration.

Selection options

When selected by	Each authentication partition
Network zone	Corresponds to one or more network zones that are selected in its configuration.
Username domain component	Is a member of one or more domains and is chosen by matching the domain name given at login time by the user. This option requires the user log in with a qualified name, for example, `domain\user` or `user@domain.com`. When a domain name is given, this option overrides the above location-based options. This option should be used for authenticating GVC, L2TP, and SSL VPN client users; see Remote Users on page 196. **NOTE:** For SSO agent authentication, one of the location-based options should be used because the SonicWall Security Appliance needs to derive the partition at the start of the process to select the SSO agent(s) to use, and at that point the Security Appliance does not yet have the user's login name; see About Inter-Partition User Roaming on page 194.

These options are configured as a set of separate selection policies, with one or more policies set per partition, to define how to select that partition. During user authentication, if no domain is given, then the partition is selected by matching the zone, interface, and IP address against the configured policies in a manner very similar to access rule matching. A default selection policy specifying the default partition is a catch-all for anything that does not match the explicitly configured policies. The default partition is initially named Default, but it can be renamed or the default selection policy can be set to another partition, after which the auto-created Default can be deleted.

Remote Users

Selection of the authentication partition to use for GVC/L2TP clients and SSL VPN users is handled differently as these remote users are connecting into the authentication partition, not coming from it. The Security Appliance needs to know the authentication partition for these users to select the correct LDAP server to look up their user group memberships, and from that, the subnets they can access. There are two options for authenticating remote users:

- Use selection by username domain component and require that remote users give a qualified name including the domain.
- Have multiple WAN interfaces and/or WAN zones and have the users for each authentication partition connect to a different public IP address. The WAN interface or zone that a remote user is coming through is then used to select the authentication partition without requiring that the remote users give qualified usernames.

  NOTE: For GVC/L2TP users, having separate WAN zones allows different Group VPN policies for each zone, thereby potentially enforcing access more securely to the correct authentication zone.

When there are multiple WAN interfaces, it is possible to set partition-selection policies to select the partition for remote access through each WAN interface. If there is only one WAN interface, then it is possible to set a special selection policy to select the default partition for remote access when that cannot be derived from the supplied username.

NOTE: If that selection policy is not set, then remote access users need to supply a qualified username unless the servers that authenticate them are assigned to the partition that is selected by default.

Appliance Notification of User Logins

If the SonicWall Security Appliance is notified of user logins by an agent/client, but does not send identification requests for them (for example, Terminal Services, RADIUS accounting, and login notifications from an SSO agent reading DC logs), the Security Appliance does not need to know the authentication partition for selection

of the agent/client as it does for sending requests to SSO agents. To select the correct LDAP server to look up their user group memberships, however, the Security Appliance does need to know the authentication partition for these users. That selection is done from the username domain component, when present, or by manually assigning each such agent/client to an authentication partition.

Web User Login

When users log in through the SonicWall Security Appliance's web login portal, they could conceivably use any account name, no matter where they are coming from. Normally, the authentication partition is selected based on where users log in from (see Selection Policies on page 195), but if the user gives a username that includes the domain, then they are able to override that by logging in with a qualified username that includes the domain to select the authentication partition.

CLI Login

When a user logs in through CLI using the built-in admin account, partitioning is not relevant because that is always authenticated locally. But when additional administrator accounts that are authenticated through LDAP or RADIUS are used, then the partition does need to be known to select the server to authenticate it. There are three separate cases for this:

Login on the console port	There is no IP address from which to derive the partition, and so when that is needed, the user needs to log in with a qualified username.
Local SSH connection from inside the firewall	The authentication partition where the user is located is selected through the source IP address of the SSH connection, as per Selection Policies on page 195.
Remote SSH connection from outside the firewall	The partition selection is not based on the user's location, but it is optionally possible to have it selected according to the WAN interface to which they are connected, as per remote client users; see Selection Policies on page 195.

If selection by username domain component has been configured (see Selection Policies on page 195), then in all cases the user is able to override that by logging in with a qualified username that includes the domain from which the authentication partition is selected. It is also possible to set a special selection policy to select the default partition for console port login when that cannot be derived from the supplied username.

NOTE: If the selection policy is not set, then users need to supply a qualified username to log in on the console port unless the servers that authenticate them are assigned to the partition that is selected by default.

Per-Partition User Authentication Settings

There can be in some cases a need to set certain settings governing user authentication differently in different partitions. For example if one partition has only RADIUS servers and another has only LDAP servers, then for user authentication RADIUS must be selection in the first partition and LDAP in the other.

By default, all such settings apply globally and are limited to the user authentication method and the single sign on methods. These settings are set only for top-level partitions; for sub-partitions, the authentication settings of their parent partitions apply.

About Extended Support for Multiple LDAP Servers

Partitioning requires multiple LDAP servers. Multiple primary LDAP servers can be configured, one for each authentication partition, plus a list of additional servers for each. For more information about multiple LDAP servers and how to configure them, see About Extended Support for Multiple LDAP Servers on page 162.

Per-Partition DNS Servers and Split DNS

With or without authentication partitions, it is usually necessary to use a domain's own DNS servers to resolve the names of devices in the domain, and occasionally there can also be a need to use different external DNS servers to resolve external host names. Multiple authentication partitions, however, usually require using different DNS servers to resolve the host names in the different partitions.

The DNS Proxy with Split DNS feature allows configuring different DNS servers associated with different domain names. The feature is separated from DNS Proxy so it can be used directly by the Security Appliance to resolve the names of devices in domains without needing to enable DNS Proxy, including for multiple, unrelated domains with authentication partitioning. For more information about Split DNS, see Managing Authentication Partitions on page 190.

About RADIUS Authentication

With RADIUS authentication there are some additional considerations as the SonicWall Security Appliance is not guaranteed a way to derive the user's domain, nor the domains of the user groups returned in RADIUS attributes, as it is with LDAP. So the Security Appliance can find the domains for selecting the correct domain user and user group objects, the Security Appliance tries the following to learn a user's domain with RADIUS authentication:

1. Have users give a qualified user name that includes the domain when they log in. If the RADIUS server returns the user groups in RADIUS attributes (Filter-ID or the SonicWall vendor-specific ones), then configure it to return them giving the fully qualified names of the groups, including their domain.
2. Use LDAP for the user group lookup after authenticating the user through RADIUS (which is the preferred method). Then, if the user does not give the domain with their user name it can be learned from the LDAP search to find their user groups.
3. Failing either of those, the domain can be looked up from the authentication partition when the user is logging in from an IP address physically located in it, but this can only definitively give the user's domain if there is only one domain per partition; hence, to use this method, it is necessary to have a separate sub-partition for every sub-domain.

   NOTE: This does not work for cross-domain user group memberships.

In summary, the best option with RADIUS authentication is to use LDAP for the user group lookup. If that is not possible (no LDAP server), then the next best option is to have the RADIUS server return qualified user group names in RADIUS attributes.

If none of these can be used to derive the domains of user groups returned from RADIUS, then it is necessary to configure the user/user group objects to match in any domain.

Upgrading from a Non-Partitioned Configuration

When starting from an existing configuration without authentication partitioning, when partitioning is enabled:

- A single authentication partition named **Default** is created with all the existing servers, agents, and clients in it.
- A single default partition selection policy is configured to select the **Default** partition as the default partition for everything.

From that base, new partitions can be added, and the relevant servers, agents and clients easily moved to them from the default partition, or new ones added.

Configuring Authentication Partitions and Policies

The **Users > Partitions** page enables you to create of a list of authentication partitions and of the policies to select them. For each partition, you can configure:

- A name for the authentication partition (for example, the name of the domain or forest to which it corresponds).
- The domains that the partition encompasses.
- How the authentication partition is to be selected for users (for example, configured as separate partition selection policies).

Before configuring the authentication partitions and partition selection policies, you can determine the users locations in the partitions from the **Monitor > Current Status > User Sessions > Active Users** page.

When authentication partitions have been configured, a selection is added in the various server/agent/client configurations so the authentication partition can be selected when adding/editing a server, agent, or client. You configure the servers, agents, and clients from the **MANAGE | System Setup > Users > Settings** page.

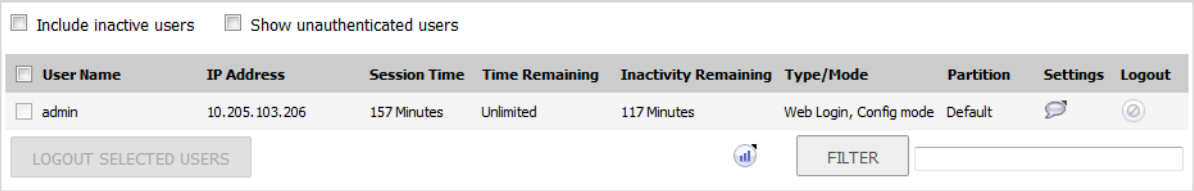
Topics:

- Displaying and Filtering Users/Partitions on page 199
- Configuring and Managing Partitions on page 201
- Configuring Partition Selection Policies on page 215
- Configuring Servers, Agents, and Clients for Authentication Partitioning on page 219

Displaying and Filtering Users/Partitions

The **Monitor | Users Sessions > Active User** page shows the partition where each user is located.

NOTE: For more information about this page, see *SonicOS 6.5 Monitoring* for your SonicWall Security Appliance.



Topics:

- Viewing User Information on page 200

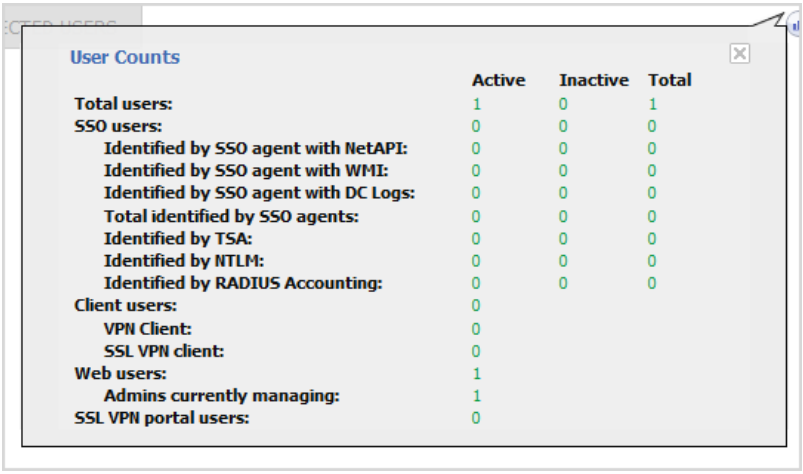
- Filtering Users on page 200

Viewing User Information

You can view the number of users by various categories:

- Active/Inactive
- SSO users by how identified
- Client users by type of client
- Web users
- SSL VPN portal users

To view this information, click the **Statistics** icon beneath the **Active User Sessions** table. The **User Counts** popup dialog displays:



Filtering Users

The **Filter** field allows filtering of a partition so that just the users in a selected partition can be shown. You search for users by specifying one or more full or partial usernames, domains, IP addresses, and/or type of user. You exclude users by prefixing entries with an exclamation point (!). When combining strings, to match:

- Any of the listed entries, separate the entries with a comma; that is, `a,b` includes users that match either a or b
- All of the listed entries, separate the entries with a semicolon (;); that is `a;b` includes users that match both a and b

To search for terminal server users, enter `user-num=`*`usernumber`*. The `type` filter matches text in the **Type/Mode** column, including any shown on mouse-over in that column. IPv6 addresses are supported, but only for full matching; for example, `ip=2012::1`, `!ip=2012::1`, or in combinations of other entries as shown in Filter examples.

Filter examples

```
name=bob                        name=bob,john,sue               domain=mydomain
ip=192.1.1.1                    ip=192.1.1.1,192.1.1.2          ip=192.1.1.0/24
type=config mode                type=sso,web                    type=sso;netapi
type=sso;from logs on domain controller 192.1.1.10
partition=somePartition         group=Trusted Users
```

Filter examples

```
name=bob;ip=192.1.1.1 (to match both name and IP address)
!name=bob !ip=192.1.1.1 (to exclude users)
```

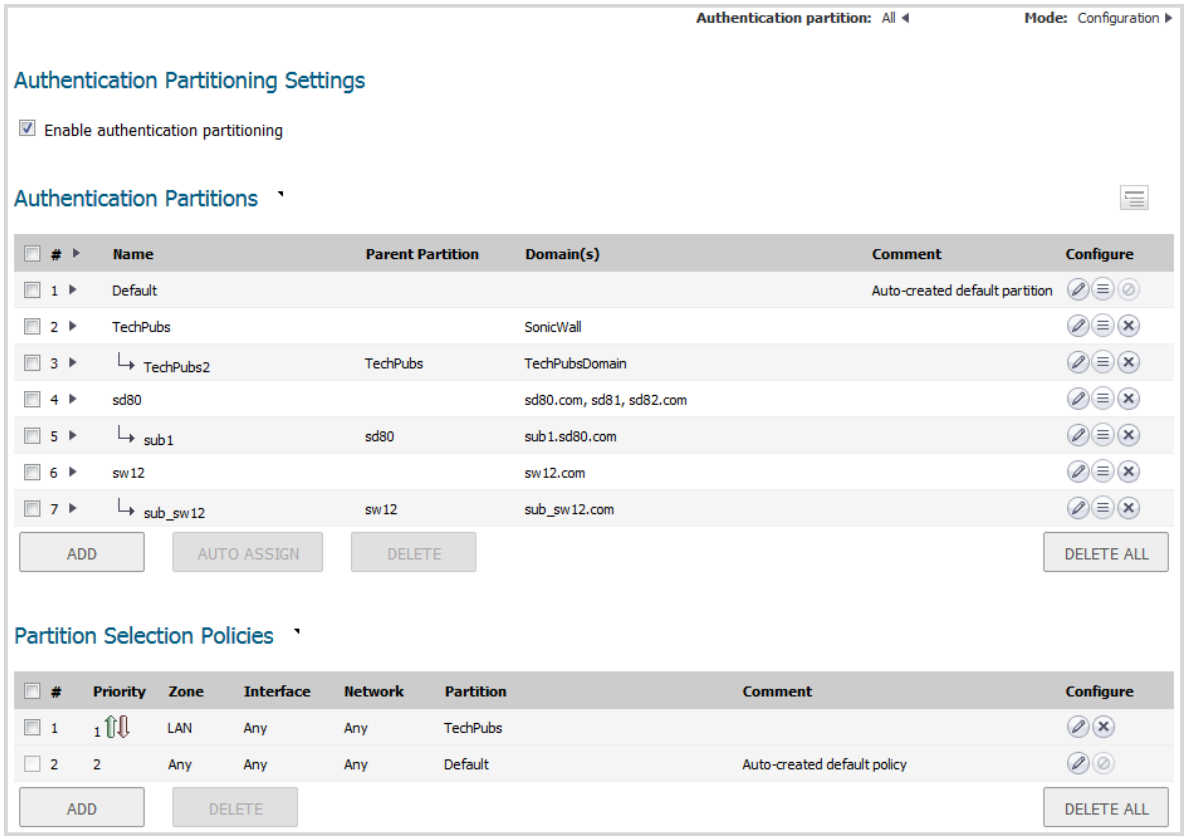
You also can use just simple strings; for example: `bob 192.1.1.1 mydomain`

Configuring and Managing Partitions

Topics:

- **Users > Partitions Page** on page 201
- Enabling/Disabling Authentication Partitioning on page 205
- Adding Partitions and Subpartitions on page 206
- Deleting Partitions and Subpartitions on page 208
- Assigning Servers, Agents, and Clients on page 209
- Editing Partitions on page 212

Users > Partitions Page



The **MANAGE | System Setup > Users > Partitions** page has three sections:

- Authentication Partitioning Settings Section on page 202
- Authentication Partitions Section on page 202

- Partition Selection Policies Section on page 204

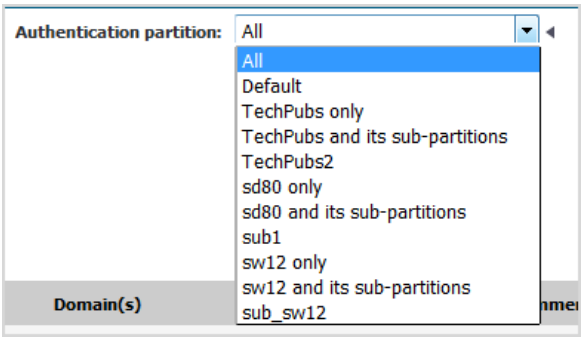
Authentication Partitioning Settings Section

This section enables/disables authentication partitioning. If authentication partitioning is disabled, the other sections do not display.



When authentication partitioning is enabled, the two sections, **Authentication Partitions** and **Authentication Selection Policies**, also display.

Also displayed when partitioning is enable is the **Authentication partitioning** drop-down menu at the top of the page from which you can select the partition to which the settings in the **MANAGE | System Setup > Users > Settings** and **MANAGE | System Setup > Users > Local Users & Groups** pages also apply. The default is **All**, that is, the settings apply to all partitions.

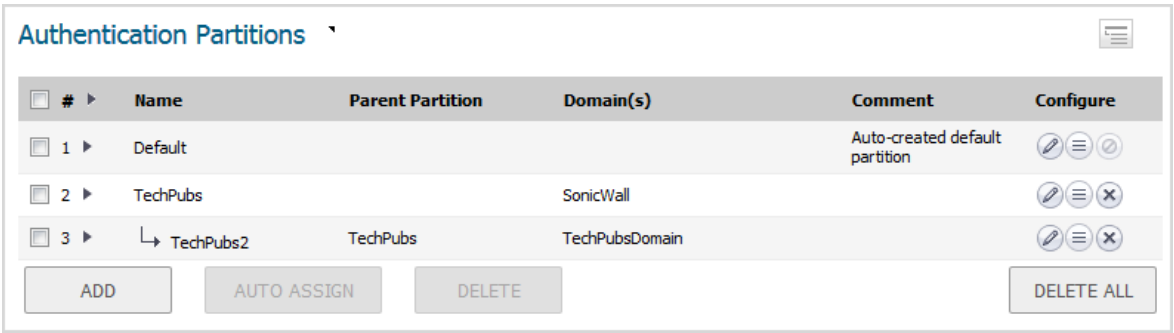


Authentication Partitions Section

NOTE: This section displays only when authentication partitioning is enabled.

This section displays a table of authentication partitions and allows you to create, edit, delete, and manage the partitions. The partitions you configure here control which authentication servers are used for which users.

You can expand a partition's tree to show the servers, agents, and clients assigned to it.

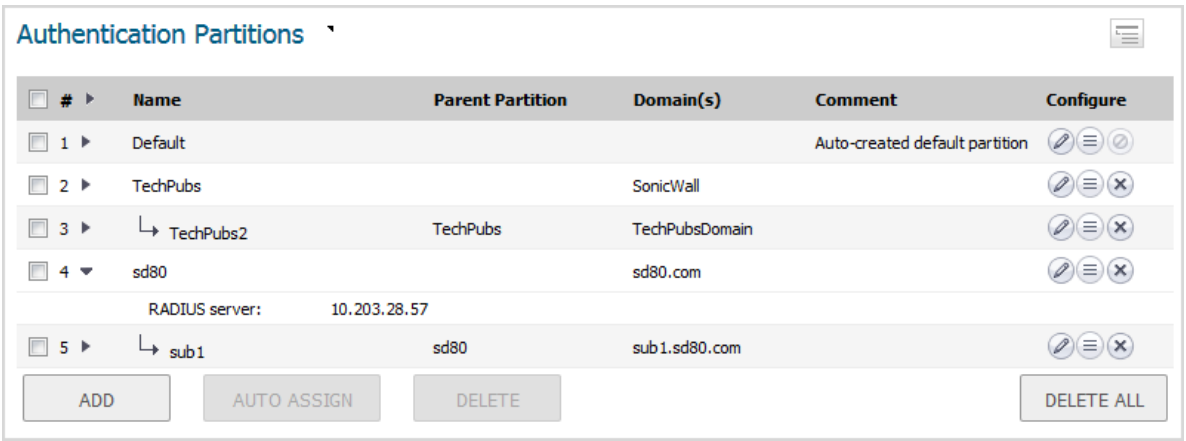


Group subpartitions icon	Toggles between grouping subpartitions with their parent authentication partitions or ungrouping the subpartitions and sorting them with the top-level partitions. **NOTE:** Grouped subpartitions are displayed immediately after their parent partition, with a **Link** ↳ icon denoting them as subpartitions.
Selection checkbox	Allows you to select one or more partitions and/or subpartitions in the table. Selecting the checkbox in the table heading selects all entries except the **Default** partition.
Name	Specifies the name of the authentication partition. Subpartitions are indicated by a **Link** ↳ icon in front of the name.
Parent Partition	Specifies the parent authentication partition for subpartitions. This column is blank for parent partitions.
Domain(s)	Specifies the domain(s) to which the partition or subpartition belongs. This column is blank for the **Default** partition.
Comment	Displays the comment included when the partition was added. The comment for the **Default** partition is **Auto-created default partition**.
Configure	Displays the **Edit**, **Selection** ≡, and **Delete** icons for the partition **NOTE:** The **Edit** and **Delete** icons are dimmed for the **Default** partition.
Add	Displays the **Add an authentication partition** popup dialog for adding an authentication partition or subpartition.
Auto assign	Assigns any unassigned LDAP servers, RADIUS servers, SSO agents, TSAs, and RADIUS accounting clients to the relevant partitions automatically, based on their IP addresses or host names. **NOTE: Auto assign** and **Delete** are dimmed unless at least one partition or subpartition has been selected.
Delete	Deletes the selected authentication partition(s) or subpartition(s). **NOTE:** You cannot delete the **Default** partition.
Delete All	Deletes all partitions and subpartitions from the table except the **Default** partition.

There is always one authentication partition in this table, the auto-created **Default** partition. You cannot delete this partition. You can, however, edit it and select servers, agents, and clients for it as well as subpartitions. If you disable authentication partitioning, all LDAP servers, SSO agents, TSAs, and RADIUS accounting clients are reassigned to the **Default** partition; when you re-enable authentication partitioning, you must reassign them. RADIUS servers are not affected and remain with their assigned partitions.

Expanding Trees

Expanding an authentication partition's tree shows the servers, clients, and agents assigned to the partition:

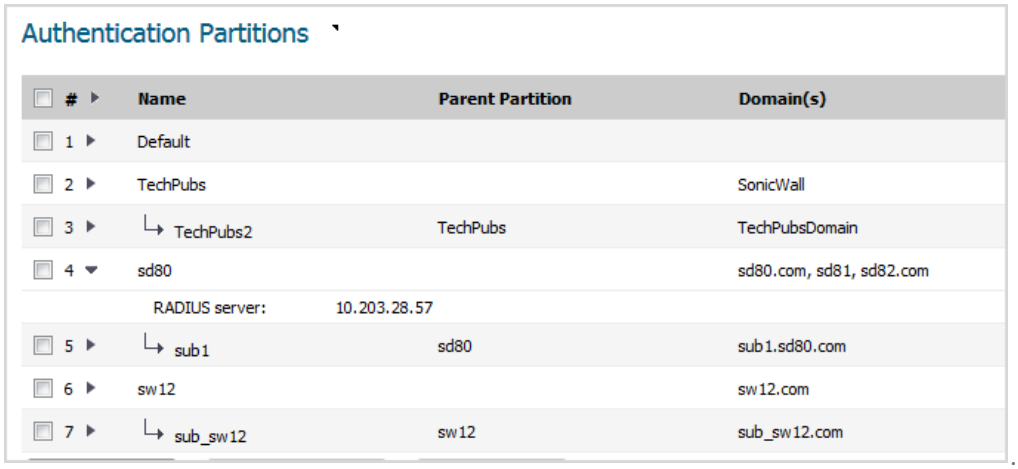


You can expand the tree of:

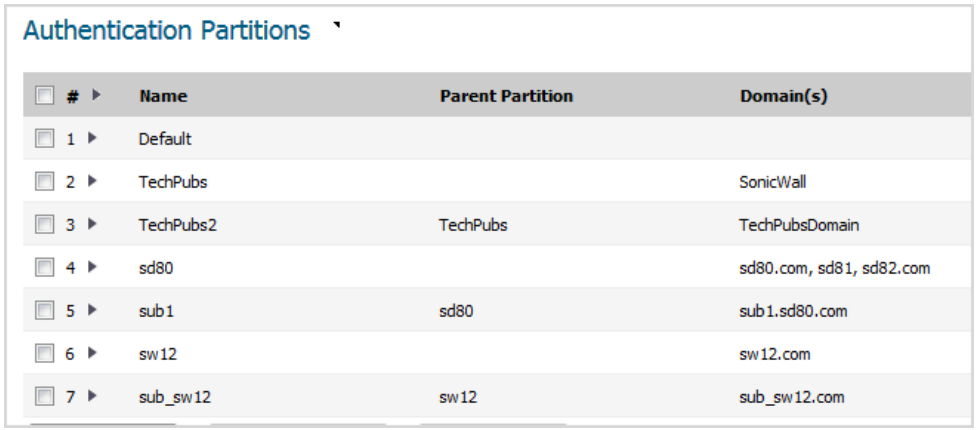
- All table entries by clicking the triangle next to the checkbox in the heading.
- One or more table entries by clicking the **Expand** icon of each.

Showing Hierarchy

By default, subpartitions are shown below their parent partition with a **Link** icon before the subpartition's name



You can show the subpartitions on the same level as their parent partition by clicking the **Group** icon.

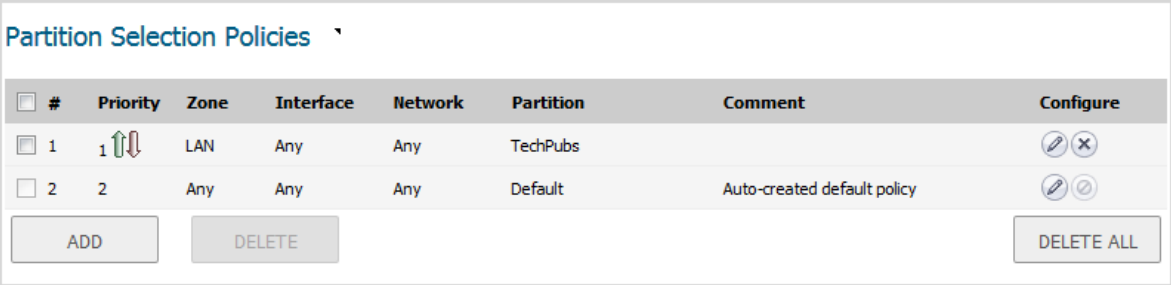


Partition Selection Policies Section

NOTE: This section displays only when authentication partitioning is enabled.

This section displays a table of policies affecting the selection of authentication partitions and allows you to create, delete, and edit the policies and change the priority of any policy you create. These policies select the partitions in the **Authentication Partitions** table based on the physical locations of the users being authenticated. When authenticating users whose domain names are not available for matching against those in the selected partitions, the users' partitions are selected base on their physical locations set by these policies. These selection policies are also used for auto-assigning authentication devices to partitions based on the physical locations of those devices.

The Default selection policy for the Default partition cannot be deleted, nor can its priority be changed; it is always the lowest priority.



Selection checkbox	Allows you to select one or more entries in the table. Selecting the checkbox in the table heading selects all entries except that of the **Default** selection policy.
Priority	Orders the partition selection policies according to the priority you assign them. Clicking on the **Priority Arrows** displays the **Change the selection policy priority** popup dialog. You cannot change the priority for the **Default** selection policy; it is always the lowest priority.
Zone	Displays the zone assigned to the partition selection policy.
Interface	Displays the interface assigned to the authentication partition selection policy.
Partition	Displays the authentication partition to which the selection policy applies.
Comment	Displays any comment you entered when creating or editing the selection policy. The selection policy for the **Default** partition has the comment **Auto-created default policy**.
Configure	Displays the **Edit** and **Delete** icons, which are dimmed for the default policy.
Add	Displays the **Add a partition selection policy** popup dialog for adding a selection policy for an authentication partition or subpartition.
Delete	Deletes the selected policy or policies. **NOTE:** You cannot delete the policy for the **Default** partition. **Delete** is dimmed unless at least one policy has been selected.
Delete All	Deletes all policies from the table except the policy for the **Default** partition.

There is always one selection policy in this table, the auto-created default policy for the **Default** partition. You cannot select this policy, delete it, change its priority, or edit it, except for choosing the partition to which it applies.

Enabling/Disabling Authentication Partitioning

To enable partitioning:

1 Navigate to the **Users > Partitions** page.



2 In the **Authentication Partitioning Settings** section, select **Enable authentication partitioning**. The **Authentication Partitions** and **Partition Selection Policies** sections display.

To disable partitioning:

1 Navigate to the **Users > Partitions** page.

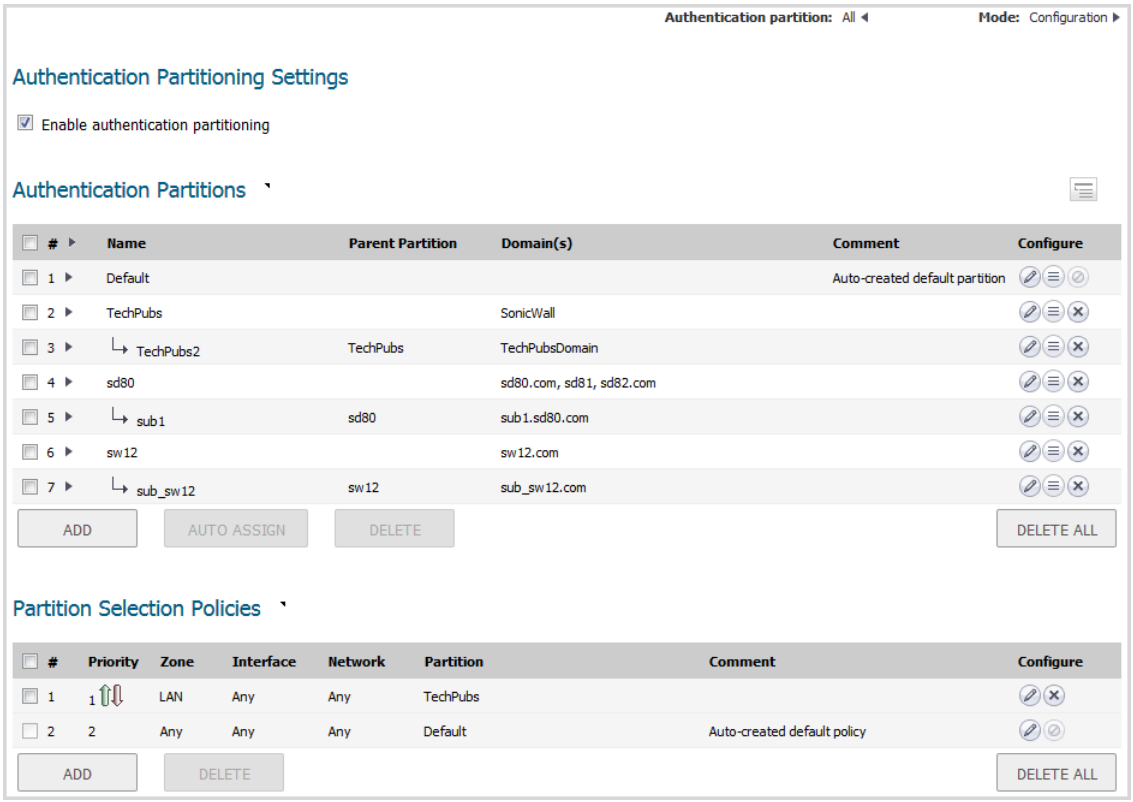
2 In the **Authentication Partitioning Settings** section, deselect the **Enable authentication partitioning** checkbox. The **Authentication Partitions** and **Partition Selection Policies** sections no longer display.

IMPORTANT: When you disable authentication partitioning, all partitioned LDAP servers, SSO agents, TSAs, and RADIUS accounting clients are moved to the **Default** authentication partition; RADIUS servers are not affected and remain in their configured authentication partitions. If you subsequently enable authentication partitioning, you need to reconfigure all other servers, agents, and clients.

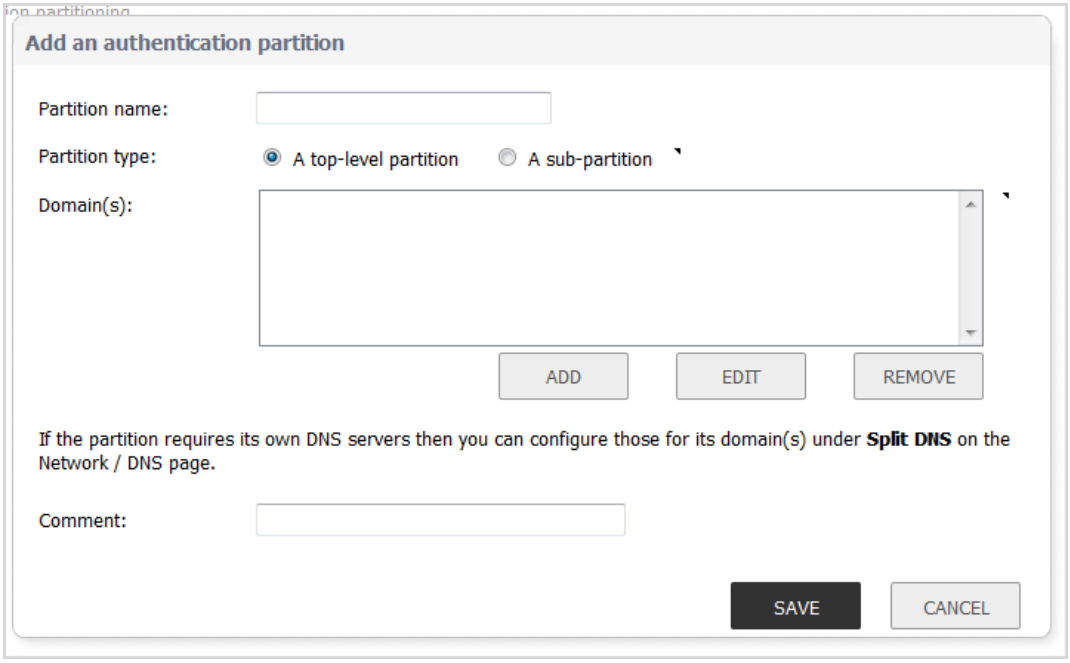
Adding Partitions and Subpartitions

To add a partition:

1 Navigate to the **Users > Partitions** page.



2 In the **Authentication Partitions** section, click **Add**. The **Add an authentication partition** popup dialog displays.



3 Enter a friendly, meaningful name in the **Partition name** field. The name can be from 1 to 32 alphanumeric characters.

4 For **Partition type**, choose whether the authentication partition is:

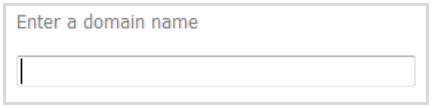
- **A top-level partition**; go to Step 6.
- **A sub-partition**; the **Parent partition** drop-down menu displays:



5 Select a parent partition from the drop-down menu. The default partition is **Default**.

TIP: If your installation does not have multiple partitions, then create subpartitions as subpartitions of the **Default** partition.

6 Under the **Domain(s)** list, click **Add**. The **Add domain** popup dialog displays.



7 Enter a domain name.

8 Click **OK**.

9 Repeat Step 6 through Step 8 for each domain you want to add.

10 Optionally, enter a comment in the **Comment** field.

11 Click **Save**. The partitions and/or subpartitions are added to the **Authentication Partitions** table. Subpartitions are positioned immediately after their parent partitions, with a **Link** icon indicating they are subpartitions.

Deleting Partitions and Subpartitions

NOTE: In this section, partition refers to both partitions and subpartitions.

You can delete a single partition, multiple partitions, or all partitions. If you delete a single partition, the servers, agents, and clients are reassigned to the **Default** partition.

NOTE: You cannot delete the **Default** partition.

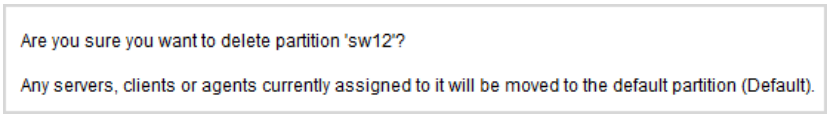
Topics:

- Deleting a Single Partition on page 208
- Deleting Multiple Partitions on page 208
- Deleting All Partitions (Except Default) on page 209

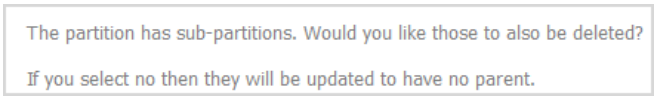
Deleting a Single Partition

To delete a single partition:

1. Navigate to **Users > Partitions**.
2. Under the **Authentication Partitions** table, click the **Delete** icon in the **Configure** column for the partition to be deleted. A verification message displays:



3. Click **OK**. If the partition:
   - Does not have subpartitions, the partition is deleted and the servers/agents/clients are reassigned to the **Default** partition.
   - Has subpartitions, this message displays:



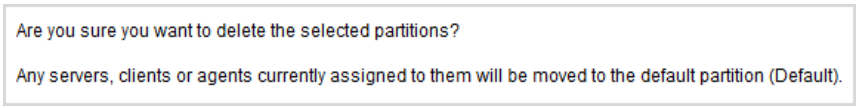
   a) Do one of these:
   - To delete the subpartitions as well as the parent partition, click **Yes**. All servers/agents/clients are reassigned to the **Default** partition.
   - To convert the subpartitions to top-level partitions while deleting the parent partition, click **No**. All servers/agents/clients are reassigned to the **Default** partition.
   - To not delete the parent subpartition, click **Cancel**.

Deleting Multiple Partitions

To delete multiple partitions

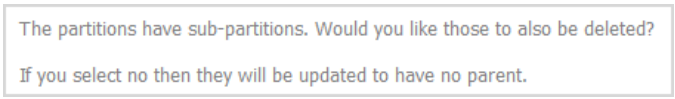
1. Navigate to **Users > Partitions**.
2. In the **Authentication Partitions** table, click the checkbox(es) of the authentication partition(s) you want to delete. You can select multiple partitions.

3 Click **Delete**. A verification message displays:



4 Click **OK**. If any partition:

- Does not have subpartitions, the partition(s) is deleted and the servers/agents/clients are reassigned to the **Default** partition.
- Has subpartitions, this message displays:



a Do one of these:

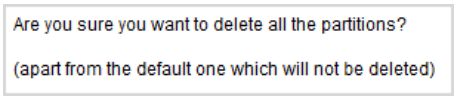
- To delete the subpartitions as well as the parent partition(s), click **Yes**. All servers/agents/clients are reassigned to the **Default** partition.
- To convert the subpartitions to top-level partitions while deleting the parent partition(s), click **No**. All servers/agents/clients are reassigned to the **Default** partition.
- To not delete the parent subpartition, click **Cancel**.

Deleting All Partitions (Except Default)

To delete all partitions (except Default)

1 Navigate to **Users > Partitions**.

2 In the **Authentication Partitions** table, click **Delete All**. A verification message displays:



3 Click **OK**. All servers/agents/clients are reassigned to the **Default** partition.

Assigning Servers, Agents, and Clients

After you have added the authentication partitions, you assign servers, agents, and/or clients to the partitions. You also can assign them to the authentication partitions at any time by following the same procedure.

You can have unassigned servers, agents, and clients auto assigned to the partition.

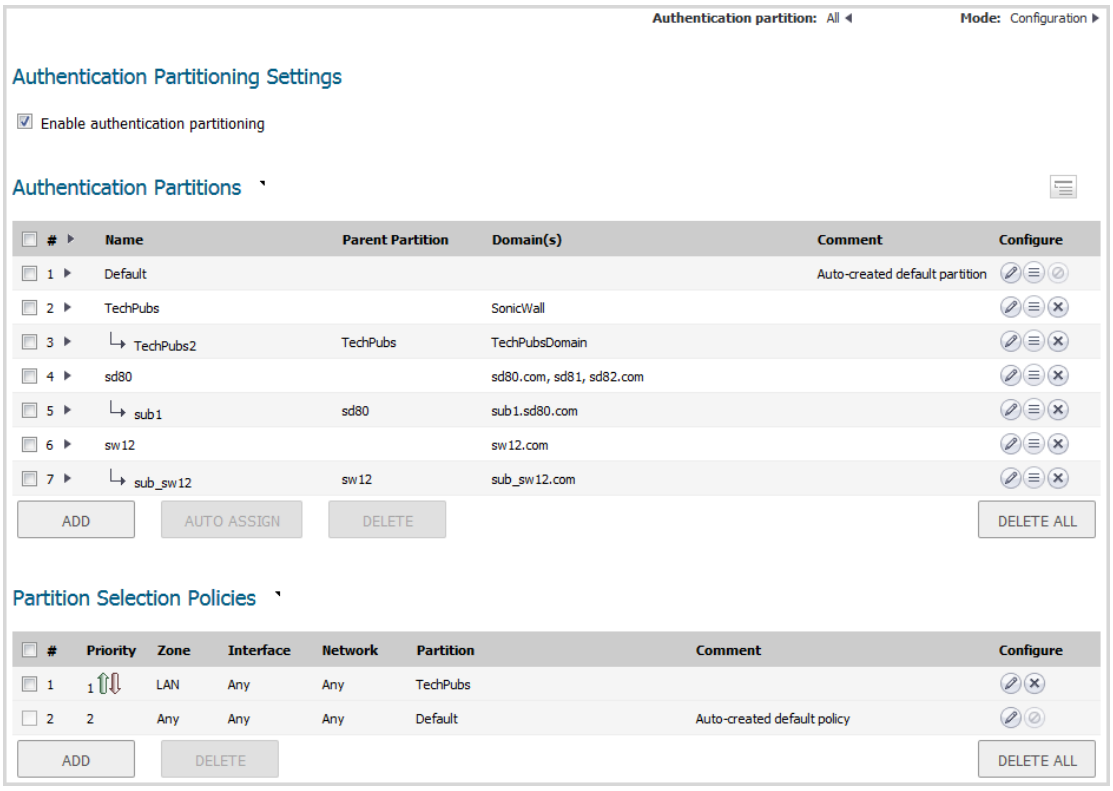
Topics:

- Assigning Manually on page 210
- Auto Assigning on page 211

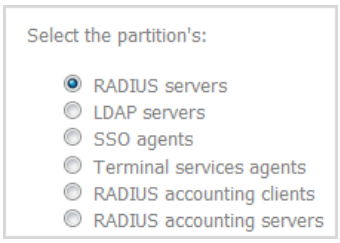
Assigning Manually

To assign servers, agents, and clients:

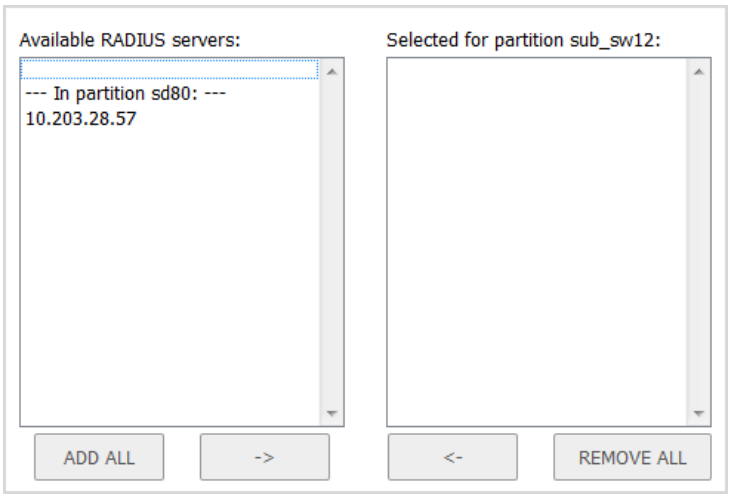
1 Navigate to **Users > Partitions**.



2 In the **Authentication Partition** table, click the partition's **Selection** icon in the **Configure** column. The **Select what?** popup dialog displays.



3 Select the type of server, agent, or client to assign. The appropriate **Select the *server/agent/client* for partition *partitionName*** popup menu displays with a list of available servers, agents, or clients.



4 Do one of the following:

- Select a server/agent/client from the **Available** list and click the **Right-arrow**.
- Select multiple items from the **Available** list by pressing the **Ctrl** key while selecting each item and then click the **Right-arrow**.
- Select all items by clicking **Add All**.

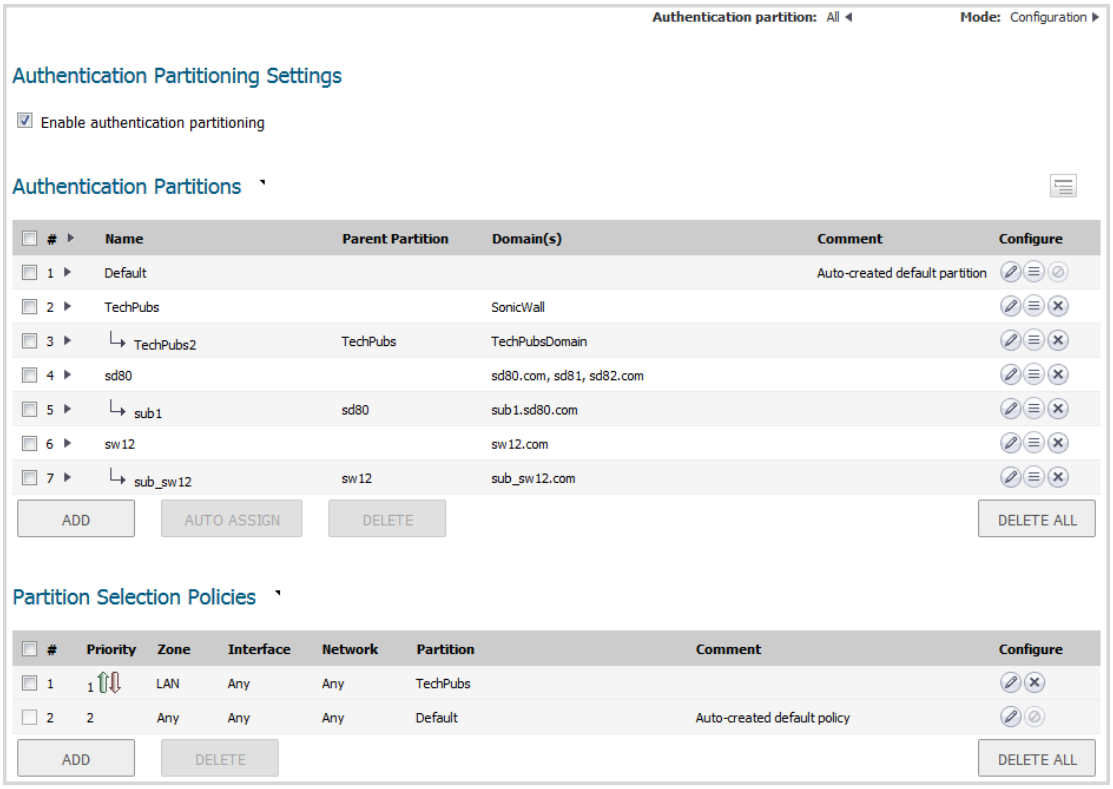
5 Click **SAVE**.

Auto Assigning

There is **Auto Assign** for assigning any unassigned servers, agents, and clients, based on their IP addresses or host names, to the relevant partitions automatically.

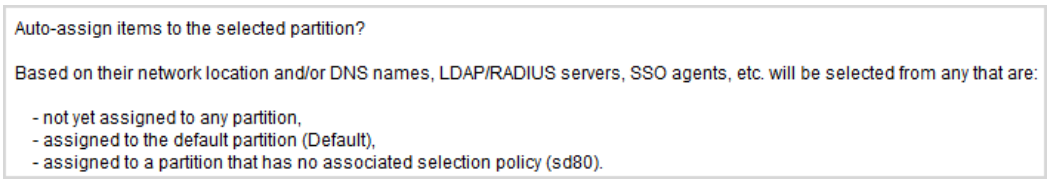
To auto assign servers, agents, and clients:

1 Navigate to **Users > Partitions**.



2 In the **Authentication Partitions** table, click the checkbox(es) of the authentication partition(s) to which you want to assign unassigned servers, agents, and/or clients. You can select more than one partition. **AUTO ASSIGN** becomes active.

3 Click **AUTO ASSIGN**. The auto-assign message appears.



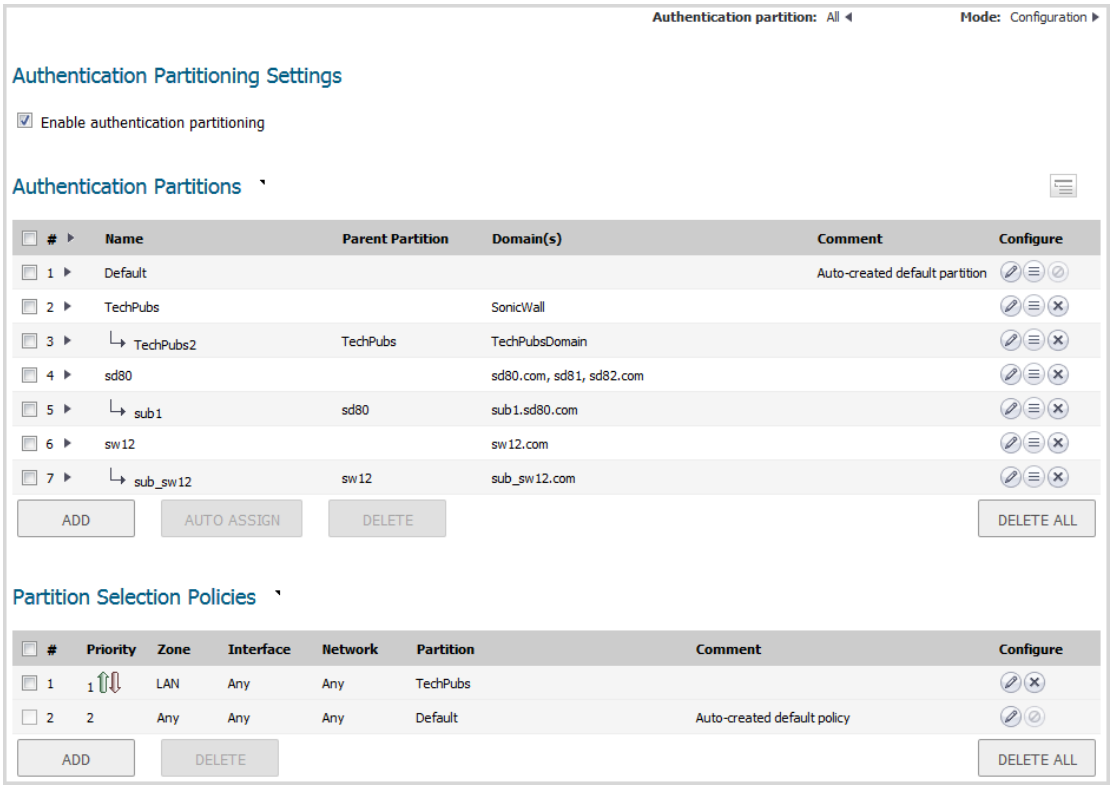
4 Click **OK**.

Editing Partitions

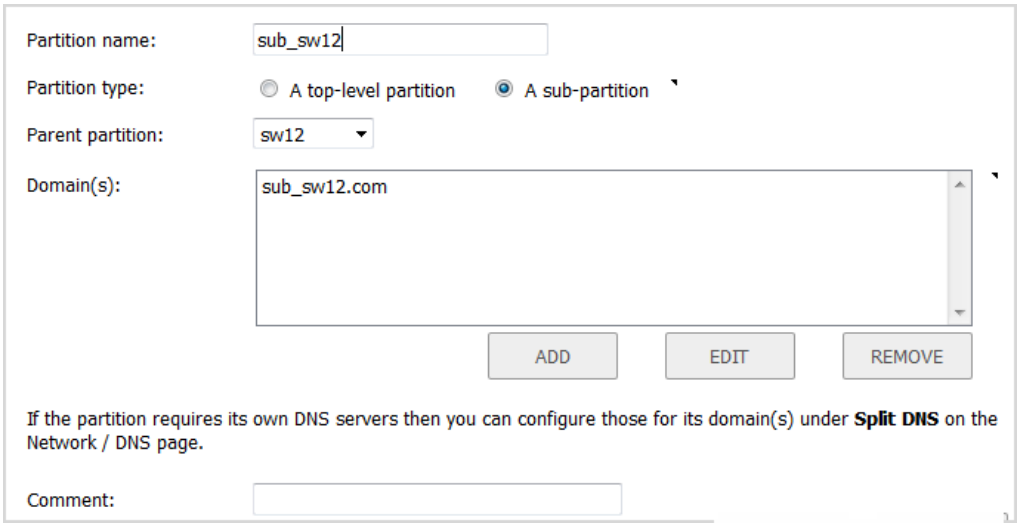
You can edit all partitions including the **Default** partition.

To edit a partition:

1 Navigate to **MANAGE | System Setup > Users > Partitions**.



2 In the **Authentication Partitions** table, click the **Edit** icon in the **Configuration** column of the authentication partition you want to modify. The **Edit authentication partition** popup displays.

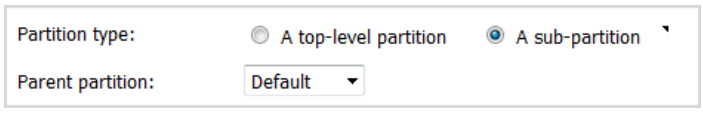


3 You can change the partition's name in the **Partition name** field. The name can be from 1 to 32 alphanumeric characters.

4 You can change a partition from a top-level partition to a subpartition or from a subpartition to a top-level partition by changing the **Partition type**; choose whether the authentication partition is now to be:

NOTE: A top-level partition that has subpartitions cannot be changed to a subpartition unless you first delete the subpartitions, reallocate them to a different top-level partition, or make them top-level partitions.

- **A top-level partition**, go to Step 6.
- **A sub-partition**; the **Parent partition** drop-down menu displays:



5 Select a parent partition from the **Parent partition** drop-down menu. The default partition is **Default**.

6 To:

- Edit a domain, go to Step 10.
- Delete a domain, go to Step 15.
- Add a domain, under the **Domain(s)** list, click **Add**. The **Add domain** popup dialog displays.

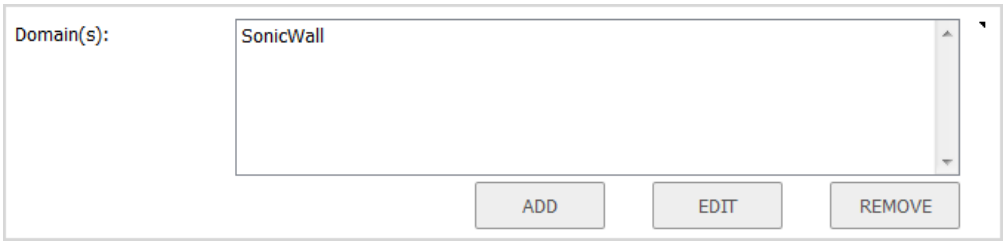


7 Enter a domain name, which can be from 1 to 32 alphanumeric characters.

8 Click **OK**.

9 Go to Step 17.

10 Select a domain to edit by clicking on it.



11 Click **Edit**. The **Edit domain** dialog displays.



12 Change the domain name.

13 Click **OK**.

14 Go to Step 17

15 Select a domain to delete.

16 Click **Remove**.

17 Repeat Step 6 for each domain you want to add, edit, or delete.

18 To change the servers used for looking up names in the partition, for **To look up names in the partition**, select either:

- **Use the default DNS servers**; go to Step 20.
- **Use the partition's DNS servers**; the **DNS Server 1**/-**2**/-**3** fields become active.

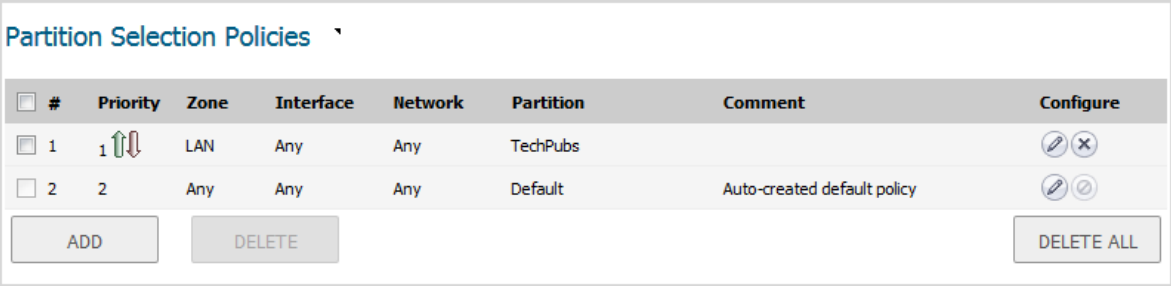
19 Enter up to three DNS servers in the **DNS Server 1**/-**2**/-**3** fields.

20 Optionally, enter a comment in the **Comment** field.

21 Click **SAVE**.

Configuring Partition Selection Policies

A partition selection policy specifies how an authentication partition is selected for a user. You add, edit, and manage authentication partition selection policies in the **Partition Selection Policies** section of the **MANAGE | System Setup > Users > Partitions** page. For a complete description of partition selection policies, see About Authentication Partition Selection on page 195.



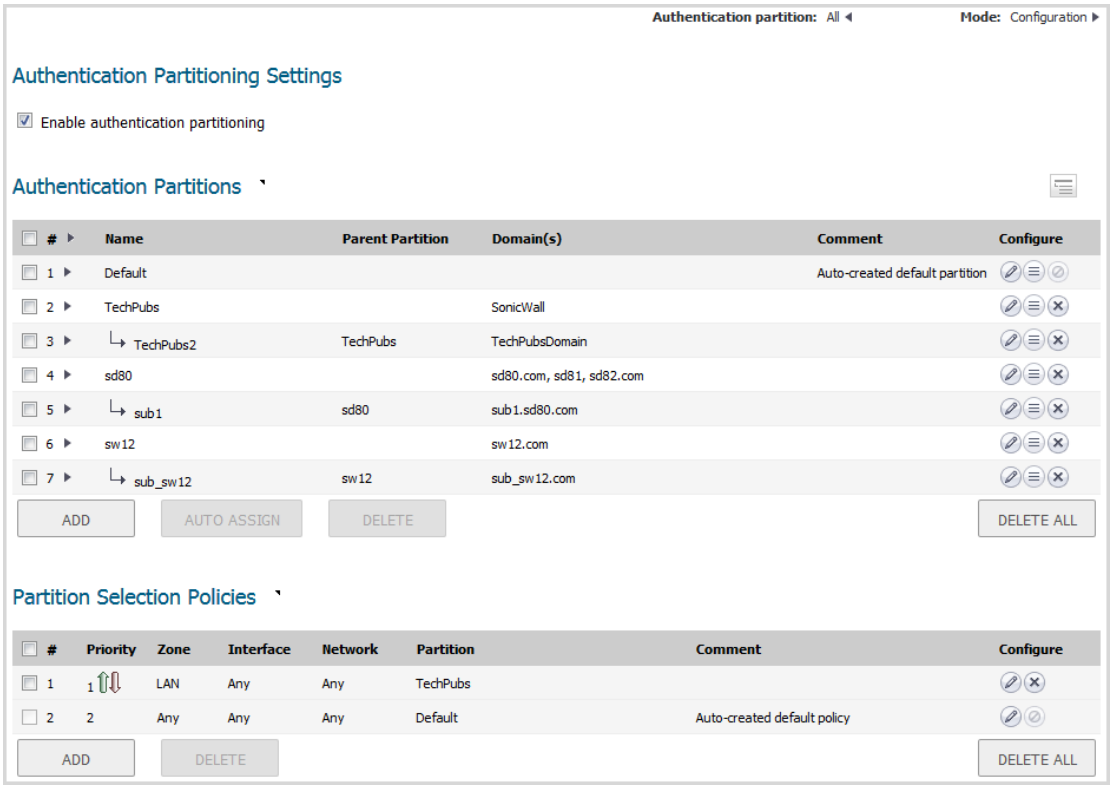
Topics:

- Adding Authentication Partition Selection Policies on page 216
- Changing Priority of a Selection Policy on page 217
- Modifying a Selection Policy on page 218
- Deleting Partition Selection Policies on page 218

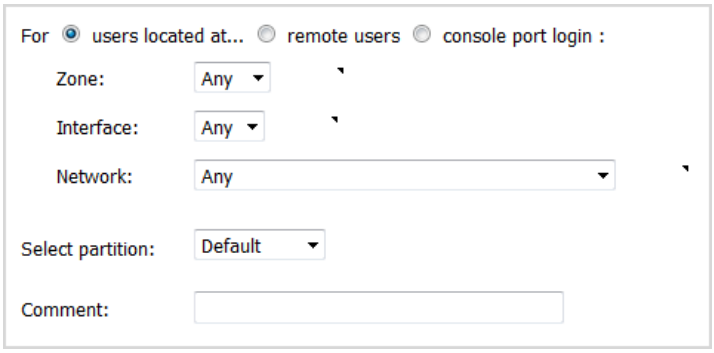
Adding Authentication Partition Selection Policies

To add a partition select policy:

1 Navigate to the **MANAGE | System Setup > Users > Partitions** page.



2 In the **Partition Selection Policies** section, click **Add**. The **Add a partition selection policy** popup dialog displays.

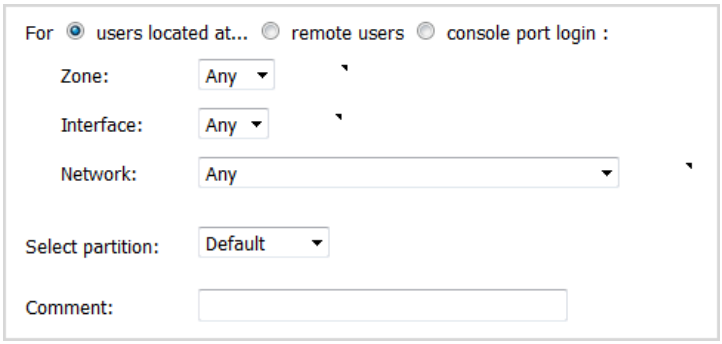


3 Choose the users' login location; what displays depends on your choice:

For this choice	Go to
users located at...	Step 4; this is the default
remote users	Step 7
console port login	Step 9

4 If you selected **users located at…**, select where the partitions are located from the **Zone**, **Interface**, and **Network** drop-down menus:

> **NOTE:** To select the partition it is typically not necessary to specify the zone, interface, and network. For optimum efficiency, it is best to give the minimum necessary.
>
> For example, if a partition is located through a specific interface, then just select that interface and leave the **Zone** as the default, **Any**. If a partition is located in a specific subnet, then just select that subnet as the **Network** and leave the **Zone** and **Interface** both set to the default, **Any**.



> **NOTE:** The choices provided in each drop-down menu vary by site.

- **Zone** – default is **Any**
- **Interface** – default is **Any**
- **Network** – default is **Any**; there are options to create a new Address Object and/or Address Group

5 Select a partition or subpartition from the **Select partition** drop-down menu. The default partition is **Default**.

6 Go to Step 10.

7 If you selected remote users, the options change; select a partition or subpartition from the **Select partition** drop-down menu. The default partition is **Default**.

8 Go to Step 10.

9 If you selected console port login, the options change; select a partition or subpartition from the **Select partition** drop-down menu. The default partition is **Default**.

10 Optionally, enter a comment in the **Comment** field.

11 Click **SAVE**.

Changing Priority of a Selection Policy

When determining an authentication partition to use, SonicOS searches the **Partition Selection Policies** table sequentially from top (1) to bottom (*n*). As you create selection policies, they are prioritized as follows:

1 Zone, with **Any** listed last in the group

2 Interface, with **Any** listed last in the group

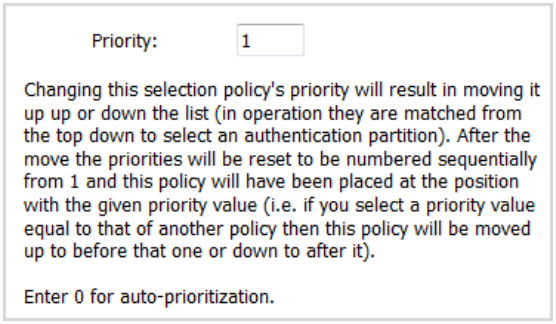
3 Network, with **Any** listed last in the group

You can change the priority of any policy except the **Default** partition selection policy, which is always the lowest priority.

Changing the priority of a selection policy moves the policy up or down the priority list. After the move, the priorities are reset to match the new ordering.

To change the priority of a policy:

1 In the **Partition Selection Policy** table, click the **Priority** icon for the selection policy. The **Change selection policy priority** popup dialog displays.



2 In the **Priority** field, enter the desired priority.

NOTE: Enter 0 for auto prioritization.

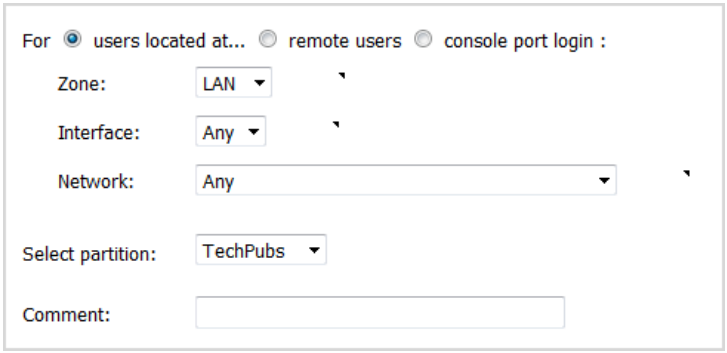
3 Click **OK**. The **Partition Selection Policies** table is updated to reflect the new ordering, including the reordering of other policies.

Modifying a Selection Policy

You can modify any partition selection policy except the auto-created **Default** policy. For the Default policy, you can only change the **Selected partition**.

To change a partition selection policy:

1 In the **Partition Selection Policies** table, click the **Edt** icon in the **Configure** column for the selection policy. The **Edit partition selection policy** popup dialog displays



2 This is the same dialog as the Add partition selection policy; for information about the dialog, see Adding Authentication Partition Selection Policies on page 216.

Deleting Partition Selection Policies

You can delete any partition selection policy except the auto-created **Default** policy for the **Default** authentication partition. You can delete a single policy, multiple policies, or all policies you created.

To delete a policy:

1 In the **Partition Selection Policies** section of the **Users > Partitions** page, click the **Delete** icon in the **Configure** column for the policy to be deleted. A verification message displays:



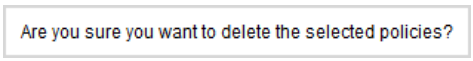
2 Click **OK**.

To delete multiple policies:

NOTE: The default partition selection policy cannot be deleted.

1 In the **Partition Selection Policies** section of the **Users > Partitions** page, select one or more policies to be deleted by clicking their checkboxes. **Delete** becomes active.

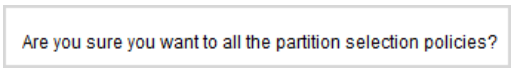
2 Click **Delete**. A verification message displays:



3 Click **OK**.

To delete all policies:

1 In the **Partition Selection Policies** section of the **Users > Partitions** page, click **Delete All**. A verification message displays:



2 Click **OK**.

Configuring Servers, Agents, and Clients for Authentication Partitioning

For each partition, you can configure:

User authentication method	Local Users
	RADIUS
	RADIUS + Local Users
	LDAP
	LDAP + Local Users
Single-sign-on method	SSO Agent
	Terminal Services Agent (TSA)
	RADIUS Accounting
	Browser NTLM Authentication

Authentication partitioning of all of the servers, agents, and clients is configured from the **Users > Settings** page; for a complete description of how to configure these entities and of the **User > Settings** page, see Configuring

Settings for Managing Users on page 116. For a description of how partitioning affects the configuration of servers and agents, see Configuring servers and agents.

NOTE: Operation of servers, agents, and clients is further described in Operation of Servers, Agents, and Clients with Subpartitions on page 193.

Configuring servers and agents

Server/agent	Partitioning configuration
RADIUS servers	A maximum of two RADIUS servers are configured as a primary/secondary redundant pair. You can configure multiple RADIUS server pairs, one primary/secondary pair per authentication partition.
LDAP servers	A number of primary LDAP servers can be configured, one for each authentication partition, plus a list of secondary servers for each (see About Extended Support for Multiple LDAP Servers on page 162). Typically, the LDAP servers for a domain or a group of inter-connected domains (forest in Active Directory terms) are allocated to each authentication partition.
SSO agents	Multiple SSO agents, in addition to supporting both load-sharing and redundancy, also support allocating agents to authentication partitions. A group of one or more agents are allocated to each authentication partition, and load-sharing and redundancy happens within each group.
TS agents	Partitioning of TSAs is required only for LDAP server selection for user group membership lookup. Configuration is optional as the TSA always supplies the full Windows NetBIOS domain name with the username. Thus, in most cases, it is possible to derive the authentication partition from the username.
RADIUS Accounting clients	Partitioning of SSO RADIUS Accounting clients is required only for LDAP server selection for user group membership lookup. Configuration is optional as some, but not all, RADIUS Accounting clients supply the domain name with the username in their accounting messages. Thus, in some cases, it is possible to derive the authentication partition from the username.

10

Configuring Local Users and Groups

Topics:

- About Authentication and Passwords on page 221
  - Using Two-Factor Authentication on page 221
  - Enforcing First Login Password Change on page 222
- Configuring Local Users on page 222
  - Quota Control for all Users on page 223
  - Viewing Local Users on page 223
  - Adding Local Users on page 224
  - Editing Local Users on page 230
  - Importing Local Users from LDAP on page 232
  - Configuring a Guest Administrator on page 232
- Configuring Local Groups on page 234
  - Creating or Editing a Local Group on page 235
  - Importing Local Groups from LDAP on page 241
  - Setting User Membership by LDAP Location on page 241

About Authentication and Passwords

Topics:

- Using Two-Factor Authentication on page 221
- Enforcing First Login Password Change on page 222

Using Two-Factor Authentication

Many user login authentication require one-time passwords (OTP). SonicOS 6.5.4 provides an additional method of OTP through email: Time-Based One-Time Password (TOTP) authentication with two-factor authentication.

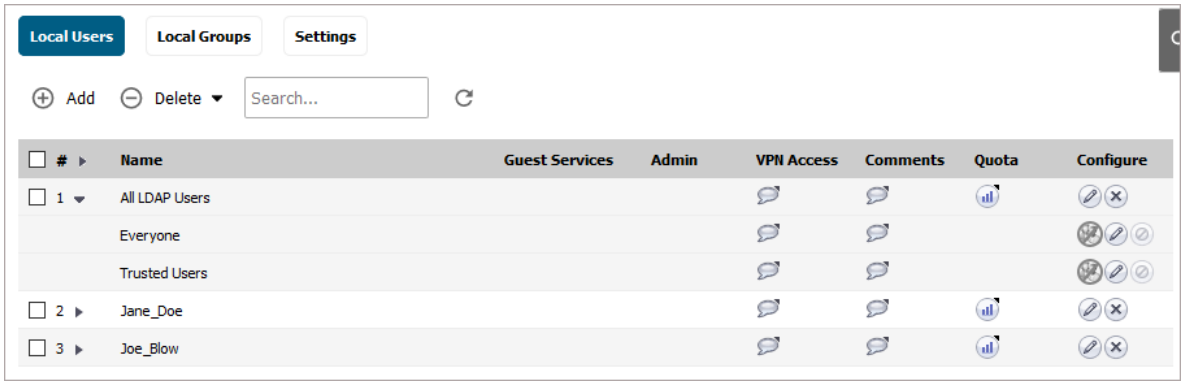
To use this feature, users must download a TOTP client app (such as Google Authentication, DUO, or Microsoft Authentication) on their smartphone. You select TOTP on the **Add/Edit User** dialog.

Enforcing First Login Password Change

Previously, when you created a user, you could allow uses to change their passwords after first logging in. SonicOS 6.5.4 allows you to force users to change their password before their first login when you create or edit a local user. You can specify the login change for users or for groups.

Configuring Local Users

Local users are users stored and managed on the SonicWall Security Appliance's local database. In **MANAGE | System Setup > Users > Local Users & Groups**, you can view and manage all local users, add new local users, and edit existing local users. You can also import users from your LDAP server.



Checkbox	Used to select individual local users.
Expand/Collapse icons	By default, only the local user's username is listed. Clicking the Expand icon displays the groups to which the local user belongs.
Name	Lists the username of the local user; when expanded lists the name(s) of the groups to which the local user belongs.
Guest Services	Indicates with a green checkmark icon whether guest services is active for the local user.
Admin	Displays the type of administration capabilities available to the local user.
VPN Access	Displays a **Comment** icon for each local user and each group to which the local user is a member. Mousing over the icon displays the status of the local group's VPN access and that of each member of the group.
Comments	Displays a **Comment** icon for each local user and each group to which the local user is a member. Mousing over the icon displays the comment entered when the local user or group was configured or edited.
Quota	For each local user, displays a **Statistics** icon. Mousing over the icon displays any usage quota for the local user.
Configure	Displays the **Edit** and **Delete** icons for each local user. If an icon is dimmed or otherwise disabled, that function is not available for that local user or local group.

For information about authentication and two-factor passwords, see About Authentication and Passwords on page 221.

Topics:

- Quota Control for all Users on page 223

- Viewing Local Users on page 223
- Adding Local Users on page 224
- Editing Local Users on page 230
- Configuring Global Settings on page 231
- Importing Local Users from LDAP on page 232
- Configuring a Guest Administrator on page 232

Quota Control for all Users

The quota control for users feature provides quota control based on the user's account. The quota can be specified as a session lifetime, or a transmit and/or receive traffic limit. With a cyclic quota, a user can not access the Internet upon meeting the account quota until the next cycle (day, week, or month) begins. If the quota cycle is Non Cyclic, the user is unable to access the Internet upon meeting the quota.

Previously, the quota control was supported only for guest users. Quota control is now specified for all local users as well.

Viewing Local Users

You can view all the groups to which a user belongs on **Users > Local Users & Groups**. Click on the **Expand** icon next to a user to view the group memberships for that user.

The columns to the right of the user's name list the privileges that the user has. In the expanded view, it displays which group the user gets each privilege from.

To:

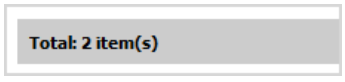
- View the network resources to which the user has VPN access, hover the mouse pointer over the **Comment** icon in the user's **VPN Access** column.
- View the quota for the user, hover the mouse over the **Statistics** icon in the **Quota** column
- Remove the user from a group, in the expanded view, click the **Remove** icon in the user's **Configure** column.

  NOTE: If the user cannot be deleted from a group, the icon is dimmed.

- Edit the user, click the **Edit** icon in the user's **Configure** column. See Editing Local Users on page 230.
- Delete the user or group in that row, click the **Delete** icon in the user's **Configure** column.

  NOTE: If the local user cannot be deleted from a group, the icon is dimmed.

The bottom of the **Users > Local Users & Groups** page displays the total number of local users:



Adding Local Users

You can add local users to the internal database on the Security Appliance from the **Users > Local Users & Groups** page.

NOTE: For the procedure for creating a user for an SSL VPN client, see *SonicOS 6.5 Connectivity*.

To add local users to the database:

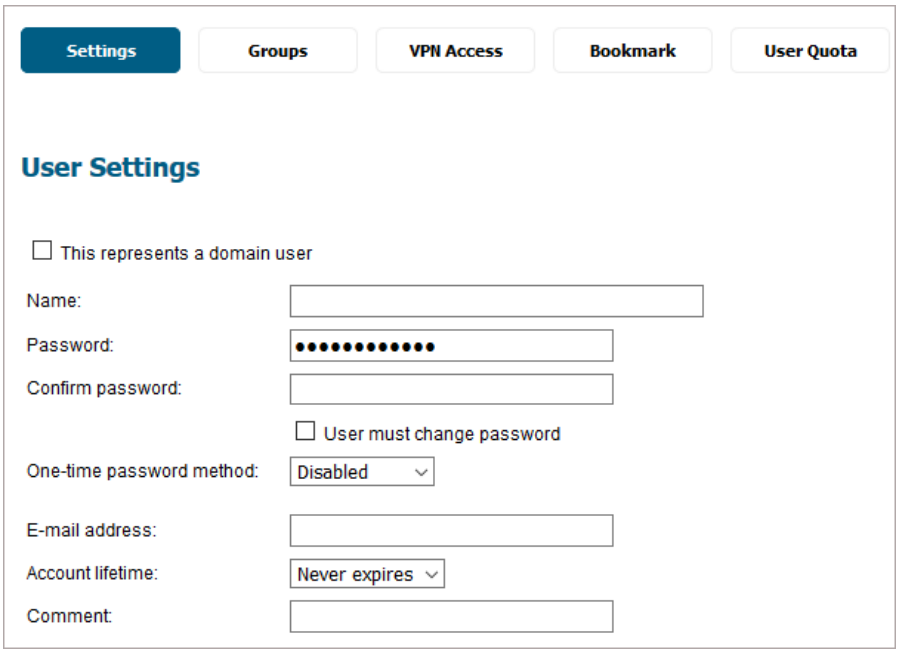
1 Navigate to **MANAGE | System Setup > Users > Local Users & Groups**.

2 If partitioning is:

- Not enabled, go to Step 3.
- Enabled, select the partition to which the settings apply from the **Authentication partitioning** drop-down menu. The default is **All**.

TIP: This menu displays only if partitioning is enabled.

3 Click the **Add** icon. The **Add User** dialog displays.



4 In **Settings**, indicate whether the group memberships, access rights, and other attributes apply to any domain user logging in using the registered domain account by selecting **This represents a domain user**. This option is not selected by default. When selected, other options display.

If **This represents a domain user** is:

- Selected, then any attributes, such as group membership and access rights, set apply for users who log in using the named domain account (authenticated through RADIUS or LDAP) or who are identified as hat domain user by SSO. You can have this attribute apply for the named user account in a specific domain or for a user with the given name in any domain.
- Not selected, the local user is a local account and anything this is set applies only for users who log in using the account and authenticated locally, in which case the password must be set in Step 8.

5 Type the user name into the **Name** field.

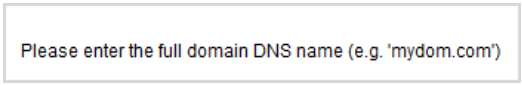
6 If the local user:

- Represents a domain user, the options change; go to Step 7.

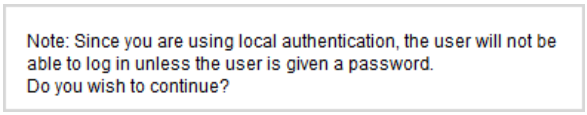


- Does not represent a domain user, go to Step 8.

7 Enter the domain name in the **Domain** field. You can select the **Domain** dome from the drop-down menu. If you enter a domain name that is not listed, you must enter the full domain name or a message is displayed:



If the domain is local, you must enter a password. If you do not, a message displays:

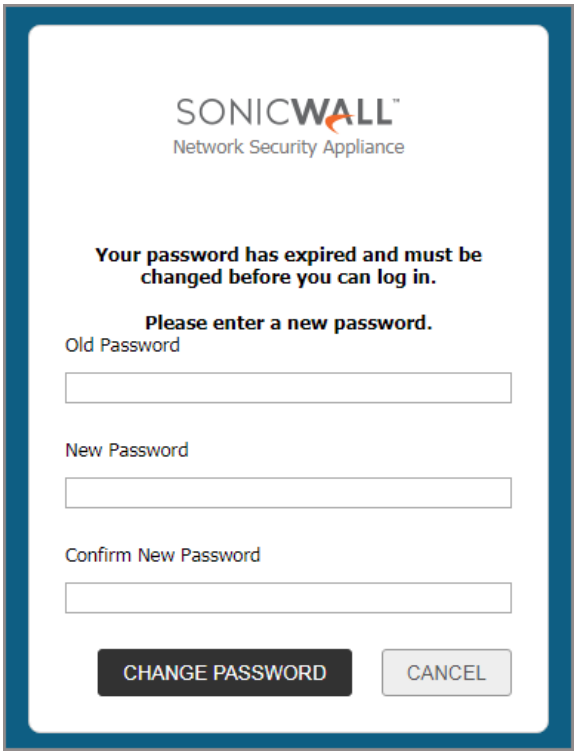


8 In the **Password** field, type a password for the user. Passwords are case-sensitive and should consist of a combination of 32 alphanumeric and special characters rather than names of family, friends, or pets. The length and type of characters are configured on **MANAGE | > Appliance > Base Settings > Login Security**.

NOTE: If **This represents a domain user** was not selected, you must enter a password.

9 Confirm the password by retyping it in the **Confirm Password** field.

10 Optionally, select **User must change password** to force users to change their passwords the first time they login. This option is not selected by default. If **User must change password** is selected, this dialog displays at the first login attempt:

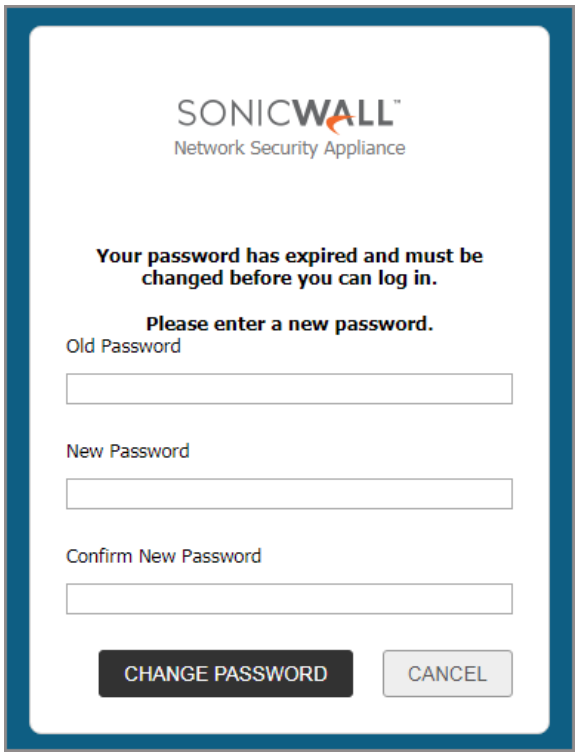


11 From **One-time password method**, select the method to require SSL VPN users to submit a system-generated password for two-factor authentication.:

TIP: If a Local User does not have one-time password enabled, while a group it belongs to does, ensure the user's email address is configured, otherwise this user cannot login.

TIP: To avoid another password change request for this user, the option applies only to the first login.

- **Disabled** (default) – If **User must change password** is selected this dialog displays at the first login attempt:



  Go to Step 13.

- **OTP via Mail** – Users receive a temporary password by email after they enter their user name and first password. After receiving the password-containing email, they can enter the second password to complete the login process. Go to Step 12.
- **TOTP** – Users receive a temporary password by email after they input their user name and first password, but to use this feature, users must download a TOTP client app (such as Google Authentication, DUO, or Microsoft Authentication) on their smartphone.

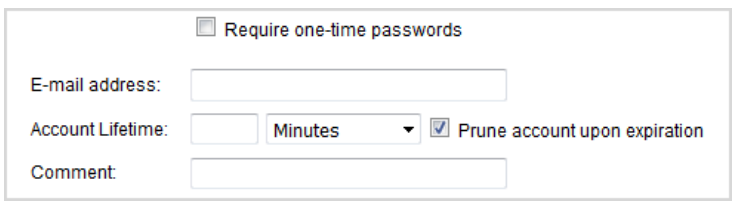
  The **UNBIND TOTP KEY** displays.



12 Enter the user's email address so they may receive one-time passwords.

13 From **Account Lifetime**, select the duration a user account exists before it is ether deleted or disabled. Depending on your selection, more options display:

- **Never expires** makes the account permanent. This is the default. Go to Step 16.
- **Minutes**, **Hours**, or **Days** specify a lifetime after which the user account is either deleted or disabled. If you choose a limited lifetime, the option changes:



14 Enter the lifetime in the **Account Lifetime** field. You can specify up to 9999 hours, minutes, or days.

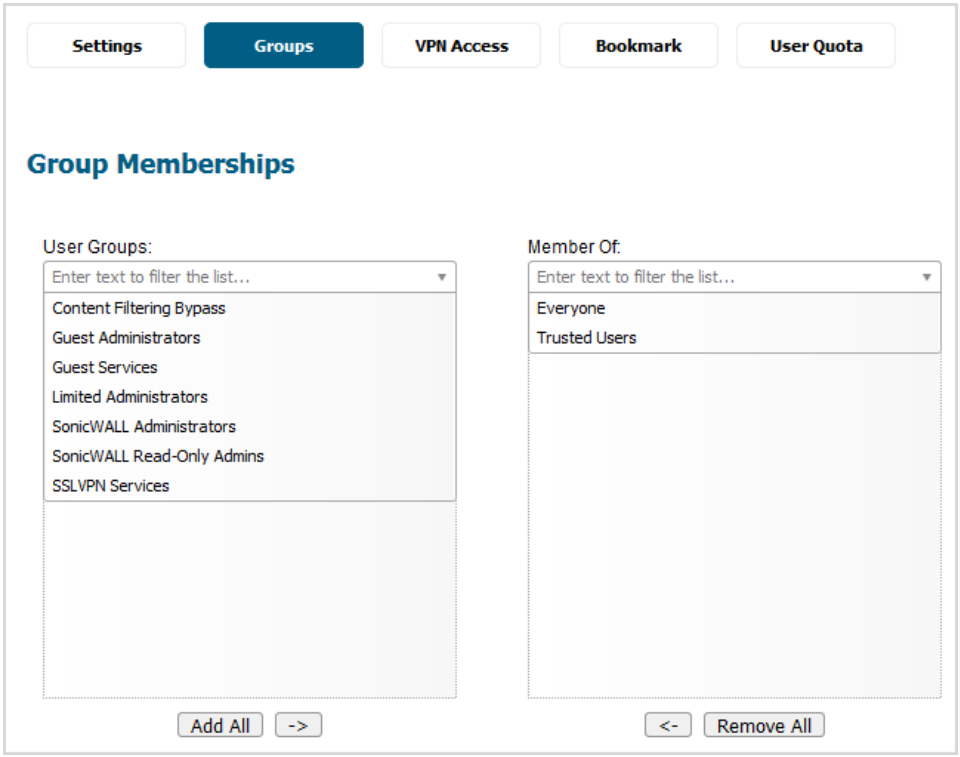
15 To

- Have the user account deleted after the lifetime expires, select **Prune account upon expiration**. This option is selected by default.
- Have the account simply be disabled after the lifetime expires, disable this option. You can then re-enable the account by resetting the account lifetime.

16 Optionally, enter a comment in the **Comment** field.

17 Click **Groups**.

Groups



1 From **User Groups**:

a Select one or more groups to which the user belongs.

b Either:

- Click the **Right Arrow ->** to move the group name(s) into the **Member of** list. The user is a member of the selected groups.
- Click **Add All**.

NOTE: To remove the user from a group:

1 Select the group from the **Member of** list

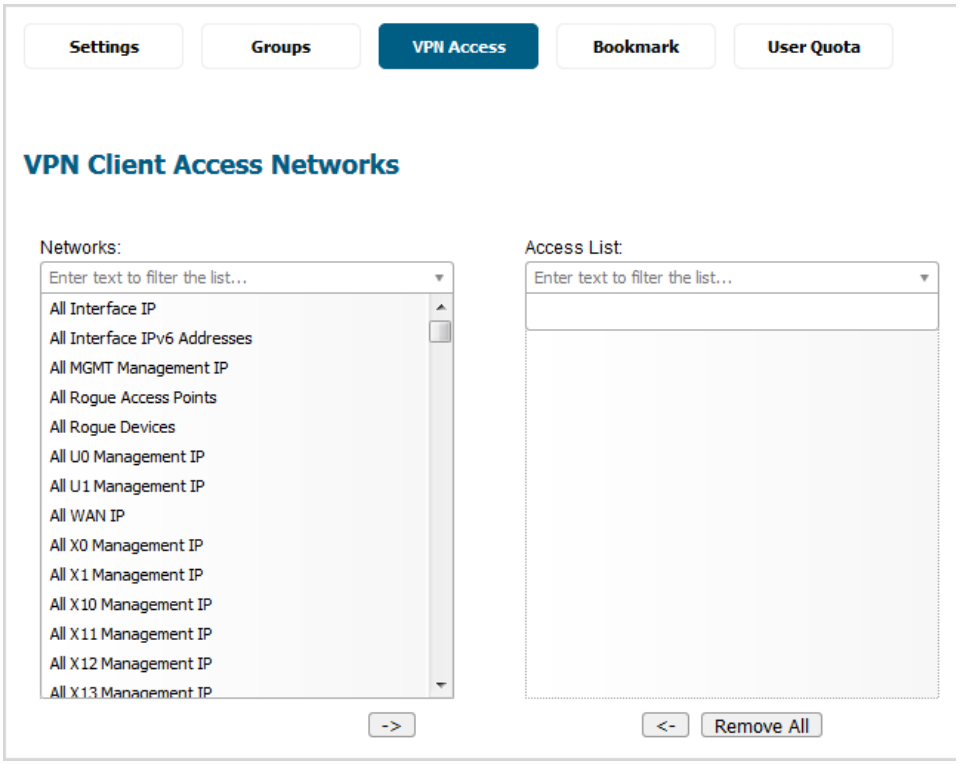
2 Either:

- Click the **Left Arrow <-**.
- Click **Remove All**.

NOTE: You cannot delete **Everyone** and **Trusted Users** from **Member Of**.

2 To configure which network resources VPN users (either GVC, NetExtender, or Virtual Office bookmarks) can access, click **VPN Access**.

VPN Access



1 Select one or more networks from **Networks**.

2 Click **Right Arrow** to move them to **Access List**.

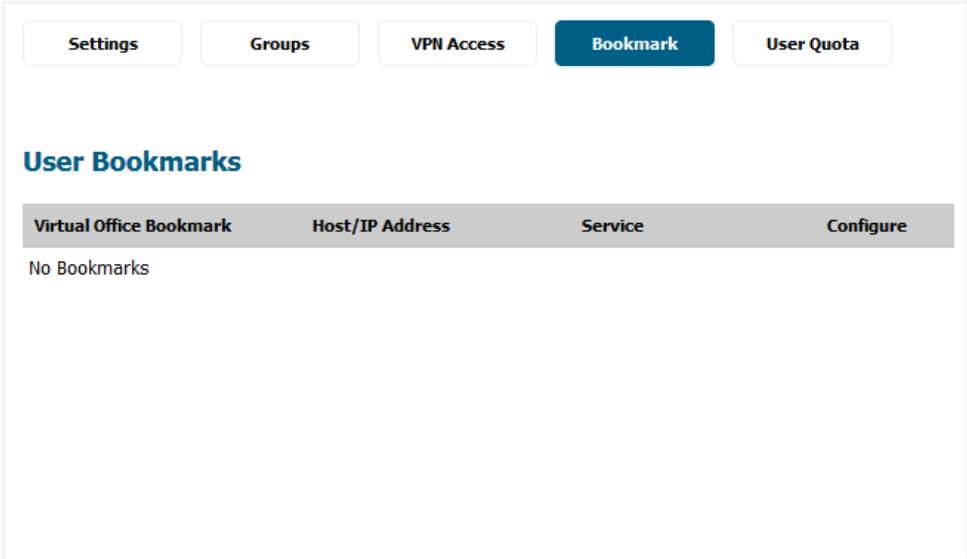
> **NOTE:** **VPN Access** affects the ability of remote clients using GVC, NetExtender, and Virtual Office bookmarks to access network resources. To allow these users to access a network resource, the network address objects or groups must be added to the **Access List**.

To remove the user's access to a network:

- Select the network(s) from the **Access List**, and then click **Left Arrow**.
- Click **Remove All**.

3 To add, edit, or delete Virtual Office bookmarks for each user who is a member of a related group, click **Bookmark**.

Bookmark



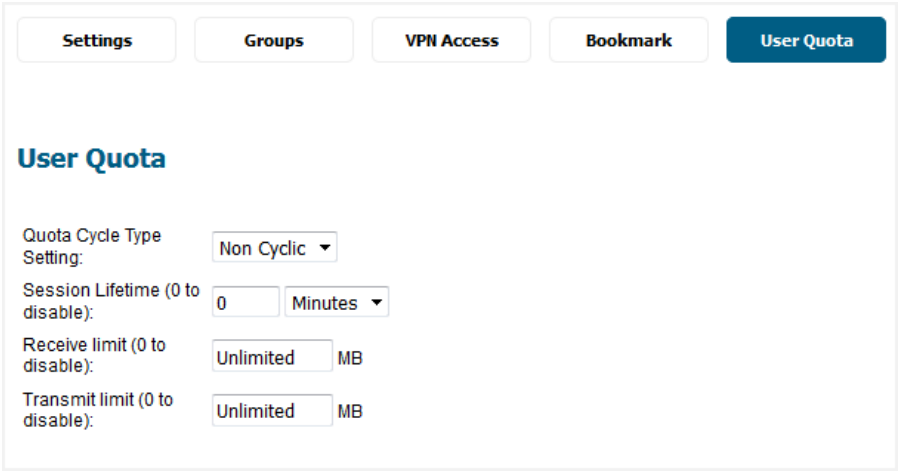
4 To add a bookmark, click **Add Bookmark**. For information on configuring SSL VPN bookmarks, see *SonicOS 6.5 Policies*.

NOTE: Users must be members of the SSL VPN Services group before you can configure Bookmarks for them. If the users are not members, you must add them to the SSL VPN Service group and submit the change to enable bookmarks.

5 Click **User Quota**.

User Quota

1 Click **User Quota**.



2 Configure the options.

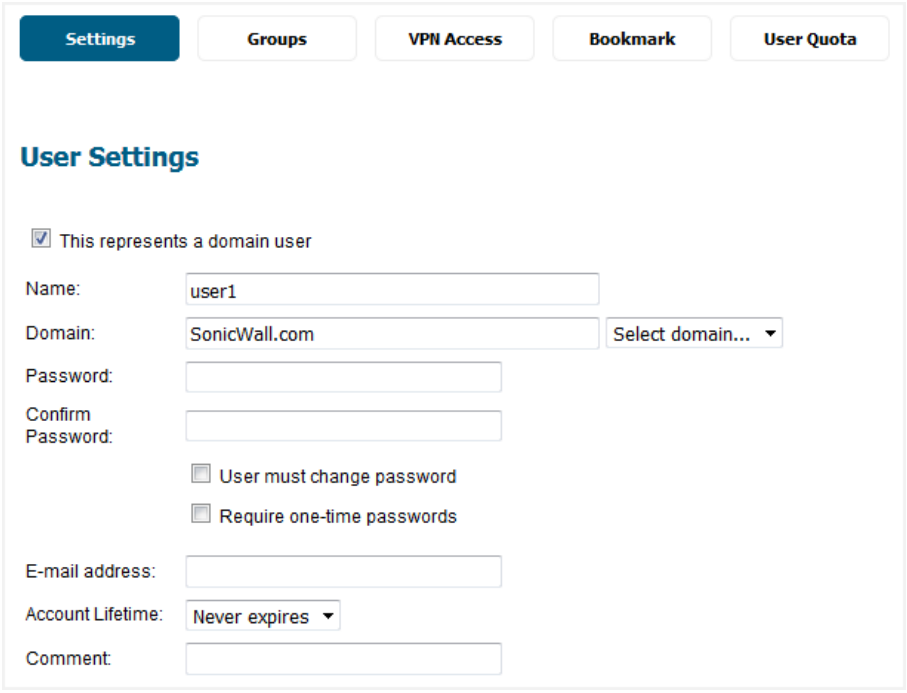
3 Click **OK** to complete the user configuration.

Editing Local Users

You can edit local users from the **Users > Local Users & Groups** page.

To edit a local user:

1 In the **Local Users** table, click the user's **Edit** icon under **Configure**. The **Edit User** dialog displays.



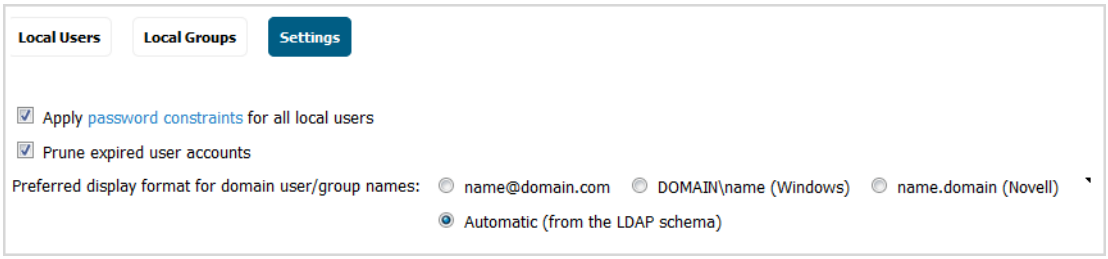
2 Configure the **Settings**, **Groups**, **VPN Access, Bookmark,** and **User Quota** options exactly as when adding a new user. See Adding Local Users on page 224.

Configuring Global Settings

To configure settings for all users:

1 Navigate to **MANAGE | System Setup** > **Users > Local Users & Groups**.

2 Click Settings.



3 Select **Apply password constraints to all local users**. This option is selected by default.

4 Select **Prune expired user accounts**. This option is selected by default.

5 To choose the display format in the management interface, chose a format from **Preferred display format for domain user/group names**:

TIP: This option affects only how the names are displayed and not how they are processed by the firewall. For example, choosing Windows format displays `user1@mydomain.com` displays as `MYDOMAIN\user1`, but the user object is processed as `user1@mydomain.com`. Also, the display is affected only for those users and user group objects that have explicit domains; those for which any domain selected is not affected.

- **name@domain.com**
- **DOMAIN\name (Windows)**
- **name.domain (Novell)**
- **Automatic (from the LDAP schema)** (default)

NOTE: If Automatic is chosen, the format is selected based on the schema of the LDAP server. LDAP, therefore, must be enabled; otherwise, the default format of `name@domain.com` is used.

If partitioning is enabled, then the display format is chosen based on the LDAP schema for the partition that hosts the domain. This allows for different display formats in partitions that have different domain types.

6 Click **ACCEPT**.

Importing Local Users from LDAP

You can configure local users on the firewall by retrieving the user names from your LDAP server. Having users on the firewall with the same name as existing LDAP/AD users allows SonicWall user privileges to be granted upon successful LDAP authentication.

The list of users read from the LDAP server can be quite long, and you probably only want to import a small number of them. **Remove from list** is provided, along with several methods of selecting unwanted users. You can use these options to reduce the list to a manageable size and then select the users to import. for information about how to import users from an LDAP server, see About Importing and Mirroring from LDAP on page 163.

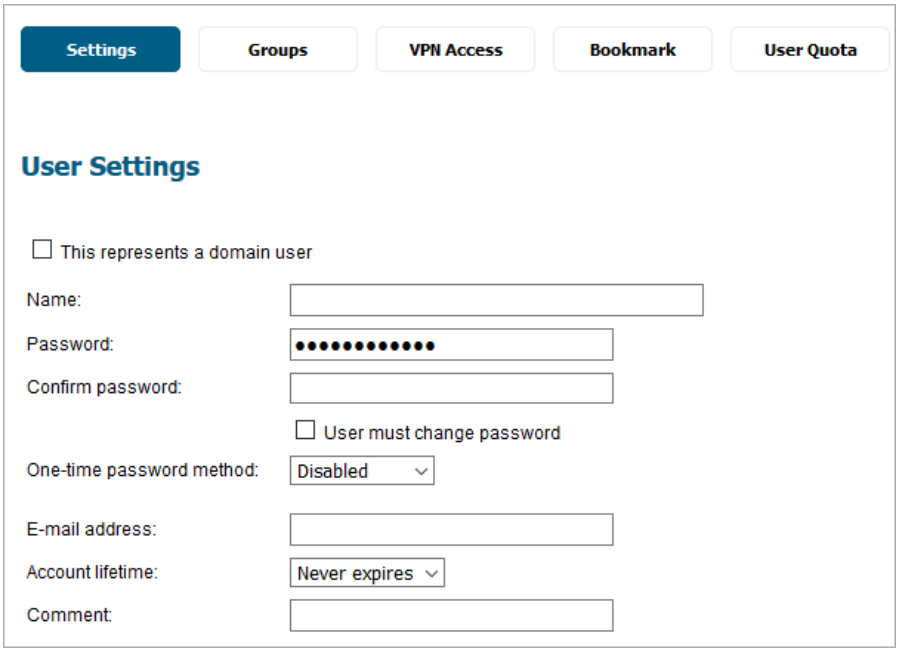
Configuring a Guest Administrator

A Guest Administrator privileges group is available to provide administrator access only to manage guest accounts and sessions.

To configure a Guest Administrator account:

1 Navigate to **Users > Local Users & Groups**.

2 Click **Add**. The **Add User** dialog displays.



3 Give the user a name in the **Name** field.

4 Click **Groups**.

5 Select **Guest Administrators** in the **User Groups** list.

6 Click Right Arrow to move **Guest Administrators** to the **Member Of** list.

7 Click **OK**.

8 Navigate to **Network > Interfaces**.

9 Click the **Edit** icon for the LAN interface. The **Edit Interface** dialog displays.

10 To allow the Guest Administrator account to login to the Security Appliance from the LAN, under **User Login** select both **HTTP** and **HTTPS**.

11 Click **OK**.

Logging on as Guest Administrator

To log on as Guest Administrator:

1 Log on to the Security Appliance as the Guest Administrator. The dialog showing access to privileged services displays.

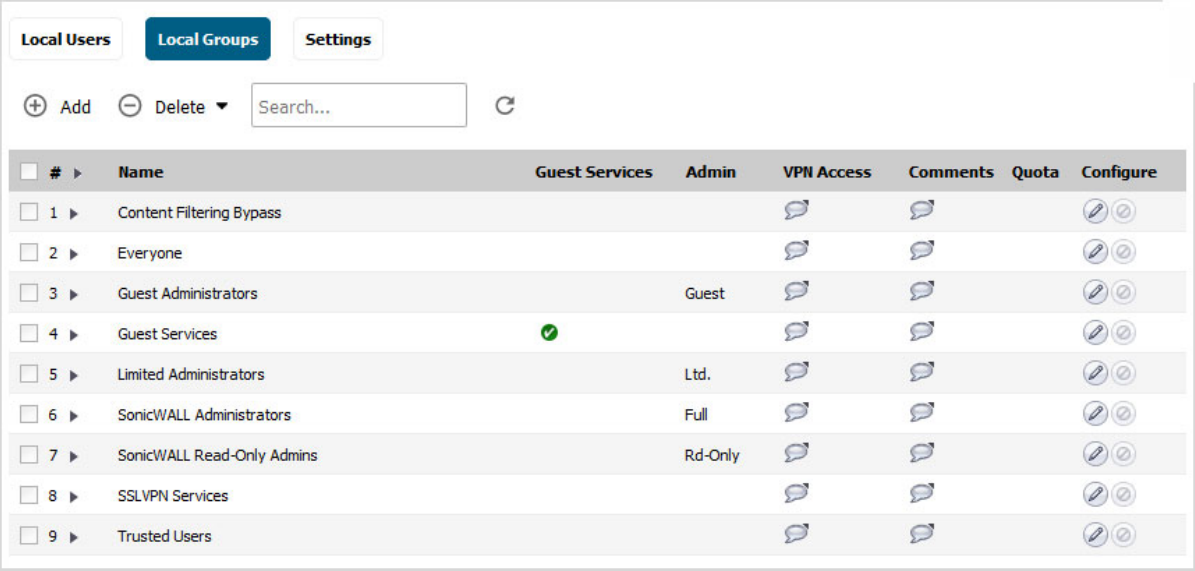
2 Click **MANAGE**.

After logging in, the Guest Administrator can:

- View all the guest accounts and sessions through the **Monitor | User Sessions > Active Guest Users** page.
- Manage the guest users (create, delete, update) on these management interface pages:
  - **MANAGE | System Setup > Users > Guest Accounts**
  - **MANAGE | System Setup > Users > Guest Services**

Configuring Local Groups

Local groups are displayed in the **Local Groups** table. Certain local groups are default groups that can be modified, but not deleted.



Checkbox	Used to select individual local groups. Default local groups cannot be changed, and, therefore, their checkboxes are dimmed.
Expand/Collapse icons	By default, only the local group's name is listed. Clicking the
Name	Lists both the default and configured local groups by name. If the **Enable Multiple Administrator Role** option has been enabled on the **System > Administration** page, the **Users > Local Groups** page lists these default role-based administrator groups: • System Administrators • Cryptographic Administrators • Audit Administrators
Guest Services	Indicates with a green checkmark icon whether guest services is active for the local group.
Admin	Displays the type of administration capabilities available to the local group. Mousing over the icon displays a tooltip regarding the listed capability.
VPN Access	Displays a **Comment** icon for each group and each member of the group. Mousing over the icon displays the status of the local group's VPN access and that of each member of the group.
Comments	Lists any comment provided for the local group.
Quota	Lists any usage quota for the local group.
Configure	Displays the **Edit** and **Delete** icons for each local group and group member, and for group members, a **Remove** icon. If an icon is dimmed, that function is not available for that local group or group member.

For information about authentication and two-factor passwords, see About Authentication and Passwords on page 221.

Topics:

- Creating or Editing a Local Group on page 235
- Importing Local Groups from LDAP on page 241

Creating or Editing a Local Group

This section describes how to create a local group, but also applies to editing existing local groups. When adding or editing a local group, you can add other local groups as members of the group.

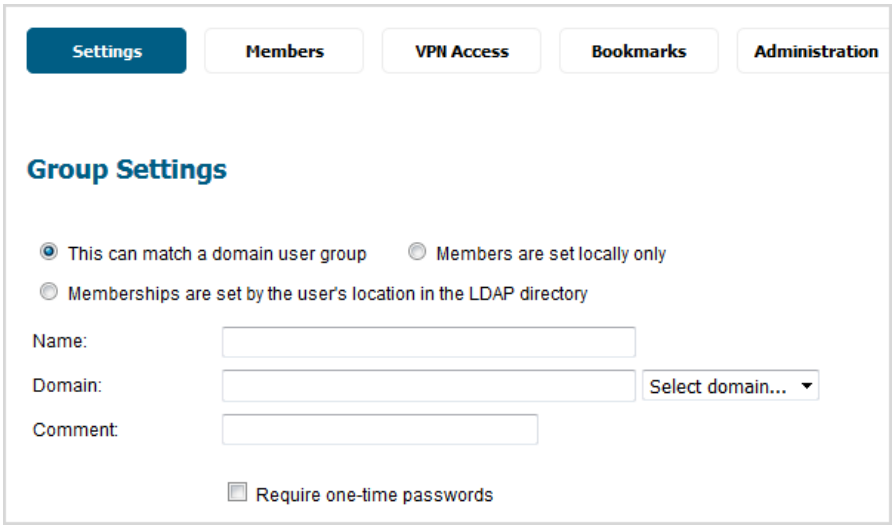
Topics:

- Adding a Local Group on page 235
- Editing a Local Group on page 241

Adding a Local Group

To add a local group:

1. Navigate to **Users > Local Users & Groups**.
2. Click **Add**. The **Add Group** dialog displays.



Topics:

- Settings on page 236
- Members on page 238
- VPN Access on page 239
- Bookmarks on page 240
- Administration on page 240

Settings

1 Choose how to set the ways that users are given membership to this group when they log in or are identified through SSO:

NOTE: Users who are given membership in this user group are given any privileges and access rights that are given to the group.

This can match a domain user group (default)	Any users who are members of a domain user group with the same name as this one are given membership to this group. You can choose to have membership: • Only for members of the domain user group in a specific domain. • For users who are members of the named group in any domain. **NOTE:** The options change when this is selected.
Members are set locally only	Local users are the only users given membership in the group. This option is not selected by default.
Memberships are set by the user's location in the LDAP directory	When users log in or are identified through SSO, if their user object on the LDAP server is at the location specified in **LDAP Location** (or under it if appropriate), they are given membership to this user group for the session. This setting is disabled by default. **NOTE:** There is no corresponding user group on the LDAP server and membership to the group is not related to any memberships set in domain user groups there. **NOTE:** The options change when this is selected.

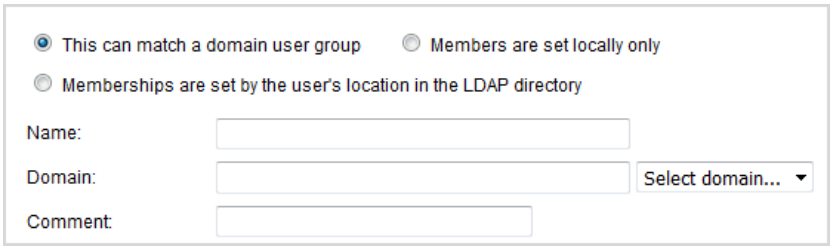
NOTE: In all cases, local users (including those representing domain users) and other user groups also can be made members of the group on the **Members** page of this dialog.

2 Enter a name for the local group in the **Name** field.

NOTE: The name of a predefined user or group cannot be edited and the field is dimmed.

3 If you selected:

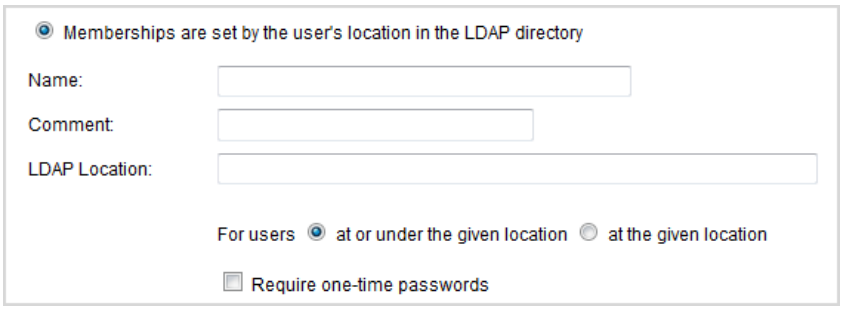
- **This can match a domain user group**, the options change. go to Step 4.



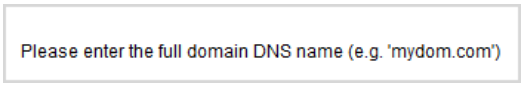
- **Members are set locally only**, go to Step 5.

- **Memberships are set by user's location in the LDAP directory**, the options change. go to Step 5.

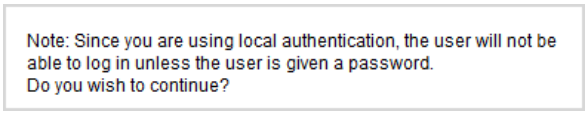
  TIP: Local users and other groups also can be made members of the group on the **Members** tab.



4 Enter the domain name in the **Domain** field. You can select the **Domain** dome from the drop-down menu. If you enter a domain name that is not listed, you must enter the full domain name or a message is displayed:



If the domain is local, you must enter a password. If you do not, a message displays:



5 Optionally, enter a descriptive comment in the **Comment** field.

6 If you selected either **This can match a domain user group** or **Members are set locally only**, go to Step 9.

7 In the **LDAP Location** field, enter the location in the LDAP directory tree. The location can be given as a path (for example, `domain.com/users`) or as an LDAP distinguished name.

NOTE: If LDAP user group mirroring is enabled, then for mirror user groups this field is read-only and displays the location in the LDAP directory of the mirrored group.

8 Choose where the location is from the **For Users** options:

- **at or under the given location** (default)
- **at the given location**

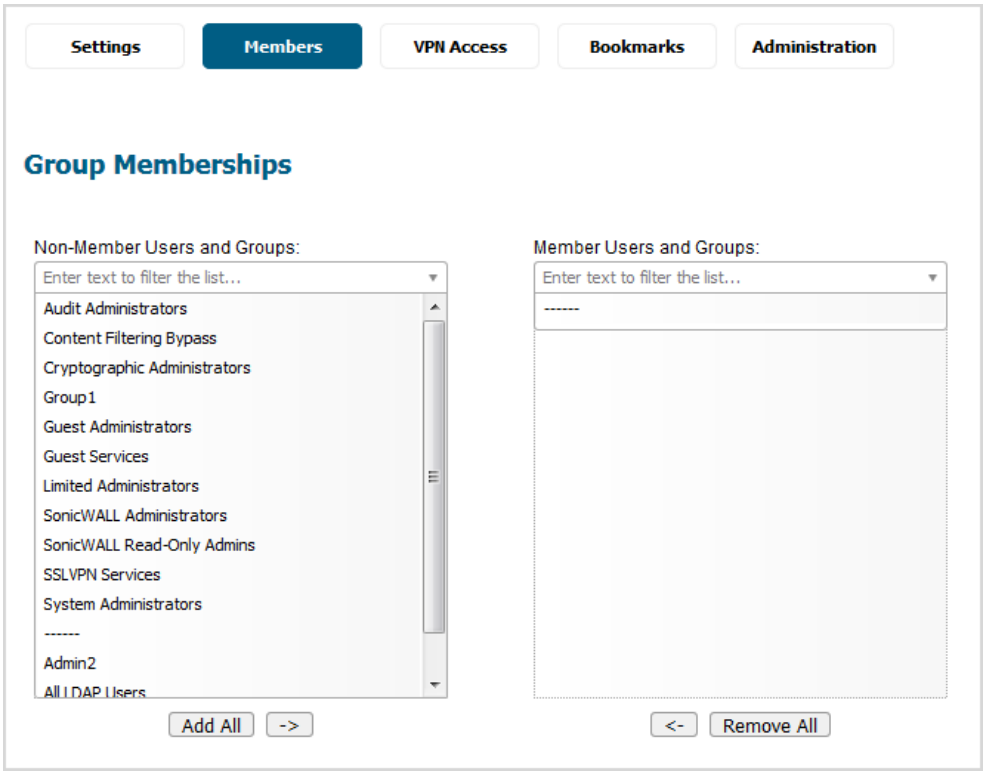
9 Optionally, to require one-time passwords for the group, select the **Require one-time passwords** checkbox. If you enable this setting, users must have their email addresses set.

10 To:

- Finish adding the group, click **OK**.
- Add members, go to Members on page 238.

Members

1 Click **Members**.



2 From the **Non-Member Users and Groups** list, select the user(s) and/or group(s) you want to add.

3 To add:

- User(s) and/or group(s) to the **Member Users and Groups** list:
  a) Select the user(s) and/or group(s) from the **Non-Member Users and Groups** list.
  b) Click the **Right Arrow ->**.
- All users and groups, click **Add All**.

NOTE: You can add any group as a member of another group except **Everybody** and **All LDAP Users**. Be aware of the membership of the groups you add as members of another group.

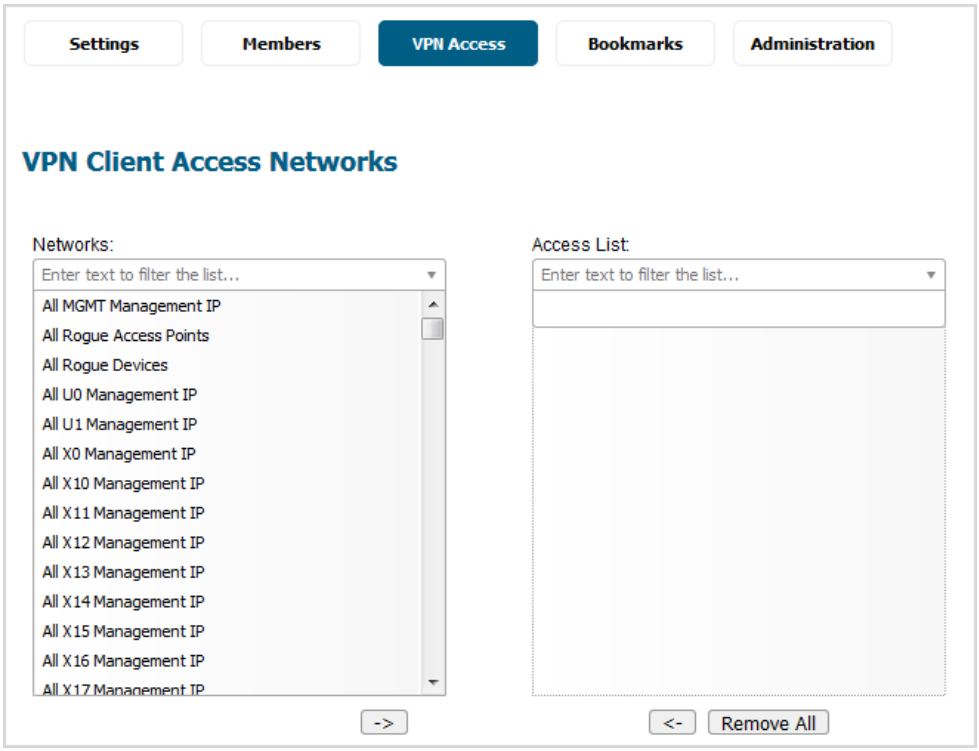
To remove users and/or groups, from the **Member Users and Groups** list, select the user(s) and/or group(s) and click the **Left Arrow <-**. To remove all users and groups, click **Remove All**.

4 To:

- Finish adding the group, click **OK**.
- Specify VPN access, go to VPN Access on page 239.

VPN Access

1 Click **VPN Access**.



2 From the **Networks** list, select the network resource(s) to which this group has VPN Access by default.

NOTE: Group VPN access settings affect remote clients and SSL VPN Virtual Office bookmarks.

3 Click the **Right Arrow ->** to add the resource(s) to **Access List**.

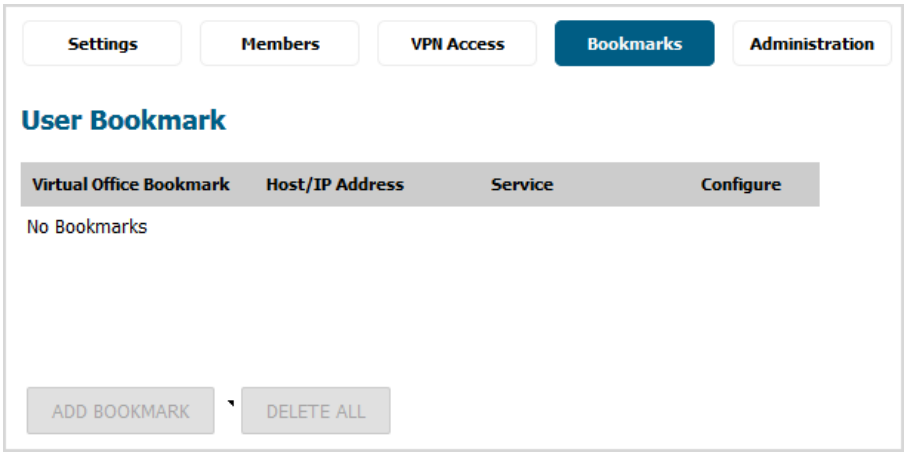
To remove resource(s), from **Access List**, select the resource(s) and click the **Left Arrow <-**. To remove resources, click **Remove All**.

4 To:

- Finish adding the group, click **OK**.
- Specify bookmarks, go to Bookmarks on page 240.

Bookmarks

1 Click **Bookmarks**.



2 You can add, edit, or delete Virtual Office bookmarks for each user who is a member of a related group. For information on configuring SSL VPN bookmarks, see the *SonicOS 6.5 Connectivity*.

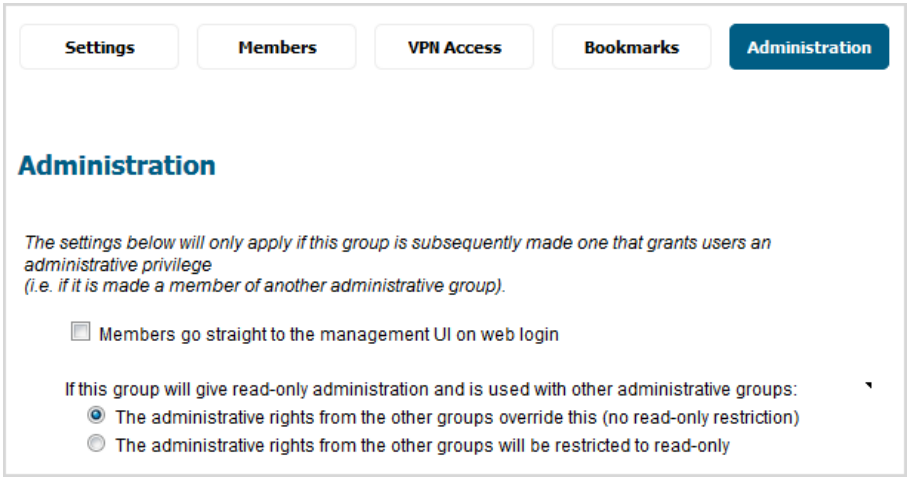
NOTE: Users must be members of the SSLVPN Services group before you can configure Bookmarks for them.

3 To:

- Finish adding the group, click **OK**.
- Specify whether the group has administration privileges, go to Administration on page 240.

Administration

1 Click **Administration**.



2 If the new group is to be made an administrative group by giving it membership in another administrative group, select **Members go straight to the management UI on web login**. This option is not selected by default.

3 The **If this read-only admin group is used with other administrative groups** options control what happens when users start with membership in a user group that gives read-only administration (that is, the SonicWall Read-Only Admins group or one with membership in it) and then are added to other administrative user groups. To give the user the:

- Admin rights set by their other administrative groups with no read-only restriction, choose **The administrative rights from the other groups override this (no read-only restriction)**. This setting allows the read-only admin group to be the default for a set of users, but then overrides the default for selected users by making them members of other administrative groups so they can do configuration. This option is selected by default. In the **Local Users** table, the **Admin** column for the user displays the other group's designation, such as *Ltd* or *"Full"*.
- To give member users the administration level set by their other groups, but restrict them to read only access, select **The administrative rights from the other groups will be restricted to read-only**. In the **Local Users** table, the **Admin** column for the user displays the dual designation, such as *Rd-Only Ltd*.

  TIP: To do a mix of both, select the first option for SonicWall Read-Only Admins, and then create another group that is a member of this group, but that has the second option selected (but not *vice versa*).

  NOTE: If a user is a member of a read-only admin group and has membership in no other administrative groups, then that member gets full level access (as per SonicWall Administrators) restricted to read-only.

4 Click **OK** to complete the configuration.

Editing a Local Group

To edit a local group:

1 Click the **Edit** icon of the group that you want to edit. The Edit Group dialog displays, and is the same as the Add Group

2 Follow the steps in Adding a Local Group on page 235.

Importing Local Groups from LDAP

Having user groups in SonicOS with the same name as existing LDAP/AD user groups allows SonicWall group memberships and privileges to be granted upon successful LDAP authentication. You can configure local user groups in SonicOS by retrieving the user group names from your LDAP server. For further information about importing local groups, see Users & Groups Page on page 157.

Setting User Membership by LDAP Location

You can set LDAP rules and policies for users located in certain Organizational Units (OUs) on the LDAP server. For more information about the LDAP Group Membership by Organizational Unit feature, see LDAP Group Membership by Organizational Unit on page 91. For the full procedure for creating new members, see Creating a New User Group for RADIUS Users on page 147.

11

Managing Guest Services

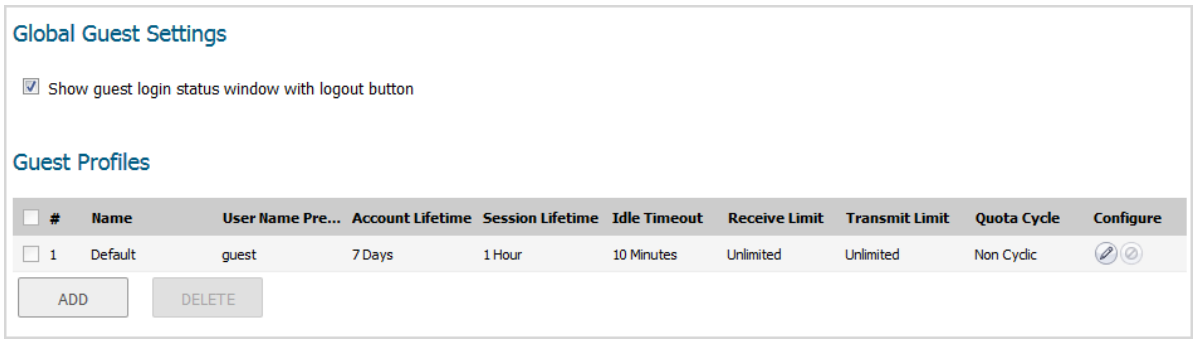
Topics:

- **Users > Guest Services** on page 242
  - Global Guest Settings on page 243
  - Guest Profiles on page 243

Users > Guest Services

Guest accounts are temporary accounts set up for users to log into your network. You can create these accounts manually, as needed or generate them in batches. SonicOS includes profiles you can configure in advance to automate configuring guest accounts when you generate them. Guest accounts are typically limited to a pre-determined life-span. After their life span, by default, the accounts are removed.

Guest Services determine the limits and configuration of the guest accounts. **MANAGE | System Setup > Users > Guest Services** page displays a list of Guest Profiles. Guest profiles determine the configuration of guest accounts when they are generated. In **Users > Guest Services**, you can add, delete, and configure Guest Profiles. In addition, you can determine if all users who log in to the Security Appliance see a user login window that displays the amount of time remaining in their current login session.

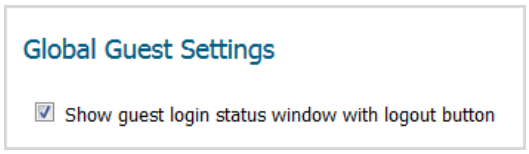


Topics:

- Global Guest Settings on page 243
- Guest Profiles on page 243

Global Guest Settings

The Global Guest Settings section provides an option for displaying the guest login status window. The window displays the time remaining in their current session. Users must keep this window open during their login session and can log out by clicking **Logout** in the login status window.



To configure the guest login status window:

1. Select **Show guest login status window with logout** to display **Logout** on the user's login window whenever the user is logged in. This option is selected by default.
2. Click **ACCEPT**.

Guest Profiles

The **Guest Profiles** table lists the profiles you have created and enables you to add, edit, and delete these profiles. There is always one guest profile, **Default**, which is generated by SonicOS and cannot be deleted, although you can edit it.



Topics:

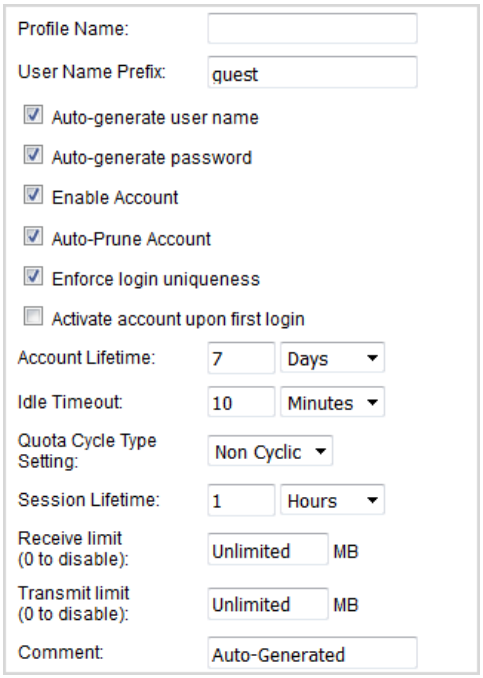
- Adding a Guest Profile on page 243
- Editing Guest Profiles on page 245
- Deleting Guest Profiles on page 245

Adding a Guest Profile

To add a profile:

1. Navigate to **Manage | System Setup > Users > Guest Services**.

2 Click **ADD** below the **Guest Profile** table. The **Add Guest Profile** dialog displays.



3 In the **Profile Name** field, enter the name of the profile.

4 In the **User Name Prefix** field, enter the first part of every user account name generated from this profile. To allow guest accounts generated from this profile to have an automatically generated user name, select **Auto-generate user name**. The user name is usually the prefix plus a two- or three-digit number. This option is selected by default.

5 To allow guest accounts generated from this profile to have an automatically generated password, select **Auto-generate password**. The generated password is an eight-character, unique alphabetic string. This option is selected by default.

6 For all guest accounts generated from this profile to be enabled upon creation, select **Enable Account**. This option is selected by default.

7 To have the account removed from the database after its lifetime expires, select **Auto-Prune Account**. This option is selected by default.

8 To allow only a single instance of an account to be used at any one time, select **Enforce login uniqueness**. By default, this feature is enabled when creating a new guest account. If you want to allow multiple users to login with a single account, disable this enforcement by clearing the **Enforce login uniqueness** checkbox.

9 To delay the Account Expiration timer until a user logs into the account for the first time, select **Activate Account Upon First Login**. This option is not selected by default.

10 To define how long an account remains on the Security Appliance before the account expires, enter the duration in **Account Lifetime**. You can specify from 1 to 9999 in the **Account Lifetime** field and select the type of duration from the drop-down menu:

- **Minutes**
- **Hours**
- **Days**

The default is **7 Days**.

11 To define the maximum period of time when no traffic is passed on an activated guest services session, enter the timeout duration in **Idle Timeout**:. Exceeding the period defined by this setting expires the session, but the account itself remains active as long as the **Account Lifetime** hasn't expired. The **Idle Timeout** cannot exceed the value set in the **Session Lifetime**.

You can specify from 1 to 9999 in the **Account Lifetime** field and select the type of duration from the drop-down menu:

- **Minutes**
- **Hours**
- **Days**

The default is **10 Minutes**.

12 To specify the quota cycle type, select from the **Quota Cycle Type Setting** drop-down menu:

- **Non Cyclic** (default)
- **Per Day**
- **Per Week**
- **Per Month**

13 To define how long a guest login session remains active after it has been activated, specify the duration in **Session Lifetime**. By default, activation occurs the first time a guest user logs into an account. The **Session Lifetime** cannot exceed the value set in the **Account Lifetime**.

You can specify from 1 to 9999 in the **Session Lifetime** field and select the type of duration from the drop-down menu:

- **Minutes**
- **Hours**
- **Days**

The default is **1 Hours**.

14 To limit the amount of data the user can receive, enter the amount, in MB, in **Receive limit (0 to disable)** field. The range is from 0 (no data can be received) to 999999999 MB to **Unlimited** (default).

15 To limit the amount of data the user can send, enter the amount, in MB, in **Transmit limit (0 to disable)** field. The range is from 0 (no data can be received) to 999999999 MB to **Unlimited** (default).

16 Optionally, enter a descriptive comment in the **Comment** field. The default is **Auto-Generated**.

17 Click **OK**.

Editing Guest Profiles

To edit guest profiles:

1 Click the **Edit** icon in the **Configure** column for the profile.

2 Follow the steps in Adding a Guest Profile on page 243.

NOTE: When editing the **Default** profile, you can edit all options except **Profile Name** and **User Name Prefix**; these options are dimmed.

Deleting Guest Profiles

You can delete all guest profiles except the **Default** profile.

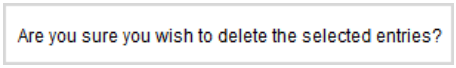
To delete guest profiles:

1 Select either:

- The checkbox(es) of the guest profile(s) to be deleted.
- The checkbox in the **Guest Profiles** table. All checkboxes (except for the **Default** profile) become selected.

DELETE becomes active.

2 Click **DELETE**. A confirmation message displays:



3 Click **OK**.

12

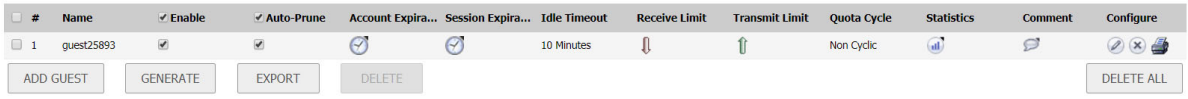
Managing Guest Accounts

Topics:

- **Users > Guest Accounts** on page 247
  - Viewing Guest Account Statistics on page 248
  - Adding Guest Accounts on page 250
  - Enabling Guest Accounts on page 255
  - Enabling Auto-Prune for Guest Accounts on page 255
  - Printing Account Details on page 257

Users > Guest Accounts

MANAGE | System Setup | Users > Guest Accounts lists the guest services accounts on the SonicWall Security Appliance. You can enable or disable individual accounts, groups of accounts, or all accounts, you can set the Auto-Prune feature for accounts, set an Account or Session Expiration date or time, and you can add, edit, delete, and print accounts.



Topics:

- Viewing Guest Account Statistics on page 248
- Adding Guest Accounts on page 250
- Enabling Guest Accounts on page 255
- Enabling Auto-Prune for Guest Accounts on page 255
- Printing Account Details on page 257

Viewing Guest Account Statistics

The **Guest Account** table displays statistics about the guest accounts.

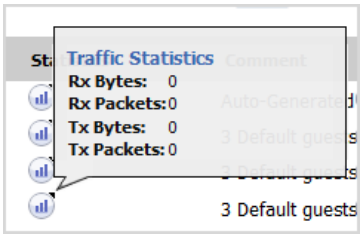
Topics:

- Viewing Traffic Statistics on page 248
- Viewing Account Expiration on page 248
- Viewing Session Expiration on page 248
- Viewing Receive and Transmit Limit Statistics on page 249
- Exporting Guest Accounts on page 249

Viewing Traffic Statistics

To view traffic statistics on a guest account:

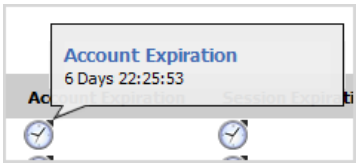
1 Hover your mouse over the **Statistics** icon in the **Statistics** column for the guest account. The **Traffic Statistics** popup displays the cumulative total bytes and packets sent and received for all completed sessions. Currently active sessions are not added to the statistics until the guest user logs out.



Viewing Account Expiration

To view the time remaining until the account expires:

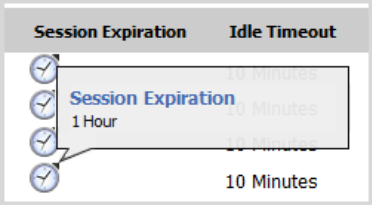
1 Hover your mouse over the **Clock** icon in the **Account Expiration** column for the guest account. The **Account Expiration** popup displays the remaining time for the guest account.



Viewing Session Expiration

To view the time remaining until the session expires:

1 Hover your mouse over the **Clock** icon in the **Account Expiration** column for the guest account. The **Account Expiration** popup displays the remaining time for the guest account.



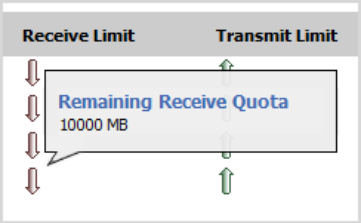
NOTE: If the user's session has not started, the **Session Expiration** popup reads `Unused`.

Viewing Receive and Transmit Limit Statistics

For each user account in the table, the **Receive Limit** column contains a red down arrow icon, and the **Transmit Limit** column contains a green up arrow icon.

To view the receive/transmit limit statistics:

1 Hover your mouse over the **Arrow** icon in the **Receive Limit/Transmit Limit** column for the guest account. The **Remaining Receive/Transmit Quota** popup displays the remaining amount of data the guest user can download or send.

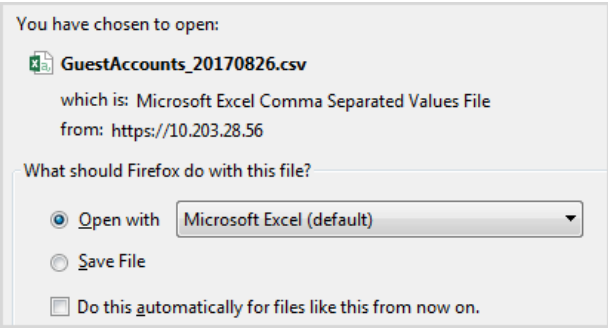


Exporting Guest Accounts

You can export the **Guest Accounts** table as a `.csv` file. This file contains not only all the displayed data, but the limit and remaining receive and transmitted data statistics.

To export the guest accounts as a .csv file:

1 Under the **Guest Accounts** table, click **Export**. The **Opening *guestaccounts_nnn*.csv** dialog displays.



2 You can:

- Open the file to view it.
- Save the file for later.

3 Click **OK**.

Adding Guest Accounts

You can add guest accounts individually or generate multiple guest accounts automatically.

Topics:

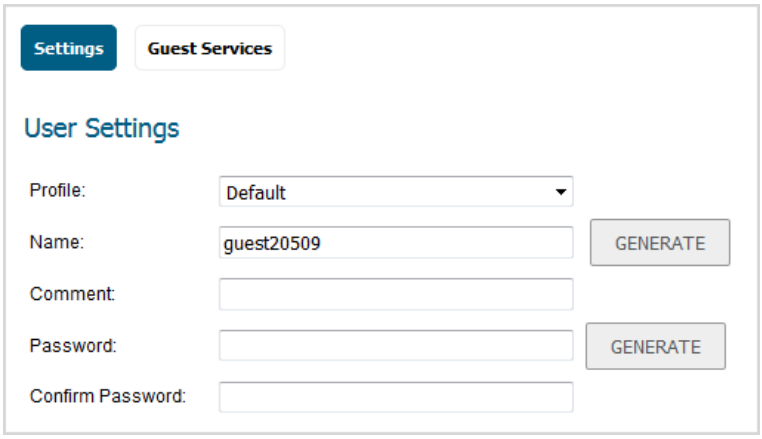
- Adding a Guest Account on page 250
- Generating Multiple Guest Accounts on page 253

Adding a Guest Account

To add an individual account:

1 Navigate to **MANAGE | System Setup > Users > Guest Accounts**.

2 Under the **Guest Accounts** table, click **Add Guest**. The **Add Guest** dialog displays.



3 From **Profile**, select the Guest Profile from which to generate this account. The default profile is **Default**.

4 Name the guest account by either:

- Entering a name for the account in the **Name** field.
- Clicking **Generate** to have SonicOS generate the name. The generated name is the first name of the profile, the word guest, and a random two- to five-digit number. For example:
  - guest1235 (for the Default profile)
  - TechPubs guest51026 (for the TechPubs Guest profile)

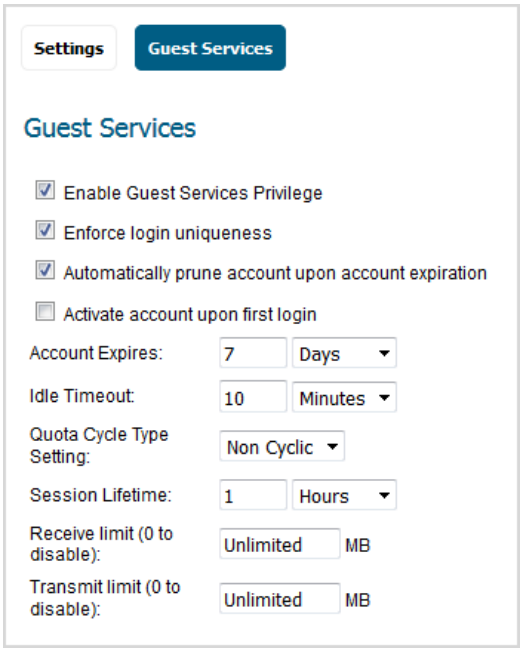
5 Enter a descriptive comment in the **Comment** field. The default comment is **Auto-Generated**.

6 Create a user account password by either:

- Entering the password in the **Password** field and the Confirm field. The password can be up to 32 alphanumeric characters.
- Clicking **Generate**. The generated password is a random string of eight alphabetic characters.

  TIP: Make a note of the password. Otherwise, you have to reset it.

7 Click **Guest Services**.



8 For the account to be enabled upon creation, select **Enable Guest Services Privilege**. This option is selected by default.

9 To allow only one instance of this account to log into the Security Appliance at one time, select **Enforce login uniqueness**. Clear it to allow multiple users to use this account simultaneously. This option is selected by default.

10 To have the account removed from the database after its lifetime expires, select **Automatically prune account upon account expiration**. This option is selected by default.

11 To begin the timing for the account expiration, select **Activate account upon first login**.

12 To define how long an account remains on the Security Appliance before the account expires, enter the expiration date in **Account Expires**. You can specify from 1 to 9999 in the **Account Expires** field and select the type of duration from the drop-down menu:

- **Minutes**
- **Hours**
- **Days**

The default is **7 Days**.

If **Automatically prune account upon account expiration** is:

- Enabled, the account is deleted when it expires.
- Disabled, the account remains in the **Guest Accounts** table with an **Expired** status to allow easy reactivation.

NOTE: This setting overrides the account lifetime set in Guest Profiles on page 243.

13 To define the maximum period of time when no traffic is passed on an activated guest services session, enter the timeout duration in **Idle Timeout**:. Exceeding the period defined by this setting expires the session, but the account itself remains active as long as the **Account Lifetime** hasn't expired. The **Idle Timeout** cannot exceed the value set in the **Session Lifetime**.

NOTE: This setting overrides the idle timeout setting in the profile.

You can specify from 1 to 9999 in the **Account Lifetime** field and select the type of duration from the drop-down menu:

- **Minutes**
- **Hours**
- **Days**

The default is **10 Minutes**.

14 To specify the quota cycle type, select from the **Quota Cycle Type Setting** drop-down menu:

- **Non Cyclic** (default)
- **Per Day**
- **Per Week**
- **Per Month**

15 To define how long a guest login session remains active after it has been activated, specify the duration in **Session Lifetime**. By default, activation occurs the first time a guest user logs into an account. The **Session Lifetime** cannot exceed the value set in the **Account Lifetime**.

NOTE: This setting overrides the session lifetime setting in the profile.

You can specify from 1 to 9999 in the **Session Lifetime** field and select the type of duration from the drop-down menu:

- **Minutes**
- **Hours**
- **Days**

The default is **1 Hours**.

16 **Receive limit (0 to disabled)**: Enter the number of megabytes the user is allowed to receive. The minimum number is 0, which disables the limit; the maximum is **Unlimited**, the default.

17 **Transmit limit (0 to disabled)**: Enter the number of megabytes the user is allowed to transmit. The minimum number is 0, which disables the limit; the maximum is **Unlimited**, the default.

18 To limit the amount of data the user can receive, enter the amount, in MB, in **Receive limit (0 to disable)** field. The range is from 0 (no data can be received) to 999999999 MB to **Unlimited** (default).

19 To limit the amount of data the user can send, enter the amount, in MB, in **Transmit limit (0 to disable)** field. The range is from 0 (no data can be received) to 999999999 MB to **Unlimited** (default).

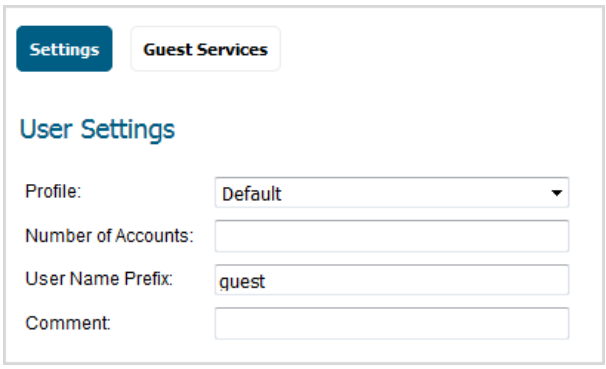
20 Click **OK** to generate the account.

Generating Multiple Guest Accounts

To generate multiple accounts:

1 Navigate to **MANAGE | System Setup > Users > Guest Accounts**.

2 Under the **Guest Accounts** table, click **Generate**. The **Generate Guest Accounts** dialog displays.



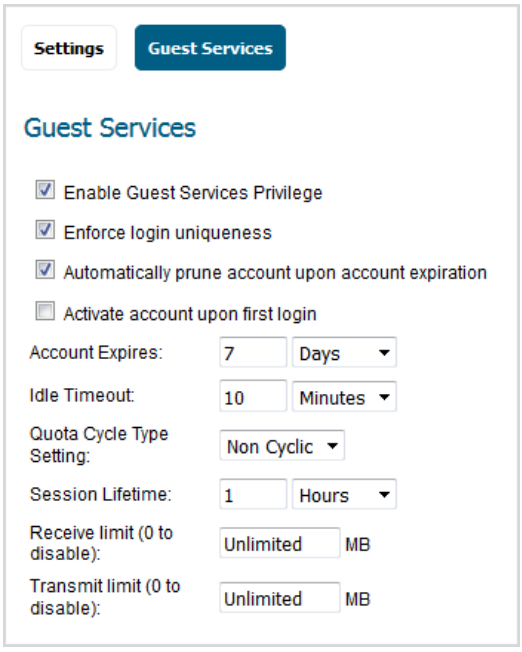
3 From **Profile**, select the Guest Profile from which to generate the accounts from. The default is **Default**.

4 Enter the number of accounts to generate in the **Number of Accounts** field. You can create from 1 to 6000 accounts

5 Enter the prefix from which account names are generated in the **User Name Prefix** field. For example, if you enter **Guest**, the generated accounts have names like `Guest123` and `Guest234`. The default prefix is **guest**.

6 Enter a descriptive comment of up to 16 alphanumeric characters in the **Comment** field.

7 Click **Guest Services**.



8 For the account to be enabled upon creation, select **Enable Guest Services Privilege**. This option is selected by default.

9 To allow only one instance of this account to log into the Security Appliance at one time, select **Enforce login uniqueness**. Clear it to allow multiple users to use this account simultaneously. This option is selected by default.

10 To have the account removed from the database after its lifetime expires, select **Automatically prune account upon account expiration**. This option is selected by default.

NOTE: This setting overrides the Auto-Prune setting in the guest profile, if they differ.

11 To begin the timing for the account expiration, select **Activate account upon first login**.

12 To define how long an account remains on the Security Appliance before the account expires, enter the expiration date in **Account Expires**. You can specify from 1 to 9999 in the **Account Expires** field and select the type of duration from the drop-down menu:

- **Minutes**
- **Hours**
- **Days**

The default is **7 Days**.

If **Automatically prune account upon account expiration** is:

- Enabled, the account is deleted when it expires.
- Disabled, the account remains in the **Guest Accounts** table with an **Expired** status to allow easy reactivation.

NOTE: This setting overrides the account lifetime set in Guest Profiles on page 243.

13 To define the maximum period of time when no traffic is passed on an activated guest services session, enter the timeout duration in **Idle Timeout**:. Exceeding the period defined by this setting expires the session, but the account itself remains active as long as the **Account Lifetime** hasn't expired. The **Idle Timeout** cannot exceed the value set in the **Session Lifetime**.

NOTE: This setting overrides the idle timeout setting in the profile.

You can specify from 1 to 9999 in the **Account Lifetime** field and select the type of duration from the drop-down menu:

- **Minutes**
- **Hours**
- **Days**

The default is **10 Minutes**.

14 To specify the quota cycle type, select from the **Quota Cycle Type Setting** drop-down menu:

- **Non Cyclic** (default)
- **Per Day**
- **Per Week**
- **Per Month**

15 To define how long a guest login session remains active after it has been activated, specify the duration in **Session Lifetime**. By default, activation occurs the first time a guest user logs into an account. The **Session Lifetime** cannot exceed the value set in the **Account Lifetime**.

NOTE: This setting overrides the session lifetime setting in the profile.

You can specify from 1 to 9999 in the **Session Lifetime** field and select the type of duration from the drop-down menu:

- **Minutes**
- **Hours**
- **Days**

The default is **1 Hours**.

16 **Receive limit (0 to disabled)**: Enter the number of megabytes the user is allowed to receive. The minimum number is 0, which disables the limit; the maximum is **Unlimited**, the default.

17 **Transmit limit (0 to disabled)**: Enter the number of megabytes the user is allowed to transmit. The minimum number is 0, which disables the limit; the maximum is **Unlimited**, the default.

18 To limit the amount of data the user can receive, enter the amount, in MB, in **Receive limit (0 to disable)** field. The range is from 0 (no data can be received) to 999999999 MB to **Unlimited** (default).

19 To limit the amount of data the user can send, enter the amount, in MB, in **Transmit limit (0 to disable)** field. The range is from 0 (no data can be received) to 999999999 MB to **Unlimited** (default).

20 Click **OK** to generate the accounts.

Enabling Guest Accounts

You can enable or disable any number of accounts at one time.

To enable one or more guest accounts:

1 Select the checkbox(es) in the **Enable** column next to the name(s) of the account(s) you want to enable. To enable all accounts, select the **Enable** checkbox in the table heading.

2 Click **Accept**.

Enabling Auto-Prune for Guest Accounts

You can enable or disable auto-prune for any number of accounts at one time. When auto-prune is enabled, the account is deleted after it expires.

NOTE: This overrides the Auto-Prune option set when configuring the user profile or the guest account.

To enable auto-prune:

1 Select the checkbox(es) in the **Auto-Prune** column next to the name(s) of the account(s). To enable it on all accounts, select the **Auto-Prune** checkbox in the table heading.

2 Click **Accept**.

Editing Guest Accounts

To edit guest accounts:

1 Click the **Edit** icon in the **Configure** column for the profile.

2 Follow the steps in Adding a Guest Profile on page 243.

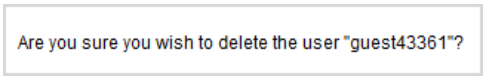
> **NOTE:** When editing the **Default** profile, you can edit all options except **Profile Name** and **User Name Prefix**; these options are dimmed.

Deleting Guest Accounts

You can delete all guest profiles except the **Default** profile.

To delete a guest account

1 Click the **Delete** icon for the guest account. A confirmation message displays:



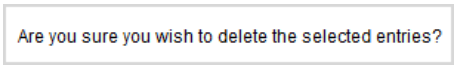
2 Click **OK**.

To delete one or more guest accounts:

1 Navigate to **MANAGE | System Setup | Users > Local Users & Groups**.

2 Select the checkbox(es) of the guest profile(s) to be deleted. **Delete** becomes active.

3 Click **DELETE**. A confirmation message displays:

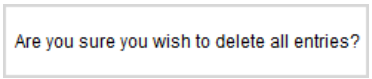


4 Click **OK**.

To delete all guest accounts:

1 Select the checkbox in header of the **Guest Accounts** table. All checkboxes (except for the **Default** profile) become selected. **DELETE ALL** becomes available.

2 Click **DELETE ALL**. A confirmation message displays:



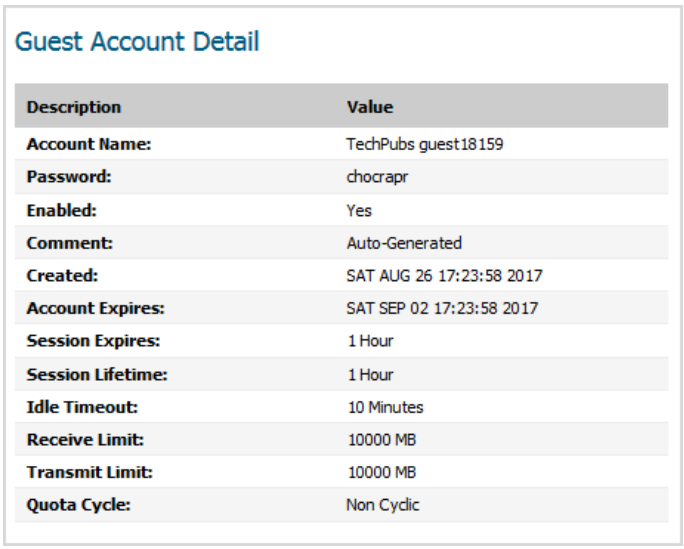
3 Click **OK**.

Printing Account Details

You can print a summary of a guest account.

To print details of a guest account.

1 Click the **Print** icon to display a summary account report and a **Print** dialog.



Guest Account Detail

Description	Value
Account Name:	TechPubs guest18159
Password:	chocrapr
Enabled:	Yes
Comment:	Auto-Generated
Created:	SAT AUG 26 17:23:58 2017
Account Expires:	SAT SEP 02 17:23:58 2017
Session Expires:	1 Hour
Session Lifetime:	1 Hour
Idle Timeout:	10 Minutes
Receive Limit:	10000 MB
Transmit Limit:	10000 MB
Quota Cycle:	Non Cyclic

2 Click **OK** to send the summary to a printer.

Part 4

System Setup | Network

- Configuring Interfaces
- Configuring PortShield Interfaces
- Configuring PoE
- Setting Up Failover and Load Balancing
- Configuring Network Zones
- Configuring Wire Mode VLAN Translation
- Configuring DNS Settings
- Configuring DNS Proxy Settings
- Configuring DNS Security
- Configuring Route Advertisements and Route Policies
- Managing ARP Traffic
- Configuring Neighbor Discovery Protocol
- Configuring MAC-IP Anti-spoof
- Setting Up the DHCP Server
- Using IP Helper
- Setting Up Web Proxy Forwarding
- Configuring Dynamic DNS
- Configuring AWS Credentials

13

Configuring Interfaces

Topics:

- About Interfaces on page 260
  - Physical Interfaces on page 260
  - Virtual Interfaces on page 262
  - SonicOS Secure Objects on page 264
  - Transparent Mode on page 264
  - IPS Sniffer Mode on page 264
  - Firewall Sandwich on page 266
  - HTTP/HTTPS Redirection on page 267
  - Enabling DNS Proxy on an Interface on page 267
- **Network > Interfaces** on page 268
  - Show/Hide PortShield Interfaces (IPv4 Only) on page 269
  - Interface Settings on page 270
  - Interface Traffic Statistics on page 271
- Configuring Interfaces on page 272
  - Configuring a Static Interface on page 272
  - Configuring Routed Mode on page 279
  - Enabling Bandwidth Management on an Interface on page 281
  - Configuring Interfaces in Transparent IP Mode (Splice L3 Subnet) on page 282
  - Configuring Wireless Interfaces on page 286
  - Configuring a WAN Interface on page 292
  - Configuring Tunnel Interfaces on page 297
  - Configuring Link Aggregation and Port Redundancy on page 301
  - Configuring Virtual Interfaces (VLAN Subinterfaces) on page 305
  - Configuring IPS Sniffer Mode on page 306
  - Configuring Security Services (Unified Threat Management) on page 309
  - Configuring Wire and Tap Mode on page 310
  - Wire Mode with Link Aggregation on page 313
  - Layer 2 Bridged Mode on page 315
  - Configuring Layer 2 Bridged Mode on page 334
  - Asymmetric Routing on page 341

- Configuring Interfaces for IPv6 on page 342
- 31-Bit Network on page 342
- PPPoE Unnumbered Interface Support on page 344

About Interfaces

- Physical Interfaces on page 260
- SonicOS Secure Objects on page 264
- Transparent Mode on page 264
- IPS Sniffer Mode on page 264
- Firewall Sandwich on page 266
- HTTP/HTTPS Redirection on page 267
- Enabling DNS Proxy on an Interface on page 267

Physical Interfaces

SonicOS Physical interfaces are bound to a single port.

Topics:

- About SFP/SFP+ Interfaces on page 261
- About 10 Gigabit Ethernet SFP+ Ports on NSA 6600 and SuperMassive 9000 Series on page 262

The front panel of a SonicWall Security Appliance has a number of physical interfaces. The number and type of interfaces depend on the model and version (for more information about the interfaces on your appliance, see the *Quick Start Guide* or *Getting Started Guide* for your appliance):

Physical Interfaces by Model Number

Interface Port		Comment
1 GE	High-speed copper Gigabit Ethernet ports	NS*a* platforms and SuperMassive 9000 series
1 GE SFP	1 Gigabit Ethernet hot-pluggable SFP interfaces	NS*a* platforms and SuperMassive 9000 series
2.5/1	GE copper ports	NS*a* platforms
2.5/1	GE SFP ports	NS*a* platforms
10/5/2.5	GE copper ports	NS*a* 5650 and higher NS*a* platforms
10 GE SFP+	10 Gigabit hot-pluggable ports	NS*a* 3600 Series and above and SuperMassive 9000 Series only; for further information, see About 10 Gigabit Ethernet SFP+ Ports on NSA 6600 and SuperMassive 9000 Series on page 262.

Physical Interfaces by Model Number (Continued)

Interface Port		Comment
10/5/2.5 GE SFP+	10/5/2.5 Gigabit Hot-pluggable ports	NS*a* 3650/4650/5650/6650/9250/9450/9650
MGMT	A 1 Gigabit Ethernet Management Interface port for secure firmware upgrading of the appliance in SafeMode. For more information about using the MGMT port for upgrading firmware in SafeMode, see the *SonicOS Upgrade Guide*. The default IP address for the MGMT port is `192.168.1.254`.	All NS*a* platforms and SuperMassive 9000 series

NOTE: For up-to-date details on port capabilities for specific SonicWall appliances, refer to the current data sheet.

Physical interfaces must be assigned to a zone to allow for configuration of Access Rules to govern inbound and outbound traffic. Security zones are bound to each physical interface where it acts as a conduit for inbound and outbound traffic. If there is no interface, traffic cannot access the zone or exit the zone. For more information on zones, see About Zones on page 390.

About SFP/SFP+ Interfaces

Overview

For an overview on using SFP/SFP+, refer to SonicWall 10 Gigabit Ethernet SFP+ Ports and 1 Gigabit Ethernet Ports.

For a list of third-party SFP/SFP+ modules, refer to Supported SFP and SFP+ Modules.

For a current list of SFP/SFP+ modules from SonicWall, see SonicWall SFP/SFP+ Transceiver Modules Reference Guide.

CAUTION: To avoid exceeding power supply limits and cooling capacity, use only Power Level I and Power Level II SFP/SFP+ modules as defined in the publication SFF-8419. Power Level II allows a maximum of roughly 1.6W (3.3VDC at 500mA).

SFP/SFP+ Ports and Modules

The port is the socket on the appliance front panel. The module is the removable interface unit that slides into the port. The following rules apply to their use.

1. If the appliance advertises SFP+ PORT with 10G capability then it will support SFP+ and SFP modules with these speeds (1G, 2.5G, 5G, 10G)
2. If the appliance advertises SFP+ PORT with 5G capability then it will support SFP+ and SFP modules with these speeds (1G, 2.5G, 5G)
3. If the appliance advertises SFP PORT with 2.5G capability then it will support SFP+ and SFP modules with these speeds (1G, 2.5G)
4. If the appliance advertises SFP PORT with 1G capability then it will support SFP+ and SFP modules with these speeds (1G)

So, for example:

- A 10G SFP+ MODULE plugs into an SFP PORT with 1G capable. It will work, but speed is limited to 1G.
- A 1G SFP MODULE plugs into an SFP+ PORT with 10G capable. It will work, but speed is limited to 1G.

About 10 Gigabit Ethernet SFP+ Ports on NSA 6600 and SuperMassive 9000 Series

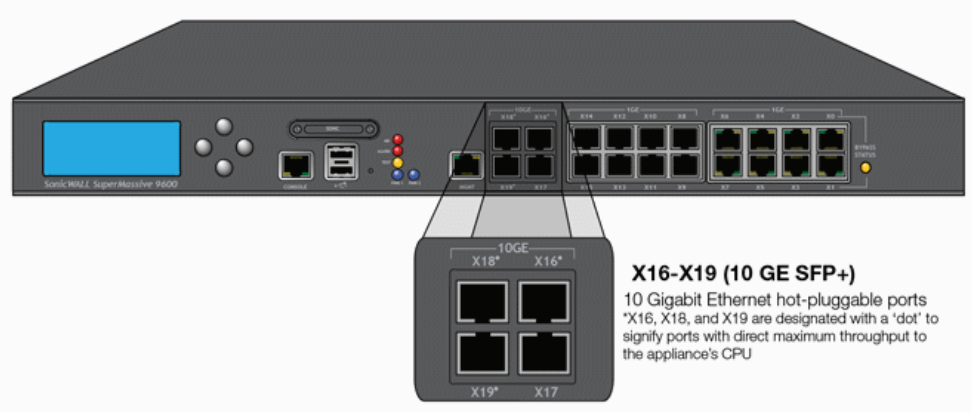
On NSA 6600 and the SuperMassive 9000 Series appliances, the enhanced small form-factor pluggable (SFP+) ports, X16, X18, and X19 are designated with a dot to signify that they have a direct maximum throughput to the CPU. These dotted ports have a dedicated (non-shared) uplink to the CPU.

This is beneficial, for example, if you have a 10Gb corporate network backbone and you are using a SuperMassive 9200 as the gateway device for your department. You should connect one of the dotted ports (X16, X18, or X19) directly to the backbone. This provides the fastest access because these ports are direct connections from the CPU to anything connected to them. You do not want the connection to the backbone sharing bandwidth with users or any other devices on your network. For maximum speed and efficiency, a dotted port should be connected directly to the backbone.

As another example, business-critical and heavily multiplexed links should also be connected to a dotted interface. An example of a business-critical use case might involve an administrative unit connecting to a 10Gb backbone network. For maximum performance, the upstream backbone connection should connect through a dotted interface. This ensures that important backbone traffic is never lost because of transient high load conditions on the other non-dotted interfaces that share a CPU uplink.

An example of a heavily multiplexed use case might involve a number of downstream enterprise switches that each have 10Gb uplinks. For maximum performance, each should be connected through a dotted interface. This ensures that differing high-level switching domains cannot starve each other of CPU resources.

10 Gigabit Ethernet hot-pluggable ports



The X17 interface is marked with an asterisk (*) in the SonicOS management interface to indicate that it is connected to the common switching domain shared with ports X0 - X15, thereby allowing X17 to participate in SonicOS advanced switching features.

Virtual Interfaces

Supported on SonicWall Security Appliances, virtual Interfaces are subinterfaces assigned to a physical interface. Virtual interfaces allow you to have more than one interface on one physical connection.

Virtual interfaces provide many of the same features as physical interfaces, including zone assignment, DHCP Server, and NAT and Access Rule controls.

Topics:

- Virtual LANs on page 263
- Subinterfaces on page 263

Virtual LANs

Virtual Local Area Networks (VLANs) can be described as a “tag-based LAN multiplexing technology” because through the use of IP header tagging, VLANs can simulate multiple LAN’s within a single physical LAN. Just as two physically distinct, disconnected LAN’s are wholly separate from one another, so too are two different VLANs; however, the two VLANs can exist on the very same wire. VLANs require VLAN aware networking devices to offer this kind of virtualization — switches, routers and firewalls that have the ability to recognize, process, remove and insert VLAN tags (IDs) in accordance with the network’s design and security policies.

VLANs are useful for a number of different reasons, most of which are predicated on the VLANs ability to provide logical rather than physical broadcast domain, or LAN boundaries. This works both to segment larger physical LAN’s into smaller virtual LAN’s, as well as to bring physically disparate LAN’s together into a logically contiguous virtual LAN. The benefits of this include:

- **Increased performance** – Creating smaller, logically partitioned broadcast domains decreases overall network utilization, sending broadcasts only where they need to be sent, thus leaving more available bandwidth for application traffic.
- **Decreased costs** – Historically, broadcast segmentation was performed with routers, requiring additional hardware and configuration. With VLANs, the functional role of the router is reversed – rather than being used for the purposes of inhibiting communications, it is used to facilitate communications between separate VLANs as needed.
- **Virtual workgroups** – Workgroups are logical units that commonly share information, such as a Marketing department or an Engineering department. For reasons of efficiency, broadcast domain boundaries should be created such that they align with these functional workgroups, but that is not always possible: Engineering and Marketing users might be commingled, sharing the same floor (and the same workgroup switch) in a building, or just the opposite – the Engineering team might be spread across an entire campus. Attempting to solve this with complex feats of wiring can be expensive and impossible to maintain with constant adds and moves. VLANs allow for switches to be quickly reconfigured so that logical network alignment can remain consistent with workgroup requirements.
- **Security** – Hosts on one VLAN cannot communicate with hosts on another VLAN unless some networking device facilitates communication between them.

Subinterfaces

VLAN support on SonicOS is achieved by means of subinterfaces, which are logical interfaces nested beneath a physical interface. Every unique (tag) requires its own subinterface. For reasons of security and control, SonicOS does not participate in any VLAN trunking protocols, but instead requires that each VLAN that is to be supported be configured and assigned appropriate security characteristics.

NOTE: VLAN IDs range from 0 – 4094, with these restrictions: VLAN 0 is reserved for QoS and VLAN 1 is reserved by some switches for native VLAN designation.

NOTE: Dynamic VLAN Trunking protocols, such as VTP (VLAN Trunking Protocol) or GVRP (Generic VLAN Registration Protocol), should not be used on trunk links from other devices connected to the firewall.

Trunk links from VLAN capable switches are supported by declaring the relevant VLAN ID’s as a subinterface on the firewall, and configuring them in much the same way that a physical interface would be configured. In other words, only those VLANs which are defined as subinterfaces are handled by the firewall, the rest are discarded as uninteresting. This method also allows the parent physical interface on the firewall to which a trunk link is connected to operate as a conventional interface, providing support for any native (untagged) VLAN traffic that might also exist on the same link. Alternatively, the parent interface may remain in an ‘unassigned’ state.

VLAN subinterfaces have most of the capabilities and characteristics of a physical interface, including zone assignability, security services, GroupVPN, DHCP server, IP Helper, routing, and full NAT policy and Access Rule controls. Multicast support is excluded from VLAN subinterfaces at this time.

SonicOS Secure Objects

The SonicOS scheme of interface addressing works in conjunction with address objects, service objects, and network zones. This structure is based on secure objects, which are utilized by rules and policies within SonicOS.

Secured objects include interface objects that are directly linked to physical interfaces and managed in the **Network > Interfaces** page. Address and Service Objects are defined in **MANAGE | Policies > Objects > Address Objects** and **MANAGE | Policies > Objects > Service Objects** respectively; for more information about address and service objects, see *SonicOS 6.5 Policies*.

Zones are the hierarchical apex of SonicOS's secure objects architecture. SonicOS includes predefined zones as well as allow you to define your own zones. Predefined zones include LAN, DMZ, WAN, WLAN, and Custom. For more information about zones, see Configuring Network Zones on page 390.

Zones can include multiple interfaces; the WAN zone, however, is restricted to a maximum of the total number of interfaces minus one. Within the WAN zone, either one or more WAN interfaces can be actively passing traffic depending on the WAN Failover and Load Balancing configuration on **Network > Failover & Load Balancing**. For more information on WAN Failover and Load Balancing on the SonicWall Security Appliance, see **Network > Failover & Load Balancing** on page 379.

At the zone configuration level, the **Allow Interface Trust** setting for zones automates the processes involved in creating a permissive intra-zone Access Rule. It creates a comprehensive Address Object for the entire zone and a inclusively permissive Access Rule from zone address to zone addresses.

Transparent Mode

Transparent Mode in SonicOS uses interfaces is the top level of the management hierarchy. Transparent Mode supports unique addressing and interface routing.

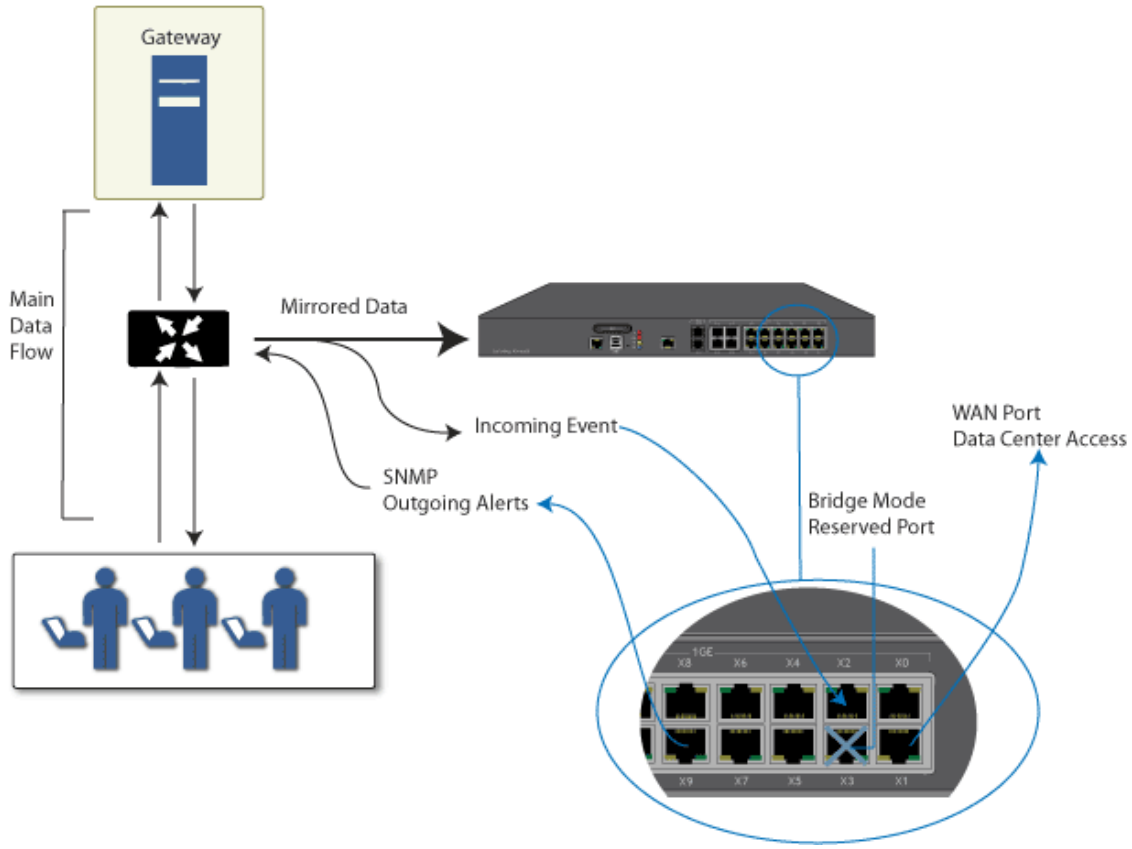
IPS Sniffer Mode

Supported on SonicWall Security Appliances, IPS Sniffer Mode is a variation of Layer 2 Bridged Mode that is used for intrusion detection. IPS Sniffer Mode configuration allows an interface on the Security Appliance to be connected to a mirrored port on a switch to examine network traffic. Typically, this configuration is used with a switch inside the main gateway to monitor traffic on the intranet.

In IPS Sniffer Mode: Network diagram, traffic flows into a switch in the local network and is mirrored through a switch mirror port into a IPS Sniffer Mode interface on the SonicWall Security Appliance. The Security Appliance inspects the packets according to the settings configured on the Bridge-Pair. Alerts can trigger SNMP traps which are sent to the specified SNMP manager through another interface on the Security Appliance. The network traffic is discarded after the Security Appliance inspects it.

The WAN interface of the Security Appliance is used to connect to the firewall Data Center for signature updates or other data.

IPS Sniffer Mode: Network diagram



In IPS Sniffer Mode, a Layer 2 Bridge is configured between two interfaces in the same zone on the Security Appliance, such as LAN-LAN or DMZ-DMZ. You can also create a custom zone to use for the Layer 2 Bridge.

Only the WAN zone is **not** appropriate for IPS Sniffer Mode. The reason for this is that SonicOS detects all signatures on traffic within the same zone such as LAN-LAN traffic, but some directional specific (client-side versus server-side) signatures do not apply to some LAN-WAN cases.

Either interface of the Layer 2 Bridge can be connected to the mirrored port on the switch. As network traffic traverses the switch, the traffic is also sent to the mirrored port and from there into the Security Appliance for deep packet inspection. Malicious events trigger alerts and log entries, and if SNMP is enabled, SNMP traps are sent to the configured IP address of the SNMP manager system. The traffic does not actually continue to the other interface of the Layer 2 Bridge. IPS Sniffer Mode does not place the Security Appliance inline with the network traffic, it only provides a way to inspect the traffic.

The **Edit Interfaces** dialog available from the **Network > Interfaces** page provides an option, **Only sniff traffic on this bridge-pair**, for use when configuring IPS Sniffer Mode. When selected, this option causes the Security Appliance to inspect all packets that arrive on the L2 Bridge from the mirrored switch port. The **Never route traffic on this bridge-pair** option should also be selected for IPS Sniffer Mode to ensure that the traffic from the mirrored switch port is not sent back out onto the network.

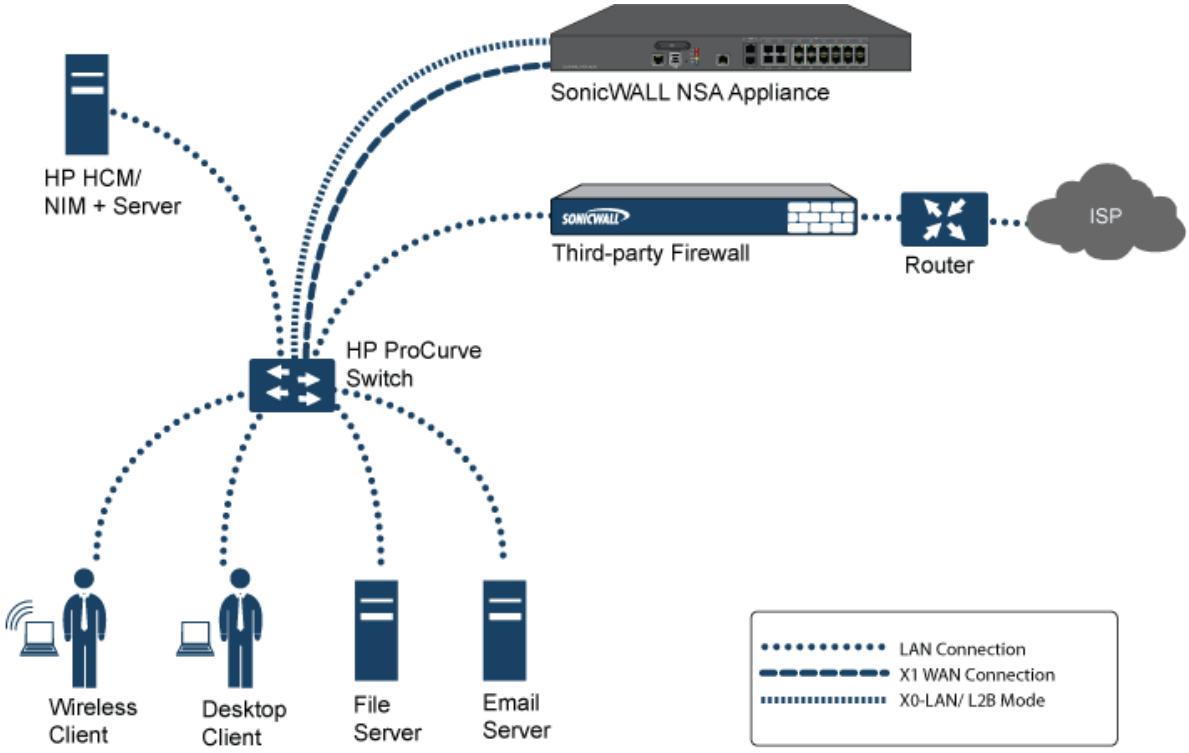
For detailed instructions on configuring interfaces in IPS Sniffer Mode, see Configuring IPS Sniffer Mode on page 306.

Sample IPS Sniffer Mode Topology

This example topology uses SonicWall IPS Sniffer Mode in a Hewlett Packard ProCurve switching environment. This scenario relies on the ability of HP's ProCurve Manager Plus (PCM+) and HP Network Immunity Manager (NIM) server software packages to throttle or close ports from which threats are emanating.

This method is useful in networks where there is an existing Security Appliance that remains in place, but you wish to use the Security Appliance's security services as a sensor.

IPS Sniffer Mode: Sample topology



In this deployment the WAN interface and zone are configured for the *internal* network's addressing scheme and attached to the internal network. The X2 port is Layer 2 bridged to the LAN port, but it is not attached to anything. The X0 LAN port is configured to a second, specially programmed port on the HP ProCurve switch. This special port is set for mirror mode: it forwards all the internal user and server ports to the "sniff" port on the firewall. This allows the firewall to analyze the entire internal network's traffic, and if any traffic triggers the security signatures it immediately traps out to the PCM+/NIM server through the X1 WAN interface, which then can take action on the specific port from which the threat is emanating.

Firewall Sandwich

You can deploy and configure a SonicWall Firewall Sandwich to improve availability, scalability, and manageability across the IT infrastructure. Deployment of the Firewall Sandwich provides the following features:

- Scalability - add more capacity as you go, reusing existing equipment
- Redundancy and resiliency – primary and secondary components
- Inline upgrades – upgrade firewalls and switches without shutting down the system
- Single point of management - manage policies for multiple firewall clusters and blades
- Full security services - including DPI-SSL capability

Firewall Sandwich deployment and configuration can be implemented using these equipment and services:

- Dell Force10 S series switches, such as the S5000, S4810, S4048, or S6000 running FTOS v9.8+
- SonicWall NSA 2600 and higher appliances or SuperMassive series appliances.

- SonicWall services, such as GAV, IPS, ASPR, DPI-SSL, and CFS in conjunction with Single Sign-On All in Wire Mode.

HTTP/HTTPS Redirection

When the Security Appliance configuration requires user authentication, HTTP/HTTPS traffic from an unauthenticated source is redirected to the SonicOS login screen for the user to enter their credentials. A problem occurs when HTTP and HTTPS traffic arrive from sources from which users do not log in, and one or more such sources repeatedly try to open new connections, which keeps triggering this redirection. These could be non-user devices that are validly trying to get access or could be malicious code attempting a Denial of Service (DoS) attack. The effect that it has on the Security Appliance is to cause high CPU load in the CP, both in the data plane task initiating the redirections and in the web server thread tasks that are serving up the target redirect pages.

To minimize this effect, ensure the **Add rule to enable redirect from HTTP to HTTPS** option is selected when adding or editing an interface. Enabling this option causes SonicOS to add an access rule that allows HTTP to the interface; a side effect of this rule is that it also allows SonicOS to be able to redirect HTTPS to HTTP in certain cases without security issues. One such case is the first step of redirecting traffic that needs to be authenticated, at which point there is no sensitive data that needs to be hidden. Then HTTP processing can occur on the data plane (DP) rather than on the CP.

NOTE: This option is not available when adding or editing VPN tunnel interfaces or when **Wire Mode (2-Port Wire)**, **Tap Mode (1-Port Tap)**, or **PortShield Switch Mode** is selected for **Mode/IP Assignment**.

HTTP/HTTPS Redirection with DP Offload

This feature improves handling of HTTP/HTTPS redirection requests that occur when user authentication is required for users to get access through the Security Appliance. HTTP/HTTPS requests received from sources that are not authenticated users are redirected to the Security Appliance's login page, which is served up by its built-in web server. This redirection happens if Single Sign-On (SSO) cannot identify the user or if SSO is not in use.

This feature improves efficiencies in both the web server and the HTTP/HTTPS redirection processes, and offloads most of the redirection processes to the Data Plane (DP) where the processing can be spread across multiple cores.

NOTE: Elements of this feature can be controlled by internal User Authentication Settings options. This includes an option to globally enable/disable redirection processing in the DP, a flush option to clear the redirect files cache, and an option to specify the internal NAT port number used for the web server. Contact SonicWall Technical Support for information about internal settings.

Enabling DNS Proxy on an Interface

When DNS Proxy is enabled globally, you can enable it on individual interfaces. This allows you to enable the feature for different network segments independently. For how to enable DNS Proxy on an interface, see Enabling DNS Proxy on page 439.

LTE Modem Support

When a LTE USB modem is connected to a SonicWall Security Appliance, SonicOS detects the model and displays a U0 interface in the **MANAGE | System Setup > Network > Interfaces** page. This interface belongs to the WAN zone by default and can be used for Failover and Load Balancing as well as for configuring the LTE connection, profiles and advanced settings. For further information about LTE modems, see *SonicOS 6.5 Connectivity*.

LAN Bypass

In the NS*a* 6650, NS*a* 9250, NS*a* 9450, and NS*a* 9650 platforms, the hardware (LAN) bypass mode is enabled in both Wire Mode and L2 Bridging. The main functionality of the LAN Bypass feature, when enabled:

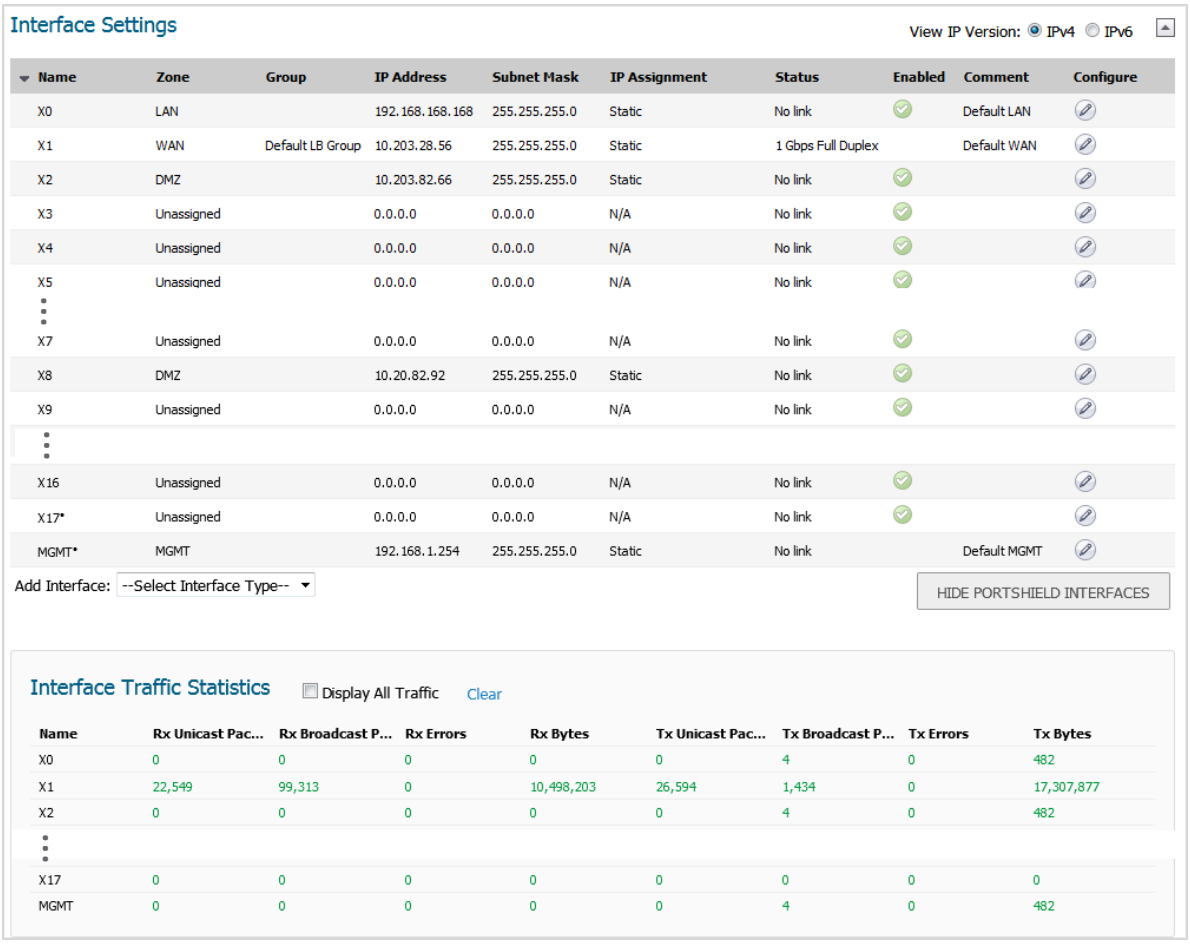
- Pass traffic in between the LBP-capable interfaces while rebooting.
- Even when the firewall is powered off, pass traffic in between those LBP-capable Interfaces.

For the NS*a* 9250, NS*a* 9450, and NS*a* 9650 platforms, the LAN Bypass feature is available between interfaces X26 and X27. For the NS*a* 6650, the feature is available between X0 and X1.

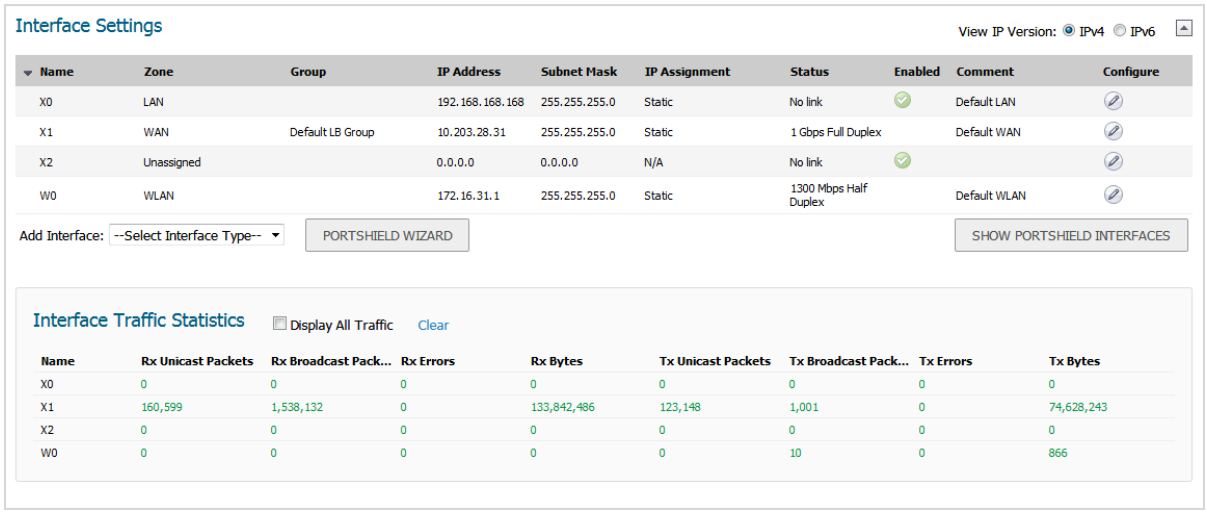
Network > Interfaces

The **Network > Interfaces** page includes interface objects that are directly linked to physical interfaces. The SonicOS scheme of interface addressing works in conjunction with network zones and address objects. There are some minor differences between the NSA 2600 and above Security Appliances and the TZ and SOHO Security Appliances. These differences are noted where they occur.

NSA 2600 and above Security Appliances



TZ Series and SOHO Security Appliances



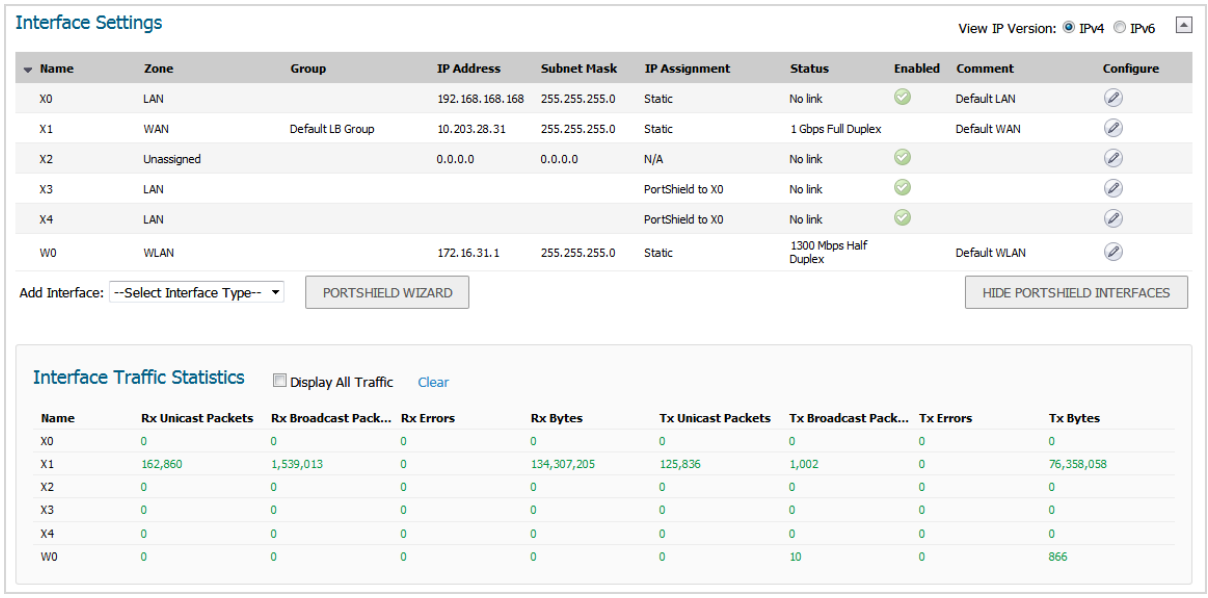
Topics:

- Show/Hide PortShield Interfaces (IPv4 Only) on page 269
- Interface Settings on page 270
- Interface Traffic Statistics on page 271
- Physical Interfaces on page 260
- SonicOS Secure Objects on page 264
- Transparent Mode on page 264
- IPS Sniffer Mode on page 264
- Configuring Interfaces on page 272
- Configuring IPS Sniffer Mode on page 306
- Configuring Wire and Tap Mode on page 310
- Wire Mode with Link Aggregation on page 313
- Layer 2 Bridged Mode on page 315
- Configuring Layer 2 Bridged Mode on page 334
- Configuring Interfaces for IPv6 on page 342
- 31-Bit Network on page 342
- PPPoE Unnumbered Interface Support on page 344

Show/Hide PortShield Interfaces (IPv4 Only)

In IPv4 mode, you can show PortShield interfaces in the **Interface Settings** and **Interface Traffic Statistics** tables by clicking **Show PortShield Interfaces**. The PortShield Interfaces display, and the button becomes **Hide PortShield Interfaces**.

Show/Hide PortShield Interfaces



To hide the PortShield interfaces, click **Hide PortShield Interfaces**.

Interface Settings

The **Interface Settings** table lists this information for each interface:

- **Name** - The name of the interface.
- **Zone** - LAN, WAN, and WLAN are listed by default as are DMZ and MGMT when applicable. As zones are configured, the names are listed in this column. Non-configured zones are designated **Unassigned**. Mousing over the zone displays zone properties:



Security Type	Displays security type selected for the zone when it was configured.
Member Interfaces	Lists interfaces assigned to this zone.
Interface Trust	Indicates whether **Allow Interface Trust** is enabled for this zone.
Anti-Virus	Indicates whether **Enable Client AV Enforcement Service** and/or **Enable Gateway Anti-Virus Service** is enabled for this zone.
SEC	
GSC	Indicates whether **Enforce Global Security Clients** (GSC) protection is enabled for this zone. For more information, see Enabling SonicWall Security Services on Zones on page 393.

- **Group** - If the interface is assigned to a Load Balancing group, it is displayed in this column.
- **IP Address** - IP address assigned to the interface.

- **Subnet Mask** - The network mask assigned to the subnet.
- **IP Assignment** - The available methods of IP assignment depend on the zone to which the interface is assigned:

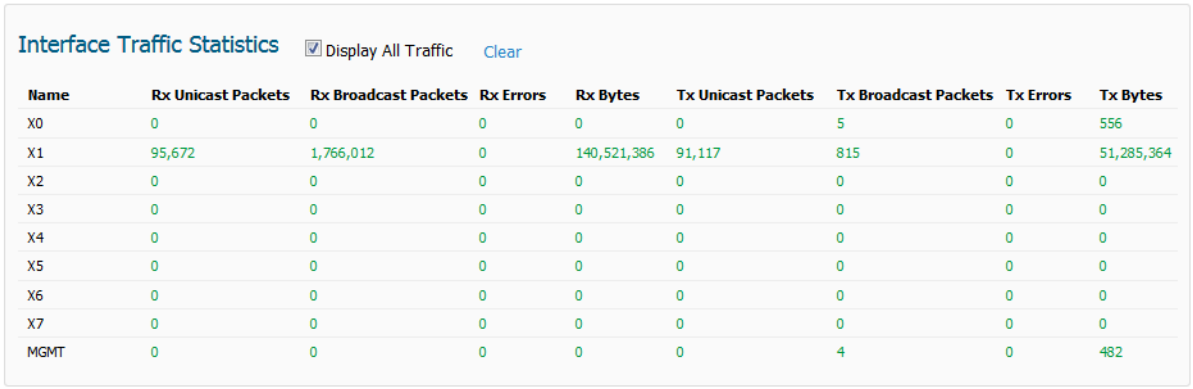
  NOTE: Wire mode is available only on NSA 2600 and higher Security Appliances.

LAN	**Static IP Mode** (default), **Transparent IP Mode (Splice L3 Subnet)**, **Layer 2 Bridged Mode (IP Route Option)**, **Wire Mode (2-Port Wire)**, **Tap Mode (1-Port Tap)**, **IP Unnumbered**, **Portshield Switch Mode**, **NativeBridge Mode**
WAN	**Static** (default), **DHCP**, **PPPoE**, **PPTP**, **L2TP**, **Wire Mode, (2-Port Wire)**, **Tap Mode (1-Port Tap)**
DMZ	**Static IP Mode** (default), **Transparent IP Mode (Splice L3 Subnet)**, **Layer 2 Bridged Mode (IP Route Option)**, **Wire Mode (2-Port Wire)**, **Tap Mode (1-Port Tap)**, **IP Unnumbered**, **Portshield Switch Mode**, **NativeBridge Mode**
WLAN	**Static IP Mode** (default), **Portshield Switch Mode**, **Layer 2 Bridged Mode**, **NativeBridge Mode**
PortShield to X*n* (IPv4 view only)	If PortShield interfaces are configured, the PortShield assignment

- **Status** - The link status and speed.
- **Enabled** - Indicates ports that can be enabled/disabled through **Network > Interfaces**. Ports that are enabled are indicated by an **Enabled** icon, those that are disabled by a **Disabled** icon. Clicking on the icon displays a message verifying you want the port enabled/disabled. Click **OK**. The port is enabled/disabled, and the icon changes.
- **Comment** - Any user-defined comments.
- **Configure** - Click the **Edit** icon to display the **Edit Interface** dialog, which allows you to configure the settings for the specified interface. For information about configuring interfaces, see Configuring Interfaces on page 272.

Interface Traffic Statistics

The **Interface Traffic Statistics** table lists, for each interface, received and transmitted information for all configured interfaces, including VLAN sub-interfaces:



Interface Traffic Statistics Display All Traffic Clear

Name	Rx Unicast Packets	Rx Broadcast Packets	Rx Errors	Rx Bytes	Tx Unicast Packets	Tx Broadcast Packets	Tx Errors	Tx Bytes
X0	0	0	0	0	0	5	0	556
X1	95,672	1,766,012	0	140,521,386	91,117	815	0	51,285,364
X2	0	0	0	0	0	0	0	0
X3	0	0	0	0	0	0	0	0
X4	0	0	0	0	0	0	0	0
X5	0	0	0	0	0	0	0	0
X6	0	0	0	0	0	0	0	0
X7	0	0	0	0	0	0	0	0
MGMT	0	0	0	0	0	4	0	482

Name	Name of the interface.
Rx Unicast Packets	Number of point-to-point communications received by the interface.

Rx Broadcast Packets or Rx Multicast Packets	Number of multipoint communications received by the interface.
RX Errors	Number of errors received by the interface.
RX Bytes	Volume of data, in bytes, received by the interface.
Tx Unicast Packets	Number of point-to-point communications transmitted by the interface.
Tx Broadcast Bytes	Number of mutlipoint communications received by the interface.
TX Errors	Number of errors transmitted b the interface.
Tx Bytes	Volume of data, in bytes, transmitted by the interface.

To clear the current statistics, click **Clear** at the top of the **Interface Traffic Statistics** table.

Configuring Interfaces

Topics:

- Configuring a Static Interface on page 272
- Configuring Routed Mode on page 279
- Enabling Bandwidth Management on an Interface on page 281
- Configuring Interfaces in Transparent IP Mode (Splice L3 Subnet) on page 282
- Configuring Wireless Interfaces on page 286
- Configuring a WAN Interface on page 292
- Configuring Tunnel Interfaces on page 297
- Configuring Link Aggregation and Port Redundancy on page 301
- Configuring Virtual Interfaces (VLAN Subinterfaces) on page 305
- Configuring IPS Sniffer Mode on page 306
- Configuring Security Services (Unified Threat Management) on page 309
- Configuring Wire and Tap Mode on page 310
- Wire Mode with Link Aggregation on page 313
- Layer 2 Bridged Mode on page 315
- Configuring Layer 2 Bridged Mode on page 334
- Asymmetric Routing on page 341
- Configuring Interfaces for IPv6 on page 342
- 31-Bit Network on page 342
- PPPoE Unnumbered Interface Support on page 344

Configuring a Static Interface

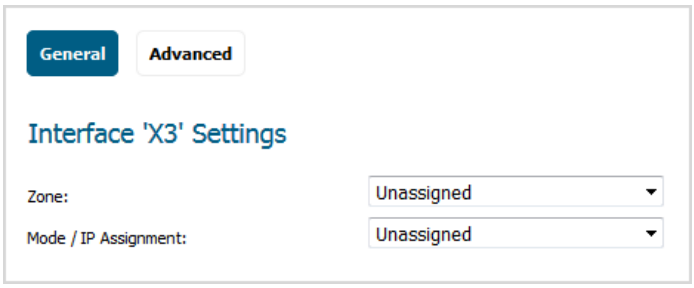
For general information on interfaces, see Physical Interfaces on page 260.

Static means that you assign a fixed IP address to the interface.

To configure a static interface:

1 Navigate to **MANAGE | Network > Interfaces**.

2 In the **Interface Settings** table, click the **Edit** icon for the interface you want to configure. The **Edit Interface** dialog displays.



3 Select a zone to assign to the interface from **Zone**:

- **LAN**
- **WAN**
- **DMZ**
- **WLAN**
- Custom zone you've created
- **Create new zone**. The **Add Zone** dialog is displayed. See About Zones on page 390 for instructions on adding a zone.

NOTE: The options displayed change, depending on the **Zone** you select.

4 From **IP Assignment** select:

- **Static** (default for WAN)
- **Static IP Mode** (default for LAN)

5 Enter the IP address and subnet mask for the interface into the **IP Address** and **Subnet Mask** fields.

NOTE: You cannot enter an IP address that is in the same subnet as another zone.

6 If configuring a:

- WAN zone interface or the MGMT interface, enter the IP address of the gateway device into the **Default Gateway** field.

  NOTE: A default gateway IP is required on the WAN interface if any destination is required to be reached through the WAN interface that is not part of the WAN subnet IP address space, regardless whether a default route is received dynamically from a routing protocol of a peer device on the WAN subnet. A default gateway IP is optional on a LAN interface.

- LAN zone interface or a DMZ zone interface, optionally enter the IP address of the gateway device into the **Default Gateway (Optional)** field.

The gateway device provides access between this interface and the external network, whether it is the Internet or a private network.

7 If configuring a:

- LAN zone interface, go to Step 8.

- WAN zone interface, enter the IP addresses of up to three DNS servers into the **DNS Server** fields. These can be public or private DNS servers. For more information, see Configuring a WAN Interface on page 292.

8 Enter any optional comment text in the **Comment** field. This text is displayed in the **Comment** column of the **Interface Settings** table.

9 If you want to enable remote management of the Security Appliance from this interface, choose the supported **Management** protocol(s): **HTTPS**, **Ping**, **SNMP**, and/or **SSH**. If **HTTPS** is chosen, **Add rule to enable redirect from HTTP to HTTPS** becomes available and selected; selecting **HTTP**, disables **Add rule to enable redirect from HTTP to HTTPS**, and it becomes dimmed.

NOTE: Elements of this feature can be controlled by internal User Authentication Settings options. For more information, see HTTP/HTTPS Redirection with DP Offload on page 267.

To allow access to the WAN interface for management from another zone on the same Security Appliance, access rules must be created. For more information about allowing WAN primary IP access from the LAN zone, see *SonicOS 6.5 Policies*.

10 If you want to allow selected users with limited management rights to log in to the Security Appliance, choose **HTTP** and/or **HTTPS** in **User Login**. If **HTTPS** is chosen, **Add rule to enable redirect from HTTP to HTTPS** becomes available and selected; selecting **HTTP**, disables **Add rule to enable redirect from HTTP to HTTPS**, and it becomes dimmed.

11 Either:

- Configure **Advanced Settings**; go to Configuring Advanced Settings for a Static Interface on page 275.
- Click **OK**.

NOTE: The administrator password is required to regenerate encryption keys after changing the Security Appliance's address.

Configuring Advanced Settings for a Static Interface

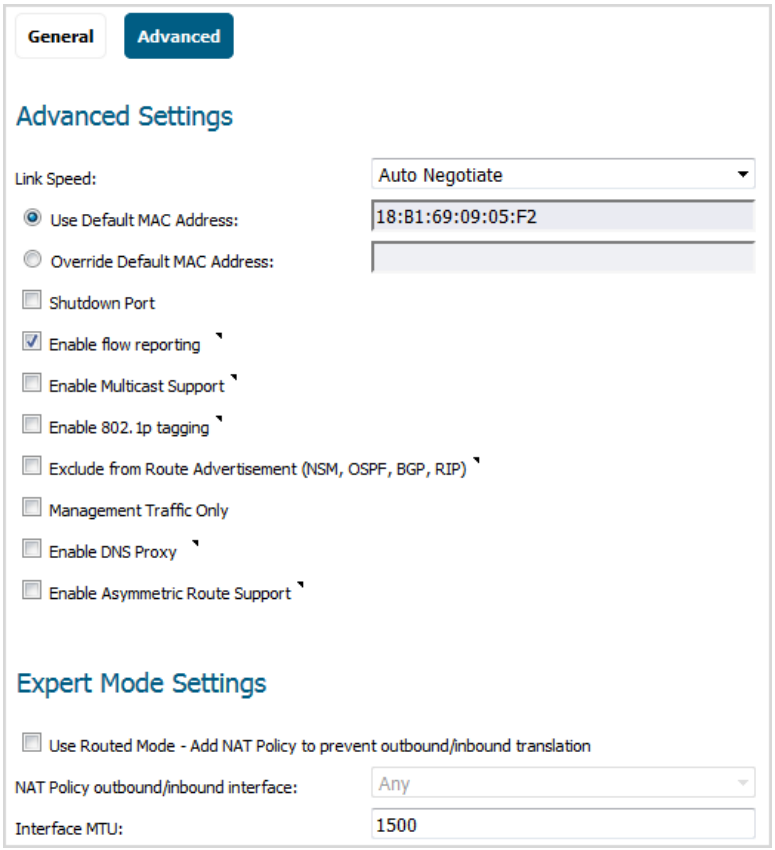
To configure advanced settings for a static interface.

1 In the **Edit Interface** dialog, click **Advanced**.

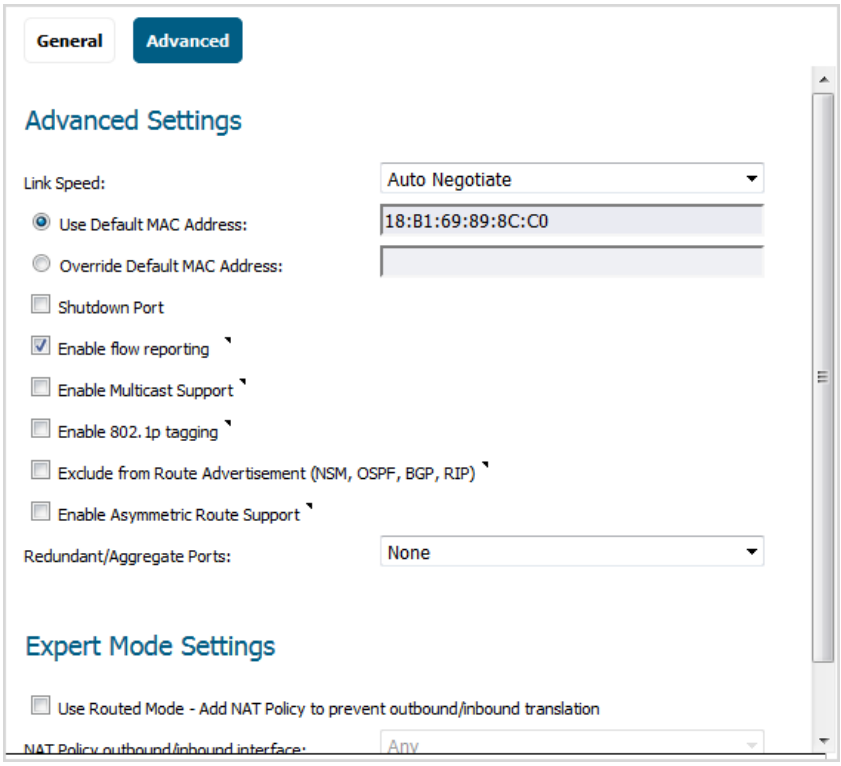
NOTE: The options available in **Advanced** for a static interface vary depending on the selected zone and the platform:

- Edit Interface Advanced settings—LAN/DMZ/WLAN: TZ Series and SOHO W Security Appliance on page 275
- Edit Interface Advanced settings—LAN/DMZ/WLAN: NSA 2600 and above Security Appliance on page 276
- Edit Interface Advanced settings—WAN: TZ Series and SOHO W Security Appliance on page 276
- Edit Interface Advanced settings—WAN: NSA 2600 and above Security Appliance on page 277

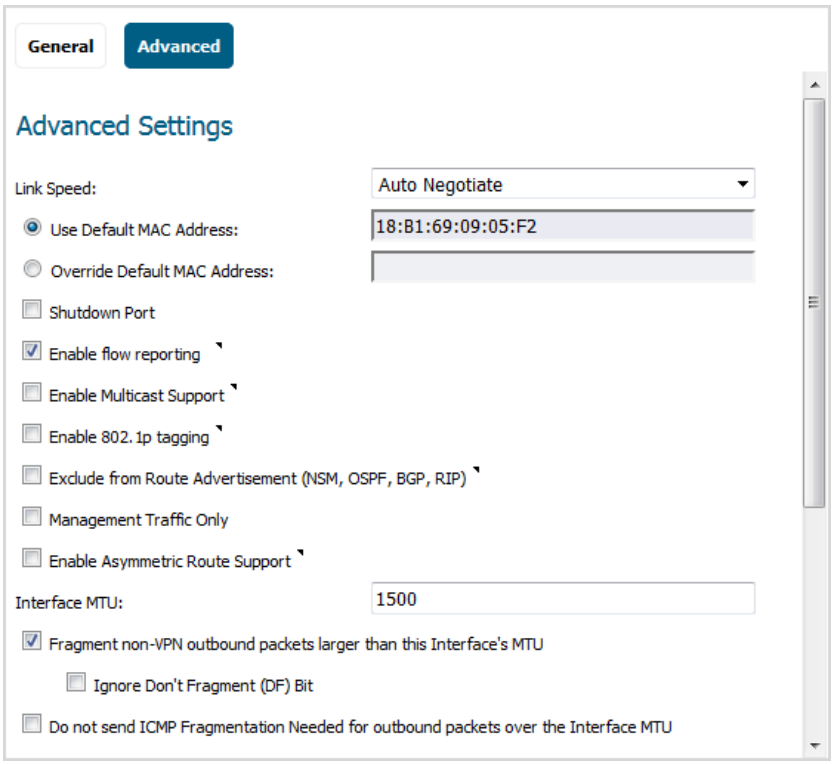
Edit Interface Advanced settings—LAN/DMZ/WLAN: TZ Series and SOHO W Security Appliance



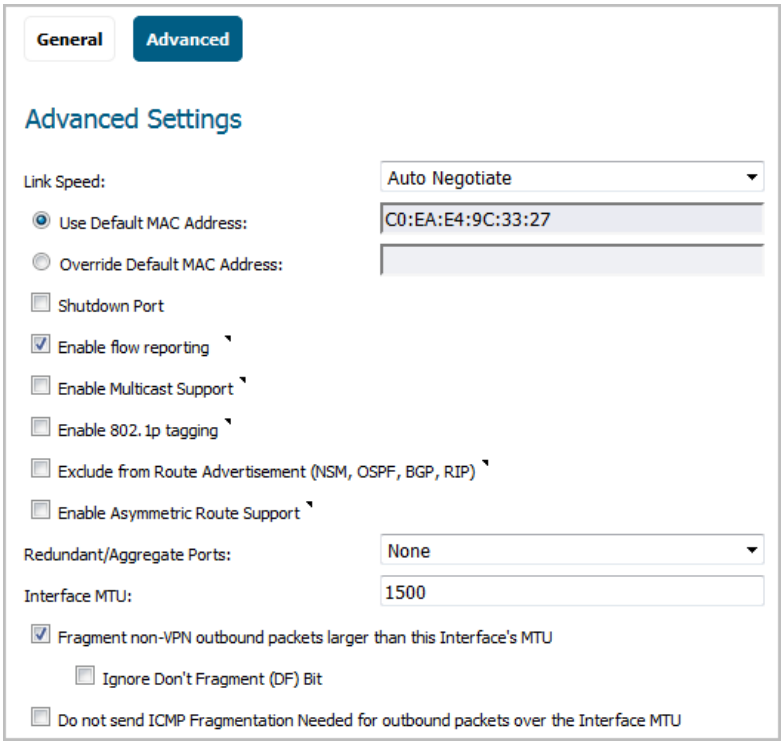
Edit Interface Advanced settings—LAN/DMZ/WLAN: NSA 2600 and above Security Appliance



Edit Interface Advanced settings—WAN: TZ Series and SOHO W Security Appliance



Edit Interface Advanced settings—WAN: NSA 2600 and above Security Appliance



2 For **Link Speed**, **Auto Negotiate** is selected by default, which causes the connected devices to negotiate the speed and duplex mode of the Ethernet connection automatically. To force Ethernet speed and duplex, select one of the following options from **Link Speed**:

For 1 Gbps interfaces	For 10 Gbps interfaces
1 Gbps - Full Duplex	10 Gbps - Full Duplex
100 Mbps - Full Duplex	
100 Mbps - Half Duplex	
10 Mbps - Full Duplex	
10 Mbps - Half Duplex	

CAUTION: If you select a specific Ethernet speed and duplex, you must force the connection speed and duplex from the Ethernet card to the Security Appliance as well.

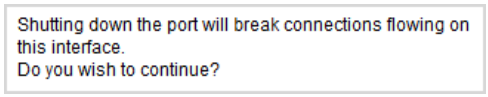
3 **Use Default MAC Address** is selected by default. You override **Use Default MAC Address** for the Interface by choosing **Override Default MAC Address** and entering the MAC address in the field.

4 Select **Shutdown Port** to temporarily take this interface offline for maintenance or other reasons. If connected, the link goes down. This option is not selected by default.

Clear the option to activate the interface and allow the link to come back up.

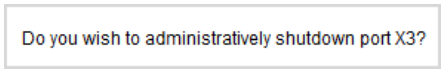
IMPORTANT: You cannot shut down the management interface or the interface you are currently using.

If you select this option, a confirmation message displays:



Click **OK** to shut down the port.

TIP: You can shut down the interface by clicking the **Enabled** icon in the **Enabled** column for the interface. A confirmation message displays:



If you click **OK**, the **Enabled** icon turns to a **Disabled** icon. To enable the interface, click the **Disabled** icon. A confirmation message displays:



If you click **OK**, the **Disabled** icon turns to an **Enabled** icon.

5 For the AppFlow feature, select **Enable flow reporting** to allow flow reporting on flows created for this interface. This option is selected by default.

6 Optionally, select **Enable Multicast Support** to allow multicast reception on this interface. This option is not selected by default.

7 Optionally, select **Enable Default 802.1p CoS** to tag information passing through this interface with 802.1p priority information for Quality of Service (QoS) management. This option is not selected by default.

NOTE: This option is available only for VLAN interfaces.

Packets sent through this interface are tagged with VLAN id=0 and carry 802.1p priority information. To make use of this priority information, devices connected to this interface should support priority frames. QoS management is controlled by access rules on **MANAGE | Policies | Rules > Access Rules**. For information on QoS and bandwidth management, see *SonicOS Policies*.

8 Optionally, to exclude the interface from Route Advertisement, select **Exclude from Route Advertisement (NSM, OSPF, BGP, RIP)** This option is not selected by default.

9 If configuring NSA 2600 and above Security Appliances, go to Step 11.

10 Optionally, select **Management Traffic Only** to restrict traffic to only SonicWall management traffic and routing protocols. This option is not selected by default.

NOTE: Only TZ series and SOHO W appliances have this option.

11 Optionally, if you have enabled DNS Proxy, the **Enable DNS Proxy** option for displays for LAN, DMZ, or WLAN interfaces. To enable DNS Proxy on the interface, select the option. This option is not selected by default.

12 Optionally, enable Asymmetric Route Support on the interface by selecting **Enable Asymmetric Route Support**. If enabled, the traffic initialized from this interface supports asymmetric routes, that is, the initial packet or response packet can pass through from other interfaces. This option is not selected by default. For more information about asymmetric routing, see Asymmetric Routing In Cluster Configurations on page 699.

13 If configuring a TZ Series or SOHO W Security Appliance for a:

- LAN/DMZ/WLAN interface, go to Configuring Routed Mode on page 279.
- WAN interface, go to Step 16.

14 Optionally, select **Link Aggregation** or **Port Redundancy** from **Redundant /Aggregate Ports**. For more information see Configuring Link Aggregation and Port Redundancy on page 301.

NOTE: This option is available only on NSA 2600 and higher appliances.

15 If configuring a LAN/DMZ/WLAN interface for an NSA 2600 and higher Security Appliance, go to Configuring Routed Mode on page 279.

16 To specify the largest packet size (MTU – maximum transmission unit) that a WAN interface can forward without fragmenting the packet, enter the size of the packets that the port receives and transmits in the **Interface MTU** field:

Standard packets (default)	1500
Jumbo frame packets	9000

NOTE: Jumbo frame support must be enabled before a port can process jumbo frames, as explained in *SonicOS 6.5 Policies*. Due to jumbo frame packet buffer size requirements, jumbo frames increase memory requirements by a factor of 4.

Jumbo frames are supported on NSA 3600 and higher appliances.

17 Optionally, to fragment non-VPN outbound packets larger than the interface's MTU, select **Fragment non-VPN outbound packets larger than this Interface's MTU**. This option is selected by default. When selected, the following option becomes available.

IMPORTANT: You specify fragmentation of outbound VPN traffic in **MANAGE | Connectivity | Advanced Settings**. For more information, see *SonicOS Connectivity*.

a Optionally, to override the Do-not-fragment packet bit, select **Ignore Don't Fragment (DF) bit**. This option is not selected by default.

18 To block notification that the WAN interface can receive fragmented packets, select **Do not send ICMP Fragmentation Needed for outbound packets over the Interface MTU**. This option is not selected by default.

19 If configuring bandwidth management for this interface, go to Enabling Bandwidth Management on an Interface on page 281.

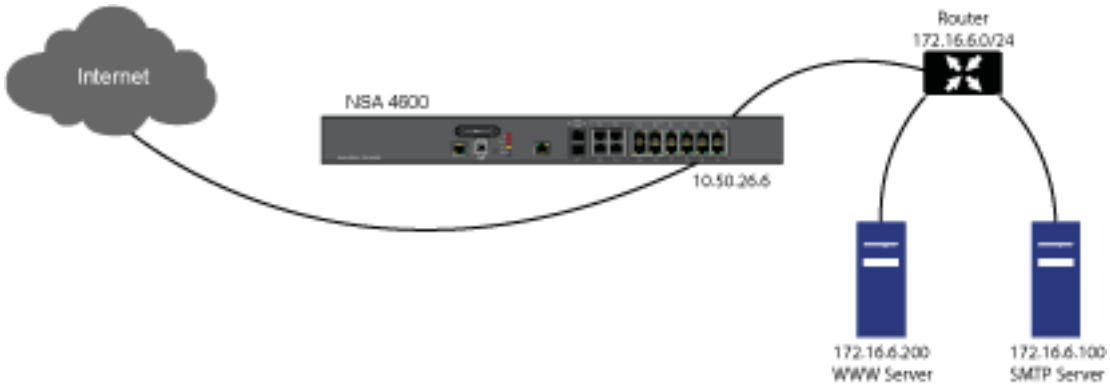
20 Click **OK**.

Configuring Routed Mode

Routed Mode provides an alternative for NAT for routing traffic between separate public IP address ranges. Consider the topology in Routed Mode Configuration, where the Security Appliance is routing traffic across two public IP address ranges:

- 10.50.26.0/24
- 172.16.6.0/24

Routed Mode Configuration



By enabling Routed Mode on the interface for the `172.16.6.0` network, NAT translations are automatically disabled for the interface, and all inbound and outbound traffic is routed to the WAN interface configured for the `10.50.26.0` network.

NOTE: Routed Mode is available when using Static IP Mode for interfaces in the LAN, DMZ, and WLAN zones. For DMZ, it is also available when using Layer 2 Bridged Mode. Routed mode is not available for WAN mode.

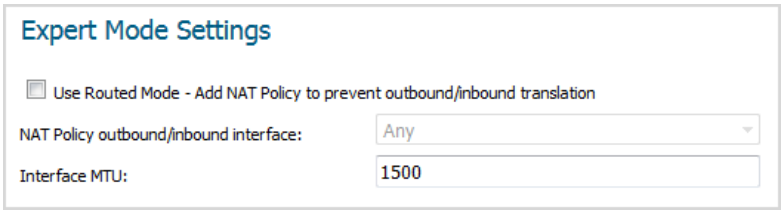
To configure Routed Mode:

1 Navigate to **MANAGE | System Setup** > **Network > Interfaces**.

2 Click on the **Configure** icon for the appropriate interface. The **Edit Interface** dialog displays.

3 Click **Advanced**.

4 Scroll to the **Expert Mode Settings** section.



5 To enable Routed Mode for the interface, select **Use Routed Mode - Add NAT Policy to prevent outbound\inbound translation**. This option is not selected by default. When you select it, the next Expert Mode setting become available.

6 From **NAT Policy outbound/inbound interface**, select the WAN interface that is to be used to route traffic for the interface. The default is **Any**.

7 To specify the largest packet size (MTU – maximum transmission unit) that the interface can forward without fragmenting the packet, enter the size of the packets that the port receives and transmits in the **Interface MTU** field:

Standard packets (default)	1500
Jumbo frame packets	9000

NOTE: Jumbo frame support must be enabled before a port can process jumbo frames, as explained in *SonicOS 6.5 Policies*. Due to jumbo frame packet buffer size requirements, jumbo frames increase memory requirements by a factor of 4.

Jumbo frames are supported on NSA 3600 and higher appliances.

8 If Bandwidth Management has been enabled on the Security Appliance, the Bandwidth Management section displays. To configure BWM for this interface, go to Enabling Bandwidth Management on an Interface on page 281.

9 Click **OK**.

IMPORTANT: The Security Appliance creates “no-NAT” policies for both the configured interface and the selected WAN interface. These policies override any more general M21 NAT policies that may be configured for the interfaces.

Enabling Bandwidth Management on an Interface

Bandwidth Management (BWM) allows you to guarantee minimum bandwidth and prioritize traffic. BWM is enabled in **MANAGE | Security Configuration | Firewall Settings > Bandwidth Management**. By controlling the amount of bandwidth to an application or user, you can prevent a small number of applications or users from consuming all available bandwidth. Balancing the bandwidth allocated to different network traffic and then assigning priorities to traffic improves network performance.

Various types of bandwidth management can be enabled:

- **Advanced**—Enables you to configure maximum egress and ingress bandwidth limitations per interface, by configuring bandwidth objects, access rules, and application policies.
- **Global**—Allows you to enable BWM settings globally and apply them to any interfaces.
- **None** (default)—Disables BWM.

For more information about configuring bandwidth management and the effect of the various BWM types, see the *SonicOS 6.5 Security Configuration* administration documentation, available at https://www.sonicwall.com/support/technical-documentation.

SonicOS can apply bandwidth management to both egress (outbound) and ingress (inbound) traffic on any interfaces. Outbound bandwidth management is done using Class-based Queuing. Inbound Bandwidth Management is done by implementing an ACK delay algorithm that uses TCP’s intrinsic behavior to control the traffic.

Class-based Queuing (CBQ) provides guaranteed and maximum bandwidth Quality of Service (QoS) for the firewall. Every packet destined to the interface is queued in the corresponding priority queue. The scheduler then dequeues the packets and transmits them on the link depending on the guaranteed bandwidth for the flow and the available link bandwidth.

Enabling Bandwidth Management on an Interface

The options described in this section are only available if Bandwidth Management is enabled in **MANAGE | Security Configuration | Firewall Settings > Bandwidth Management**.

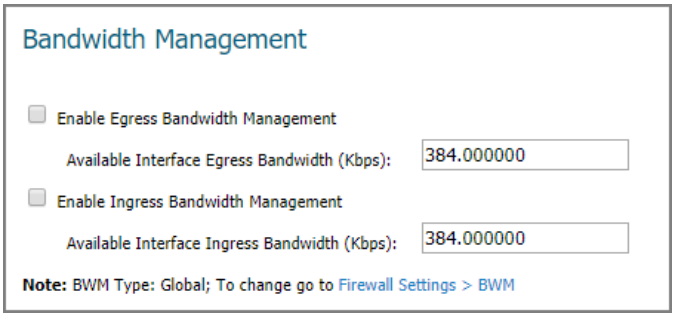
To enable or disable ingress and egress BWM:

1 Navigate to **MANAGE | System Setup > Network > Interfaces**.

2 Click the **Edit** icon of an interface. The **Edit Interface** dialog displays.

3 If this is an unassigned interface, configure the interface according to the sections contained in Configuring Interfaces on page 272.

4 On the **Advanced** screen, scroll to **Bandwidth Management**.



NOTE: Advanced Settings may differ, depending on the Security Appliance model and the type of zone selected.

5 Enable Bandwidth Management for this interface. For more information about Bandwidth Management, see *SonicOS 6.5 Security Configuration*.

a To limit outgoing traffic to a maximum bandwidth on the interface, select **Enable Egress Bandwidth Management**. This option is not selected by default.

- Specify the maximum bandwidth, in kbps, in the **Available Interface Egress Bandwidth** field. The minimum is 20 Kbps, the maximum is 1000000, and the default is **384.000000**.

b To limit incoming traffic to a maximum bandwidth on the interface, select **Enable Ingress Bandwidth Management**. This option is not selected by default.

- Specify the maximum bandwidth, in kbps, in the **Available Interface Ingress Bandwidth** field. The minimum is 20 Kbps, the maximum is 1000000, and the default is **384.000000**.

When either of these options is:

- Selected, the maximum available egress BWM is defined, but as advanced BWM is policy-based, the limitation is not enforced unless there is a corresponding Access Rule or App Rule.
- Not selected, no bandwidth limitation is set at the interface level, but traffic can still be shaped using other options.

6 Optionally, select **Enable Default 802.1p tagging** to tag information passing through this interface with 802.1p priority information for Quality of Service (QoS) management. This option is not selected by default.

Packets sent through this interface are tagged with VLAN id=0 and carry 802.1p priority information. To make use of this priority information, devices connected to this interface should support priority frames. QoS management is controlled by access rules on **MANAGE | Policies > Rules > Access Rules**. For information on QoS and bandwidth management, see the *SonicOS 6.5 Policies* administration documentation.

7 Click **OK**.

Configuring Interfaces in Transparent IP Mode (Splice L3 Subnet)

Transparent IP Mode enables the SonicWall Security Appliance to bridge the WAN subnet onto an internal interface.

To configure an interface for transparent mode:

1 Navigate to **MANAGE | System Setup > Network > Interfaces**.

2 Click on the **Configure** icon for the **Unassigned** Interface you want to configure. The **Edit Interface** dialog is displayed.

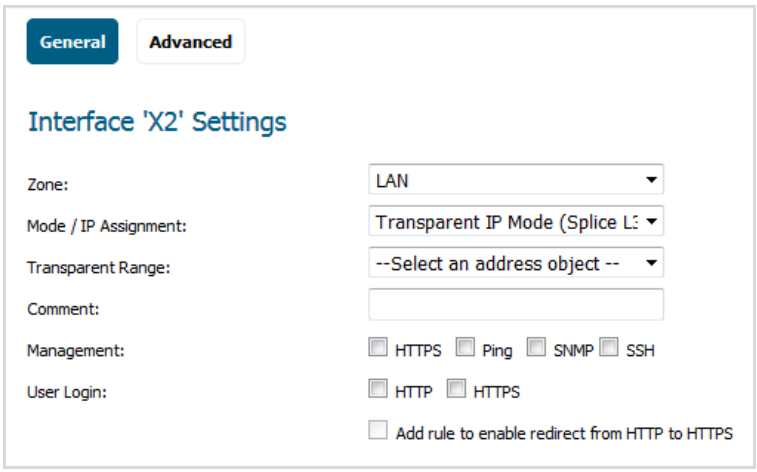
3 Either select

- **LAN** or **DMZ** for **Zone**.

  NOTE: The options available change according to the type of zone you select.

- Create a new zone for the configurable interface by selecting **Create a new zone.** The **Add Zone** dialog displays. See About Zones on page 390 for instructions on adding a zone.

4 Select **Transparent IP Mode (Splice L3 Subnet)** from **Mode / IP Assignment**. The options change.



5 From **Transparent Range**, select an address object that contains the range of IP addresses you want to have access through this interface. The address range must be within an internal zone, such as LAN, DMZ, or another trusted zone matching the zone used for the internal transparent interface.

If you do not have an address object configured that meets your needs, select **Create New Address Object**. The **Add Address Object** dialog displays. For information about creating an address object, see *SonicOS 6.5 Policies*.

6 Enter any optional comment text in the **Comment** field. This text is displayed in the **Comment** column of the **Interface** table. This option is not selected by default.

7 To enable remote management of the Security Appliance from this interface, choose the supported management protocol(s): **HTTPS**, **Ping**, **SNMP**, and/or **SSH**. This option is not selected by default.

To allow access to the WAN interface for management from another zone on the same Security Appliance, access rules must be created. For how to allow WAN primary IP access from a LAN zone, see *SonicOS 6.5 Policies*.

8 To allow selected users with limited management rights to log directly into the Security Appliance through this interface, choose **HTTP** and/or **HTTPS** in **User Login**.

9 If you selected **HTTPS** for either **Management** and/or **User Login** protocol, the **Add rule to enable redirect from HTTP to HTTPS** becomes available and selected. To prevent redirection of HTTP to HTTPS, deselect the option.

TIP: Selecting **HTTP** for **User Login** protocol disables redirection.

NOTE: Elements of this feature can be controlled by internal User Authentication Settings options. For more information, see HTTP/HTTPS Redirection with DP Offload on page 267.

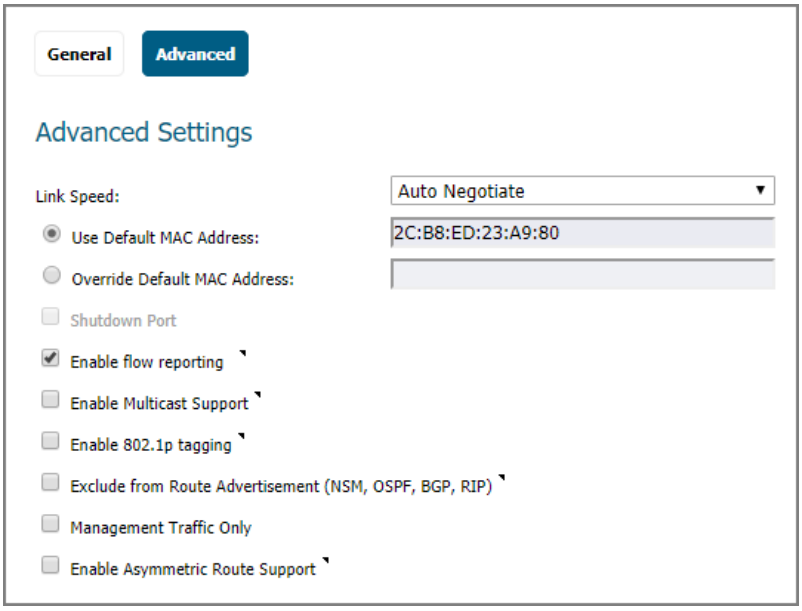
10 Click **OK**.

NOTE: The administrator password is required to regenerate encryption keys after changing the Security Appliance's address.

Configuring Advanced Settings for a Transparent IP Mode Interface

To configure advanced settings for a transparent IP mode interface:

1 In the **Edit Interface** dialog, click **Advanced**.



2 For **Link Speed**, **Auto Negotiate** is selected by default, which causes the connected devices to negotiate the speed and duplex mode of the Ethernet connection automatically. To force Ethernet speed and duplex, select one of the following options from **Link Speed**:

For 1 Gbps interfaces	For 10 Gbps interfaces
1 Gbps - Full Duplex	10 Gbps - Full Duplex
100 Mbps - Full Duplex	
100 Mbps - Half Duplex	
10 Mbps - Full Duplex	
10 Mbps - Half Duplex	

CAUTION: If you select a specific Ethernet speed and duplex, you must force the connection speed and duplex from the Ethernet card to the Security Appliance as well.

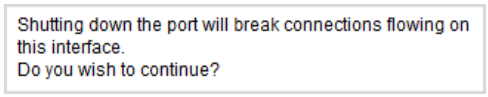
3 **Use Default MAC Address** is selected by default. You override **Use Default MAC Address** for the Interface by choosing **Override Default MAC Address** and entering the MAC address in the field.

4 Select **Shutdown Port** to temporarily take this interface offline for maintenance or other reasons. If connected, the link goes down. This option is not selected by default.

Clear the option to activate the interface and allow the link to come back up. This option is not selected by default.

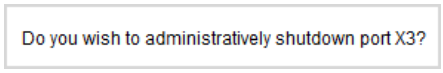
NOTE: You cannot shut down the management interface or the interface you are currently using.

If you select this option, a confirmation message displays:



Click **OK** to shut down the port.

TIP: You can shut down the interface by clicking the **Enabled** icon in the **Enabled** column for the interface. A confirmation message displays:



If you click **OK**, the **Enabled** icon turns to a **Disabled** icon. To enable the interface, click the **Disabled** icon. A confirmation message displays:



If you click **OK**, the **Disabled** icon turns to an **Enabled** icon.

5 For the AppFlow feature, select **Enable flow reporting** to allow flow reporting on flows created for this interface. This option is selected by default.

6 Optionally, select **Enable Multicast Support** to allow multicast reception on this interface. This option is not selected by default.

7 Optionally, select **Enable Default 802.1p CoS** to tag information passing through this interface with 802.1p priority information for Quality of Service (QoS) management. This option is not selected by default.

NOTE: This option is available only for VLAN interfaces.

Packets sent through this interface are tagged with VLAN id=0 and carry 802.1p priority information. To make use of this priority information, devices connected to this interface should support priority frames. QoS management is controlled by access rules on **MANAGE | Policies > Rules > Access Rules**. For information on QoS and bandwidth management, see the *SonicWall SonicOS 6.5 Policies Administration Guide* located on the Support portal at https://www.sonicwall.com/support/technical-documentation/ and choose NSa Series, NSA Series, SuperMassive 9000 Series, and TZ Series in the **Select A Product** field.

8 Optionally, to exclude the interface from Route Advertisement, select **Exclude from Route Advertisement (NSM, OSPF, BGP, RIP)** This option is not selected by default.

9 Optionally, select **Management Traffic Only** to restrict traffic to only SonicWall management traffic and routing protocols. This option is not selected by default.

NOTE: Only TZ series and SOHO W appliances have this option.

10 Optionally, if you have enabled DNS Proxy, the **Enable DNS Proxy** option for displays. To enable DNS Proxy on the interface, select the option. This option is not selected by default.

11 Optionally, enable Asymmetric Route Support on the interface by selecting **Enable Asymmetric Route Support**. If enabled, the traffic initialized from this interface supports asymmetric routes, that is, the initial packet or response packet can pass through from other interfaces. This option is not selected by default. For more information about asymmetric routing, see Asymmetric Routing In Cluster Configurations on page 699.

12 If configuring TZ series and SOHO series Security Appliances, go to Step 14.

13 Optionally, select **Link Aggregation** or **Port Redundancy** from **Redundant /Aggregate Ports**. For more information see Configuring Link Aggregation and Port Redundancy on page 301.

NOTE: This option is available only on NSA 2600 and higher appliances.

14 To specify the largest packet size (MTU – maximum transmission unit) that the interface can forward without fragmenting the packet, enter the size of the packets that the port receives and transmits in the **Interface MTU** field:

Standard packets (default)	1500
Jumbo frame packets	9000

NOTE: Jumbo frame support must be enabled before a port can process jumbo frames, as explained in the *SonicWall SonicOS 6.5 Policies Administration Guide* located on the Support portal at https://www.sonicwall.com/support/technical-documentation/ and choose NSa Series, NSA Series, SuperMassive 9000 Series, and TZ Series in the **Select A Product** field. Due to jumbo frame packet buffer size requirements, jumbo frames increase memory requirements by a factor of 4.

Jumbo frames are supported on NSA 3600 and higher appliances.

15 If bandwidth management has been enabled, to configure BWM for this interface, go to Enabling Bandwidth Management on an Interface on page 281.

16 Click **OK**.

Configuring Wireless Interfaces

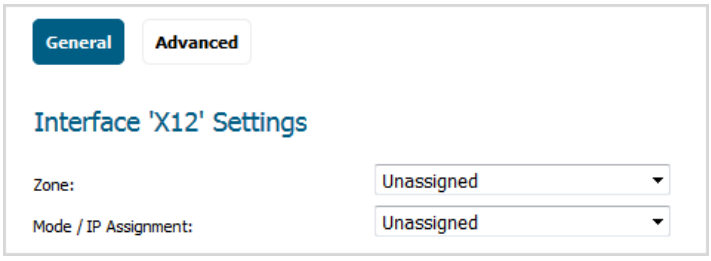
A wireless interface is an interface that has been assigned to a Wireless zone and is used to support SonicWall SonicPoint and SonicWave secure access points.

NOTE: SonicPoints can only be provisioned and managed on the interfaces of security-type wireless (WLAN by default).

To configure wireless interfaces:

1 Navigate to **MANAGE | System Setup > Network > Interfaces**.

2 Click the **Edit** icon in the **Configure** column for the interface you want to configure. The **Edit Interface** dialog displays.

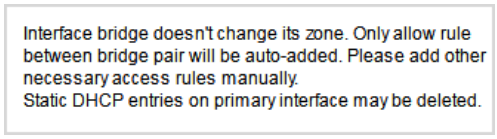


3 From **Zone**, select **WLAN** or a previously defined custom Wireless zone.

4 For **Mode / IP Assignment**, select either:

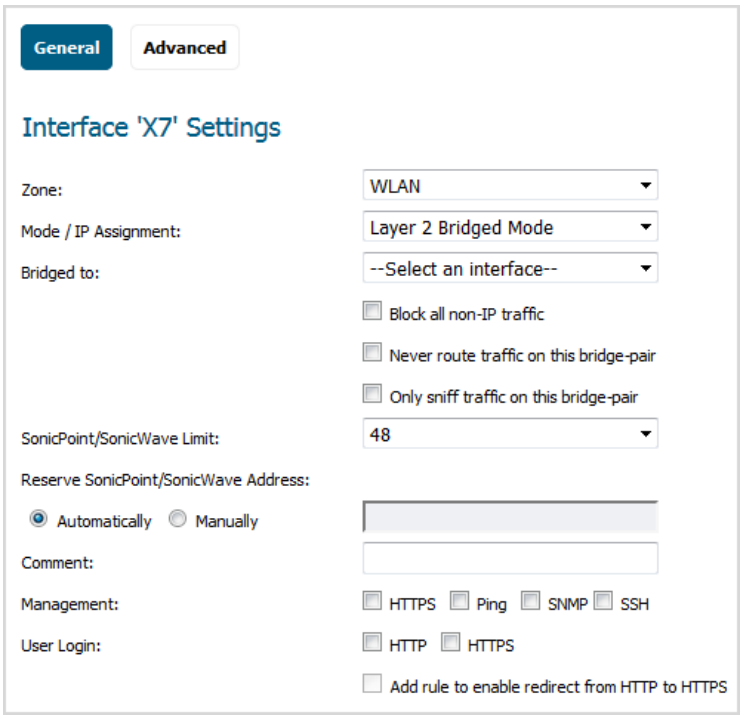
- **Static IP Mode** (default); go to Step 12.

- **Layer 2 Bridged Mode**; a message displays:



IMPORTANT: Choosing this mode requires configuring access rules for the bridged pair. for more information, see the SonicWall SonicOS 6.5 Policies Administration Guide located on the Support portal at https://www.sonicwall.com/support/technical-documentation/ and choose NSa Series, NSA Series, SuperMassive 9000 Series, and TZ Series in the **Select A Product** field. For more information about Layer 2 bridge mode, see Layer 2 Bridged Mode on page 315.

5 Click **OK**. The options change.



6 Select an interface to bridge to from **Bridged to**. Only those interfaces to which this interface can be bridged are displayed.

7 To block all non-IP traffic, select **Block all non-IP traffic**.

8 To never route traffic on the bridged pair, select **Never route traffic on this bridge-pair**.

9 For TZ Series or SOHO W Security Appliance, go to Step 16.

10 To sniff only traffic on the bridged pair, select **Only sniff traffic on this bridge-pair**.

11 Go to Step 14.

12 Enter the IP address of the zone in the **IP Address** field.

13 Enter the subnet mask of the zone in the **Subnet Mask** field.

IMPORTANT: The upper limit of the subnet mask is determined by the number of SonicPoints you select from **SonicPoint/SonicWave Limit**. If you are configuring several interfaces or subinterfaces as Wireless interfaces, you may want to use a smaller subnet (higher) to limit the number of potential DHCP leases available on the interface. Otherwise, if you use a class C subnet (a subnet mask of `255.255.255.0`) for each Wireless interface, you may exceed the limit of DHCP leases available on the Security Appliance.

14 From **SonicPoint/SonicWave Limit**, select the maximum number of SonicPoints allowed on this interface. This value determines the highest subnet mask you can enter in the **Subnet Mask** field. Maximum subnet mask sizes allowed shows the subnet mask limit for each **SonicPoint/SonicWave Limit** selection and the number of DHCP leases available on the interface if you enter the maximum allowed subnet mask.

Available Client IPs assumes 1 IP for the firewall gateway interface, in addition to the presence of the maximum number of SonicPoints allowed on this interface, each consuming an IP address.

Maximum subnet mask sizes allowed

SonicPoints/SonicWaves per Interface	Maximum Subnet Mask	Total Usable IP addresses	Available Client IP address
None	30 bits – 255.255.255.252	2	2
2	29 bits – 255.255.255.248	6	3
4	29 bits – 255.255.255.248	6	1
8	28 bits – 255.255.255.240	14	5
16	27 bits – 255.255.255.224	30	13
24	26 bits – 255.255.255.192	62	29
32	26 bits – 255.255.255.192	62	29
48	25 bits - 255.255.255.128	126	77
64	25 bits - 255.255.255.128	126	61
96	24 bits - 255.255.255.0	190	93
128	23 bits - 255.255.254.0	254	125

TIP: Maximum subnet mask sizes allowed depicts the maximum subnet mask sizes allowed. You can still use classful subnetting (class A, class B, or class C) or any variable-length subnet mask that you wish on WLAN interfaces. You are encouraged to use a smaller subnet mask (for example, 24-bit class C: 255.255.255.0 - 254 total usable IPs), thus allocating more IP addressing space to clients if you have the need to support larger numbers of wireless clients.

15 From SonicPoint/SonicWave Address, choose how SonicOS determines the reserve address of the SonicPoint/SonicWave appliance:

- Automatically – This option is selected by default. Go to Step 16.
- Manually – When selected the address field becomes available. Enter the reserve address.

16 Enter any optional comment text in the **Comment** field. This text is displayed in the **Comment** column of the **Interface** table.

17 If you want to enable remote management of the firewall from this interface, select the supported **Management** protocol(s): **HTTPS**, **Ping**, **SNMP**, and/or **SSH**.

18 If you want to allow selected users with limited management rights to log in to the Security Appliance, select **HTTP** and/or **HTTPS** in **User Login**.

19 If you selected **HTTPS** for either **Management** or **User Login** protocol, the **Add rule to enable redirect from HTTP to HTTPS** becomes available and selected. Selecting **HTTP** for **User Login** deselects the option even if **HTTPS** is also selected.

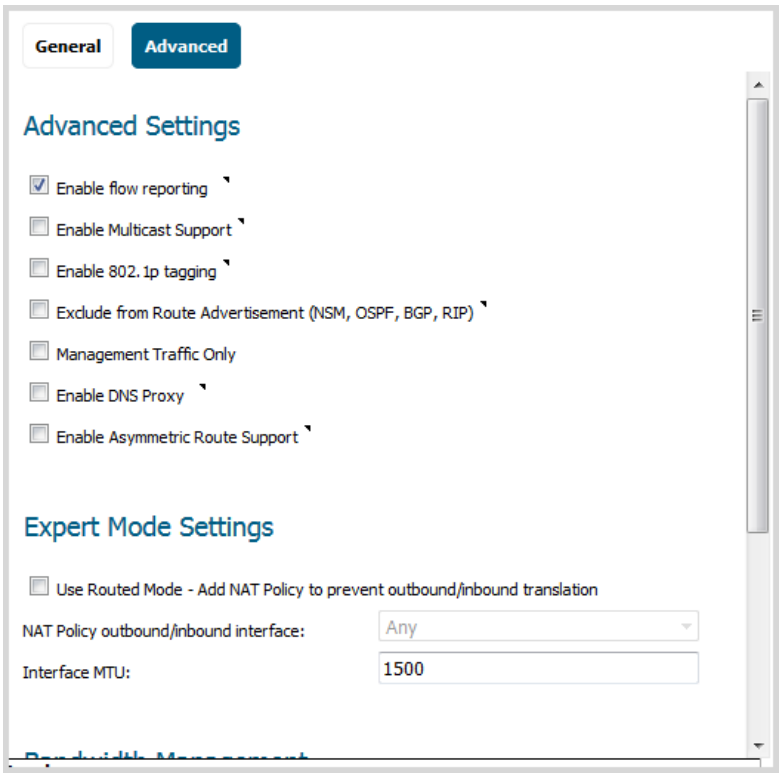
20 Click **OK**.

Configuring Advanced Settings for a Wireless Interface

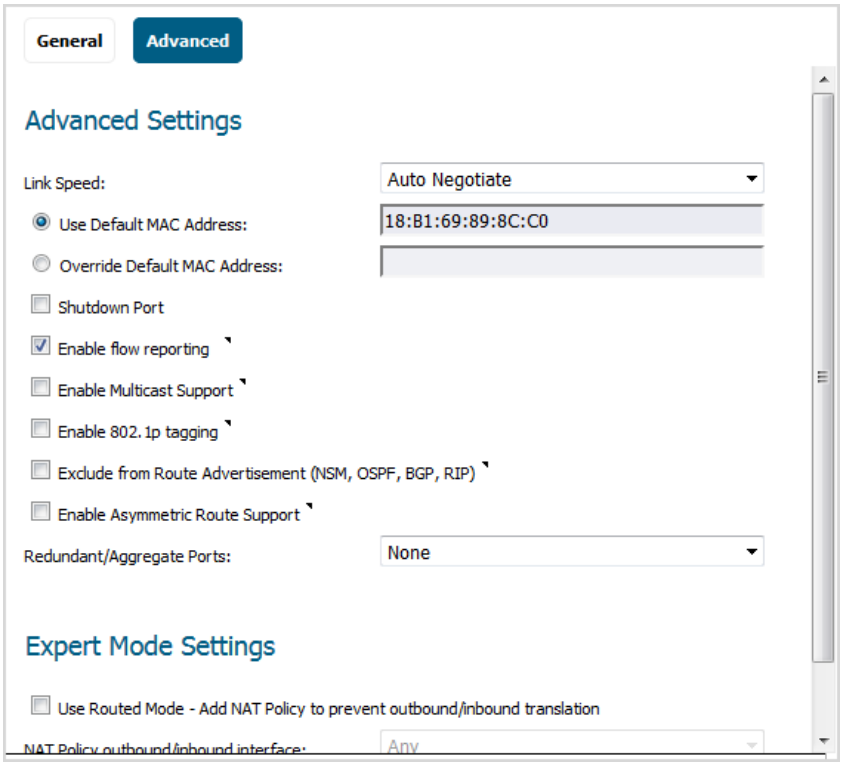
To configure advanced settings for a wireless interface:

1 In the **Edit Interface** dialog, click the **Advanced** tab. The options depend on the platform of the Security Appliance.

Advanced tab: TZ Series and SOHO W Security Appliance



Advanced tab: NSA and above Security Appliances



2 For configuring TZ Series and SOHO W appliances, go to Step 6.

3 For **Link Speed**, **Auto Negotiate** is selected by default, which causes the connected devices to negotiate the speed and duplex mode of the Ethernet connection automatically. To force Ethernet speed and duplex, select one of the following options from **Link Speed**:

For 1 Gbps interfaces	For 10 Gbps interfaces
1 Gbps - Full Duplex	10 Gbps - Full Duplex
100 Mbps - Full Duplex	
100 Mbps - Half Duplex	
10 Mbps - Full Duplex	
10 Mbps - Half Duplex	

CAUTION: If you select a specific Ethernet speed and duplex, you must force the connection speed and duplex from the Ethernet card to the Security Appliance as well.

4 For the default MAC address, choose either:

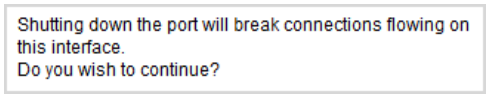
- **Use Default MAC Address** – The address is chosen automatically; the address is displayed in the dimmed field and cannot be changed. This option is selected by default.
- **Override Default MAC Address** – To specify a different default MAC address for the Interface; the address field becomes available. Enter the MAC address in the field.

5 Select **Shutdown Port** to temporarily take this interface offline for maintenance or other reasons. If connected, the link goes down. This option is not selected by default.

Clear the option to activate the interface and allow the link to come back up. This option is not selected by default.

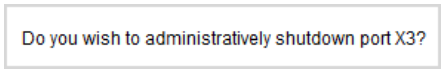
NOTE: You cannot shut down the management interface or the interface you are currently using.

If you select this option, a confirmation message displays:



Click **OK** to shut down the port.

TIP: You can shut down the interface by clicking the **Enabled** icon in the **Enabled** column for the interface. A confirmation message displays:



If you click **OK**, the **Enabled** icon turns to a **Disabled** icon. To enable the interface, click the **Disabled** icon. A confirmation message displays:



If you click **OK**, the **Disabled** icon turns to an **Enabled** icon.

6 For the AppFlow feature, select **Enable flow reporting** to allow flow reporting on flows created for this interface. This option is selected by default.

7 Optionally, select **Enable Multicast Support** to allow multicast reception on this interface. This option is not selected by default.

8 Optionally, select **Enable Default 802.1p CoS** to tag information passing through this interface with 802.1p priority information for Quality of Service (QoS) management. This option is not selected by default.

NOTE: This option is available only for VLAN interfaces.

Packets sent through this interface are tagged with VLAN id=0 and carry 802.1p priority information. To make use of this priority information, devices connected to this interface should support priority frames. QoS management is controlled by access rules on **MANAGE | Policies | Rules > Access Rules**. For information on QoS and bandwidth management, see *SonicOS Policies*.

9 Optionally, to exclude the interface from Route Advertisement, select **Exclude from Route Advertisement (NSM, OSPF, BGP, RIP)** This option is not selected by default.

10 If configuring a SuperMassive or NSA series appliance, go to Step 12.

11 Optionally, select **Management Traffic Only** to restrict traffic to only SonicWall management traffic and routing protocols. This option is not selected by default.

NOTE: Only TZ series and SOHO series Security Appliances have this option.

12 Optionally, if you have enabled DNS Proxy, the **Enable DNS Proxy** option displays. To enable DNS Proxy on the interface, select the option. This option is not selected by default.

13 Optionally, enable Asymmetric Route Support on the interface by selecting **Enable Asymmetric Route Support**. If enabled, the traffic initialized from this interface supports asymmetric routes, that is, the initial packet or response packet can pass through from other interfaces. This option is not selected by default. For more information about asymmetric routing, see Asymmetric Routing In Cluster Configurations on page 699.

14 If configuring TZ series and SOHO series Security Appliances, go to Step 14.

15 Optionally, select **Link Aggregation** or **Port Redundancy** from **Redundant /Aggregate Ports**. For more information see Configuring Link Aggregation and Port Redundancy on page 301.

NOTE: This option is available only on NSA 2600 and higher appliances.

16 Select **Enable Gratuitous ARP Forwarding Towards WAN** to forward gratuitous ARP packets received on this interface towards the WAN, using the hardware MAC address of the WAN interface as the source MAC address.

17 Select **Enable Automatic Gratuitous ARP Generation Towards WAN** to automatically send gratuitous ARP packets towards the WAN whenever a new entry is added to the ARP table for a new machine on this interface. The hardware MAC address of the WAN interface is used as the source MAC address of the ARP packet.

18 To specify the largest packet size (MTU – maximum transmission unit) that the interface can forward without fragmenting the packet, enter the size of the packets that the port receives and transmits in the **Interface MTU** field:

Standard packets (default)	1500
Jumbo frame packets	9000

NOTE: Jumbo frame support must be enabled before a port can process jumbo frames, as explained in *SonicOS Policies*. Due to jumbo frame packet buffer size requirements, jumbo frames increase memory requirements by a factor of 4.

Jumbo frames are supported on NSA 3600 and higher appliances.

19 If configuring routed mode for this interface, go to Configuring Routed Mode on page 279.

20 If bandwidth management has been enabled, to configure BWM for this interface, go to Enabling Bandwidth Management on an Interface on page 281.

21 Click **OK**.

Configuring a WAN Interface

NOTE: A default gateway IP is required on the WAN interface if any destination is required to be reached through the WAN interface that is not part of the WAN subnet IP address space, regardless whether we receive a default route dynamically from a routing protocol of a peer device on the WAN subnet.

Configuring a WAN interface enables Internet connectivity. You can configure up to *N minus 2* WAN interfaces on the SonicWall Security Appliance, where *N* is the number of interfaces defined on the unit (both physical and VLAN). Only the X0 and MGMT interfaces cannot be configured as WAN interfaces.

To configure your WAN interface:

1 Navigate to **MANAGE | System Setup | Network > Interfaces**.

2 Click on the **Edit** icon in the **Configure** column for the Interface you want to configure. The **Edit Interface** dialog displays.

3 If you are configuring an unassigned Interface, select **WAN** from the **Zone** menu. If you selected the **Default WAN** interface, **WAN** is already selected in the **Zone** menu.

4 Select one of the following WAN Network Addressing Modes from **IP Assignment**.

NOTE: Depending on the option you choose from the IP Assignment drop-down menu, the options available change. Complete the corresponding fields that are displayed after selecting the option.

- **Static** - configures the Security Appliance for a network that uses static IP addresses.

- **DHCP** - configures the Security Appliance to request IP settings from a DHCP server on the Internet. NAT with DHCP Client is a typical network addressing mode for cable and DSL customers.
- **PPPoE** - uses Point to Point Protocol over Ethernet (PPPoE) to connect to the Internet. If a username and password is required by your ISP, enter them into the **User Name** and **User Password** fields. This protocol is typically found when using a DSL modem.
- **PPTP** - uses PPTP (Point to Point Tunneling Protocol) to connect to a remote server. It supports older Microsoft Windows implementations requiring tunneling connectivity.
- **L2TP** - uses IPsec to connect a L2TP (Layer 2 Tunneling Protocol) server and encrypts all data transmitted from the client to the server. However, it does not encrypt network traffic to other destinations.
- **Wire Mode (2-Port Wire)** - allows insertion of the Security Appliance into a network, in Bypass, Inspect, or Secure mode. For detailed information, see Configuring Wire and Tap Mode on page 310.
- **Tap Mode (1-Port Tap)** - allows insertion of the Security Appliance into a network for use with network taps, port mirrors, or SPAN ports. For detailed information, see Configuring Wire and Tap Mode on page 310.

5 If using **DHCP**, optionally enter a descriptive name in the **Host Name** field and any desired comments in the **Comment** field.

6 If using **PPPoE**, **PPTP**, or **L2TP**, additional fields display:

- If **Schedule** is displayed, select the desired schedule from the drop-down list during which this interface should be connected.
- In **User Name** and **User Password**, type in the account name and password provided by your ISP.
- If the **Server IP Address** field is displayed, enter the server IP address provided by your ISP.
- If the **(Client) Host Name** field is displayed, enter the host name of the appliance. This is the Firewall Name from **MANAGE | System Setup | Appliance > Base Settings**.
- If the **Shared Secret** field is displayed, enter the value provided by your ISP.

7 If you want to enable remote management of the Security Appliance from this interface, select the supported management protocol(s): **HTTPS**, **Ping**, **SNMP**, and/or **SSH**.

To allow access to the WAN interface for management from another zone on the same Security Appliance, access rules must be created. For information about creating access rules, see *SonicOS Policies*.

8 If using PPPoE, PPTP, or L2TP, additional fields display:

- For PPPoE, choose one of the following:
  - **Obtain IP Address Automatically** to get the IP address from the PPPoE server.
  - **Specify IP Address** and enter the desired IP address into the field to use a static IP address for this interface.
  - **Unnumbered interface** and either:
    - Select an unnumbered interface.
    - Create a new unnumbered interface by selecting **Create new Unnumbered Interface**.

      NOTE: The interface must be unassigned.

- For PPTP or L2TP, configure these options:

- Select **Inactivity Disconnect** and enter the number of minutes of inactivity after which the connection is terminated. Clear this option to disable inactivity timeouts.
- From **IP Assignment**, select either:
  - **DHCP**; the IP Address, Subnet Mask, and Gateway Address fields are automatically provisioned by the server.
  - **Static**, enter the appropriate values for these fields.

9 If using DHCP, optionally choose:

- **Request renew of previous IP on startup** to request the same IP address for the WAN interface that was previously provided by the DHCP server.
- **Renew DHCP lease on any link up occurrence** to send a lease renewal request to the DHCP server every time this WAN interface reconnects after being disconnected.

The fields displayed below these options are provisioned by the DHCP server. After provisioning, these buttons are available; choose:

- **Renew** to restart the DHCP lease duration for the currently assigned IP address.
- **Release** to cancel the DHCP lease for the current IP address. The connection is dropped. You need to obtain a new IP address from the DHCP server to reestablish connectivity.
- **Refresh** to obtain a new IP address from the DHCP server.

10 To allow selected users with limited management rights to log directly into the Security Appliance from this interface, select **HTTP** and/or **HTTPS** in **User Login**.

11 Check **Add rule to enable redirect from HTTP to HTTPS**, if you want an HTTP connection automatically redirected to a secure HTTPS connection to the Security Appliance. For more information about this option, see HTTP/HTTPS Redirection on page 267.

12 Continue the configuration on the **Advanced** and **Protocol** tabs (if displayed) as described in Configuring Advanced Settings for a WAN Interface on page 294.

13 To continue with advanced settings; go to Configuring Advanced Settings for a WAN Interface on page 294.

14 If you selected **PPPoE**, **PPTP**, or **L2TP** for **IP Assignment**, go to Configuring Protocol Settings for a WAN Interface on page 296.

15 Click **OK.**

Configuring Advanced Settings for a WAN Interface

To configure advanced settings for a WAN interface:

1 In the **Edit Interface** dialog, click the **Advanced** tab.

2 For **Link Speed, Auto Negotiate** is selected by default, which causes the connected devices to automatically negotiate the speed and duplex mode of the Ethernet connection. If you want to specify the forced Ethernet speed and duplex, select one of the following options from the **Link Speed** menu:

- For 1 Gbps interfaces, select:
  - **1 Gbps - Full Duplex**
  - **100 Mbps - Full Duplex**
  - **100 Mbps - Half Duplex**
  - **10 Mbps - Full Duplex**

- **10 Mbps - Half Duplex**

- For 10 Gbps interfaces, the only selection is **10 Gbps - Full Duplex**.

IMPORTANT: If you select a specific Ethernet speed and duplex, you must force the connection speed and duplex from the Ethernet card to the firewall as well.

3 You can choose to override the **Use Default MAC Address** for the Interface by selecting **Override Default MAC Address** and entering the MAC address in the field.

4 Select the **Shutdown Port** checkbox to temporarily take this interface off line for maintenance or other reasons. If connected, the link goes down. Clear the checkbox to activate the interface and allow the link to come back up.

5 For the AppFlow feature, select the **Enable flow reporting** checkbox to allow flow reporting on flows created for this interface.

6 Select the **Enable Multicast Support** checkbox to allow multicast reception on this interface.

7 Select the **Enable 802.1p tagging** checkbox to tag information passing through this interface with 802.1p priority information for Quality of Service (QoS) management. Packets sent through this interface are tagged with VLAN id=0 and carry 802.1p priority information. In order to make use of this priority information, devices connected to this interface should support priority frames. QoS management is controlled by access rules on **MANAGE | Security Configuration | Firewall Rules > Quality of Service Mapping**. For information on QoS and bandwidth management, see *SonicOS Security Configuration*.

8 Optionally select **Link Aggregation** or **Port Redundancy** from the **Redundant /Aggregate Ports** drop-down list. For more information see Configuring Link Aggregation and Port Redundancy on page 301.

9 **Interface MTU** - Specifies the largest packet size that the interface can forward without fragmenting the packet. Identify the size of the packets that the port receives and transmits:

Standard packets (default)	1500
Jumbo frame packets	9000

NOTE: Jumbo frame support must be enabled before a port can process jumbo frames. For more information about jumbo frames, see *SonicOS Security Configuration*. Due to jumbo frame packet buffer size requirements, jumbo frames increase memory requirements by a factor of 4.

Jumbo frames are supported by NSA 3600 and higher appliances.

- **Fragment non-VPN outbound packets larger than this Interface's MTU** - Specifies all non-VPN outbound packets larger than this Interface's MTU be fragmented. Specifying the fragmenting of VPN outbound packets is set in **MANAGE | Connectivity | VPN**; for further information about VPN traffic, see *SonicOS Connectivity*.
- **Ignore Don't Fragment (DF) Bit** - Overrides DF bits in packets.
- **Suppress ICMP Fragmentation Needed message generation -** blocks notification that this interface can receive fragmented packets.

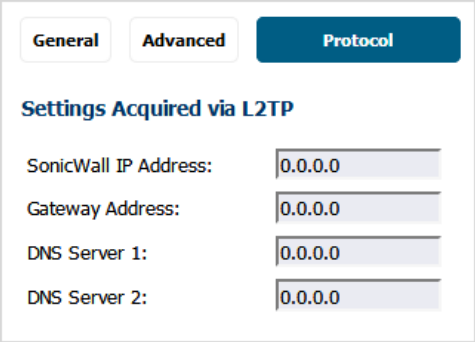
10 If using DHCP, the following options are displayed:

- Select **Initiate renewals with a Discover when using DHCP** if the server might change.
- Select **Use an interval of _ seconds between DHCP Discovers** during lease acquisition and adjust the number of seconds for the interval if the DHCP server might not respond immediately.

11 Optionally enable Bandwidth Management for this interface. For more information about Bandwidth Management, see Enabling Bandwidth Management on an Interface on page 281.

Configuring Protocol Settings for a WAN Interface

If you specified a **PPPoE**, **PPTP**, or **L2TP** for **IP assignment** when configuring the WAN interface, the **Edit Interface** dialog displays the **Protocol** tab.

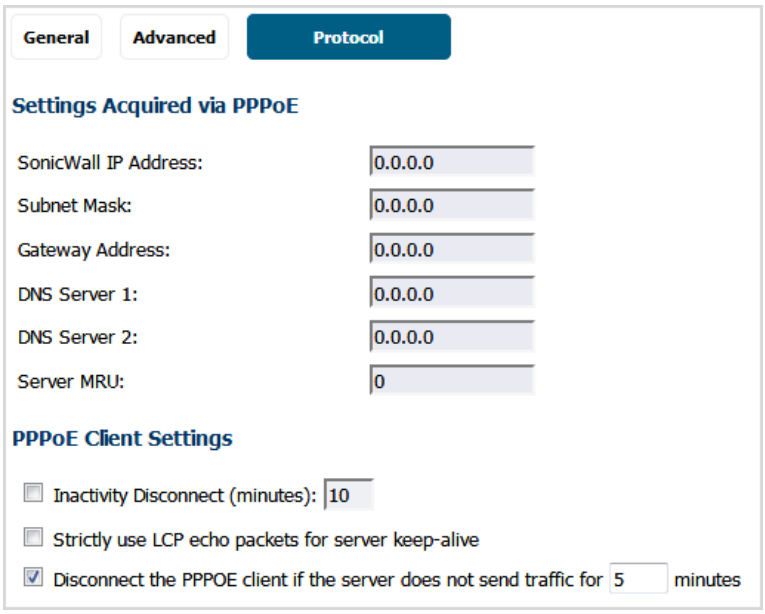


The Internet Service Provider (ISP) provisions the fields (for example, **SonicWall IP Address**, **Subnet Mask**, and **Gateway Address**) in the **Settings Acquired via** section of the **Protocol** tab. These fields show actual values after you connect the Security Appliance to the ISP.

Additionally, specifying PPPoE causes SonicOS to set the Interface MTU option in the **Advanced** tab to **1492** and provides additional settings in the **Protocol** tab.

To configure additional settings for PPPoE:

1 In the **Edit Interface** dialog, click **Protocol**.



2 Enable the following options in the **PPPoE Client Settings** section:

- **Inactivity Disconnect (minutes)**: Enter the number of minutes (the default is 10) after which SonicOS terminates the connection when it detects that packets are not being sent. This option is not selected by default.
- **Strictly use LCP echo packets for server keep-alive**: Select this to have SonicOS terminate the connection if it detects that the PPoE server has not sent a `ppp LCP echo request` packet within a minute. Select this option only if your PPPoE server supports the `send LCP echo` function. This option is not selected by default.

- **Reconnect the PPPOE client if the server does not send traffic for __ minutes**: Enter the number of minutes (the default is 5) after which SonicOS terminates the PPPoE server's connection, and then reconnects, if the server does not send any packets (including the LCP echo request). This option is selected by default.

Configuring Tunnel Interfaces

You can configure several types of tunnel interfaces in SonicOS:

- Numbered and unnumbered tunnel interfaces, WLAN tunnel interfaces, and IPv6 6to4 tunnel interfaces are configured on **Network > Interfaces**.
- Drop tunnel interfaces and VPN tunnel interfaces are configured from **MANAGE | Network > Routing**; for more information, see Configuring Route Advertisements and Route Policies on page 450.
- Unnumbered tunnel interfaces are configured as part of a VPN policy from **MANAGE | Connectivity > VPN**; for information about VPN policies, see *SonicOS Connectivity*.

Numbered and unnumbered tunnel interfaces are used with VPNs. A numbered tunnel interface is assigned its own IP address, but an unnumbered tunnel interface borrows an IP address from an existing physical or virtual (VLAN) interface.

Both numbered and unnumbered tunnel interface types support static routing and dynamic routing with RIP and OSPF, while numbered tunnel interfaces can also be used with BGP.

Also, both numbered VPN and unnumbered tunnel interfaces can support advanced routing, and unnumbered tunnel interfaces have no restrictions. See Maximum VPN policies and tunnels by platform.

Maximum VPN policies and tunnels by platform

Platform	Maximum VPN policies	Maximum VPN Tunnel	Platform	Maximum VPN policies	Maximum VPN Tunnel
SM 9200	10000	192	NSA 2600	75	64
SM 9400	10000	192	NSA 3600	1000	96
SM 9600	10000	192	NSA 4600	3000	96
			NSA 5600	4000	128
TZ 600	50	50	NSA 6600	6000	128
TZ 500	25	25			
TZ 500 W	25	25	NS*a* 2650	75	64
TZ 400	20	20	NS*a* 3650	1000	96
TZ 400 W	20	20	NS*a* 4650	3000	96
TZ 300	10	10	NS*a* 5650	4000	128
TZ 300 W	10	10	NS*a* 6650	6000	128
			NS*a* 9250	10000	192
SOHO W	10	10	NS*a* 9450	10000	192
			NS*a* 9650	10000	192

See these sections for configuring the various types of tunnel interfaces:

- Numbered Tunnel Interfaces; see Configuring VPN Tunnel Interfaces on page 298
- Unnumbered Tunnel Interfaces; see *SonicOS Connectivity*.
- Drop Tunnel Interfaces; see Drop Tunnel Interface on page 463

- IPv6 6to4 Tunnel Interfaces; see Configuring the 6to4 Auto Tunnel on page 949

Configuring VPN Tunnel Interfaces

You can create a numbered tunnel interface by selecting VPN Tunnel Interface from the Add Interface drop-down list. VPN tunnel interfaces are added to the Interface Settings table and then can be used with dynamic routing, including RIP, OSPF, and BGP, or a static route policy can use the VPN tunnel interface as the interface in a configuration for a static route-based VPN.

A VPN Tunnel Interface (TI) can be configured like a standard interface, including options to enable appliance management or user login using HTTP, HTTPS, Ping, or SSH in addition to multicast, flow reporting, asymmetric routing, fragmented packet handling, and Don't Fragment (DF) Bit settings.

NOTE: A similar VPN policy and numbered tunnel interface must be configured on the remote gateway. The IP addresses assigned to the numbered tunnel interfaces (on the local gateway and the remote gateways) must be on the same subnet.

VPN tunnel interface deployment lists how a VPN Tunnel Interface can be deployed.

VPN tunnel interface deployment

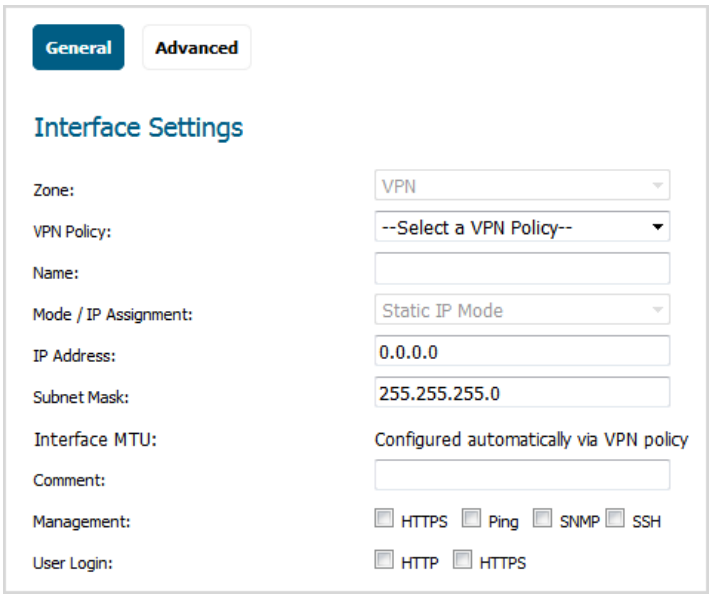
TI can be configured as an interface in	TI cannot be configured as
Static Route	Static ARP entries interface
NAT	HA interface
ACL (Virtual Access Point Access Control List)	WLB (WAN Load Balancing) interface
	Static NDP (Neighbor Discovery Protocol) entries interface
OSPF	OSPFv3/RIPnG: currently not supported for IPv6 advanced routing
RIP	MAC_IP Anti-spoof interface
BGP	DHCP server interface

For all platforms, the maximum supported number of VPN Tunnel Interfaces (numbered tunnel interfaces) is 64. The maximum number of unnumbered tunnel interfaces differs by platform and directly corresponds to the maximum number of VPN policies supported on each platform.

To configure a VPN Tunnel Interface:

1 Navigate to **MANAGE | System Setup > Network > Interfaces**.

2 From **Add Interface** under the **Interface Settings** table, select **VPN Tunnel Interface**. The **Add Tunnel Interface** dialog displays.



The zone is defined as VPN and cannot be changed.

3 From **VPN Policy**, select a VPN policy.

4 In the **Name** field, enter a friendly name for this interface. The name can contain alphanumeric characters, periods (dots), or underscores; it cannot contain spaces or hyphens.

5 Enter an IP address in the **IP Address** field. The default is **0.0.0.0**, but you need to enter an explicit IP address or an error message displays.

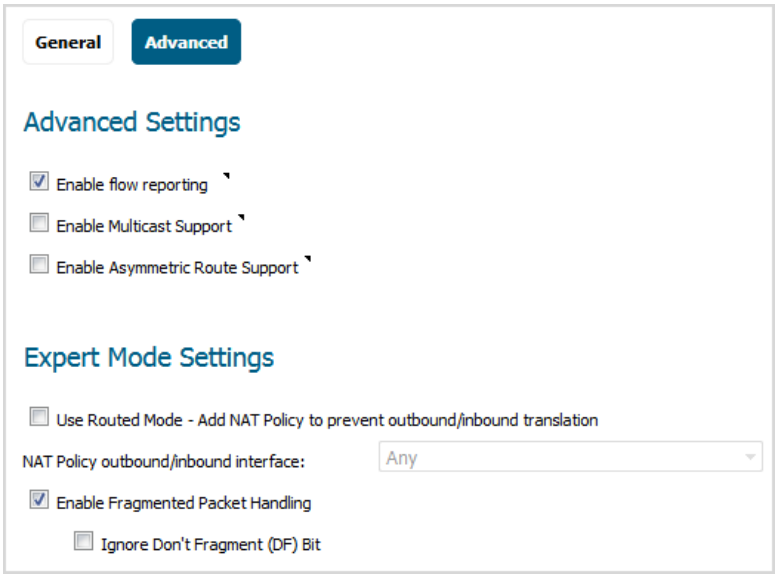
6 In the **Subnet Mask** field, enter the subnet mask. The default is **255.255.255.0**.

7 Optionally, add a comment in the **Comment** field.

8 Optionally, specify the **Management** protocol(s) allowed on this interface: **HTTPS**, **Ping**, **SNMP**, and/or **SSH**.

9 Optionally, specify the **User Login** protocol(s) allowed on this interface: **HTTP** and/or **HTTPS**.

10 Click **Advanced**.



11 To enable flow reporting on flows created for the tunnel interface, select **Enable flow reporting**. This option is selected by default.

12 Optionally, enable multicast reception on the interface by selecting **Enable Multicast Support**. This option is not selected by default.

13 Optionally, enable Asymmetric Route Support on the tunnel interface by selecting **Enable Asymmetric Route Support**. This option is not selected by default. For more information about asymmetric routing, see Asymmetric Routing In Cluster Configurations on page 699.

14 To use Routed Mode and add a NAT policy o prevent outbound/inbound translation, select **User Routed Mode – Add NAT Policy to prevent outbound/inbound translation**. When selected, the following option becomes available. This option is not selected by default.

15 If Routed Mode is selected, to specify an interface for the NAT policy, select an interface from **NAT Policy outbound/inbound interface**. The available interfaces depend on your Security Appliance. The default is **ANY**.

16 To enable fragmented packet handling on this interface, select **Enable Fragmented Packet Handling**. If this option is not selected, fragmented packets are dropped and the VPN log report shows the log message `Fragmented IPsec packet dropped`. This option is selected by default.

If this option is selected, the **Ignore Don't Fragment (DF) Bit** option is available.

17 Select **Ignore Don't Fragment (DF) Bit** to ignore the DF bit in the packet header. Some applications can explicitly set the Don't Fragment option in a packet, which tells all Security Appliances to not fragment the packet. This option, when enabled, causes the Security Appliance to ignore the DF bit and fragment the packet regardless.

18 Click **OK**. The numbered VPN tunnel interface is added to the **Interface Settings** table.

Configuring Link Aggregation and Port Redundancy

NOTE: Link Aggregation and Port Redundancy is supported on NSA 2600 and above Security Appliances.

Both Link Aggregation and Port Redundancy are configured on the **Advanced** tab of the **Edit Interface** dialog in the SonicOS Management Interface.

- Link Aggregation on page 301 - Groups multiple Ethernet interfaces together forming a single logical link to support greater throughput than a single physical interface could support. This provides the ability to send multi-gigabit traffic between two Ethernet domains.

  NOTE: Link Aggregation is supported on NSA 2600 and higher Security Appliances. The NSA 2600 supports Link Aggregation for Network Interfaces, but the NSA 2600 does not support Switching and, therefore, does not support Link Aggregation for Switching, which is covered in **Switching > Link Aggregation** on page 627.

  Link Aggregation is not supported in Layer 2 Bridged Mode.

- Port Redundancy on page 304 - Configures a single redundant port for any physical interface that can be connected to a second switch to prevent a loss of connectivity in the event that either the primary interface or primary switch fail.

  NOTE: Port Redundancy is supported on NSA 2600 and higher Security Appliances. Link Aggregation and Port Redundancy are not supported for the HA Control Interface.

Topics:

- Link Aggregation on page 301
- Link Aggregation Configuration on page 302
- Port Redundancy on page 304
- Port Redundancy Configuration on page 304

Link Aggregation

Link Aggregation is used to increase the available bandwidth between the firewall and a switch by aggregating up to four interfaces into a single aggregate link, referred to as a Link Aggregation Group (LAG). All ports in an aggregate link must be connected to the same switch. The Security Appliance uses a round-robin algorithm for load balancing traffic across the interfaces in a Link Aggregation Group. Link Aggregation also provides a measure of redundancy, in that if one interface in the LAG goes down, the other interfaces remain connected.

Link Aggregation is referred to using different terminology by different vendors, including Port Channel, Ether Channel, Trunk, and Port Grouping.

Topics:

- Link Aggregation Failover on page 301
- Link Aggregation Limitations on page 302

Link Aggregation Failover

SonicWall provides multiple methods for protecting against loss of connectivity in the case of a link failure, including High Availability (HA), Load Balancing Groups (LB Groups), and now Link Aggregation. If all three of these features are configured on a Security Appliance, the following order of precedence is followed in the case of a link failure:

1 High Availability

2 Link Aggregation

3 Load Balancing Groups

HA takes precedence over Link Aggregation. Because each link in the LAG carries an equal share of the load, the loss of a link on the Active firewall forces a failover to the Idle firewall (if all of its links remain connected). Physical monitoring needs to be configured only on the primary aggregate port.

When Link Aggregation is used with a LB Group, Link Aggregation takes precedence. LB takes over only when all the ports in the aggregate link are down.

Link Aggregation Limitations

- Currently only static addressing is supported for Link Aggregation. Static port channel, which is referred to as PAG (port aggregation), is one way of configuring Ethernet port channels. No LACP or PAGP packets are sent out to form an EtherChannel with the partnering device (switch or server etc).
- A static Link Aggregation Group (LAG) configured with Ethernet port channels must be manually configured/bundled for NSA 3600 or higher Security Appliances.
- The dynamic Link Aggregation Control Protocol (LACP) is currently not supported. Dynamic, through a protocol to bundle Ethernet ports such as IEEE LACP or Cisco's PAGP, is another way of configuring Ethernet port channels. In this method, LACP or PAGP packets are sent out on the port.

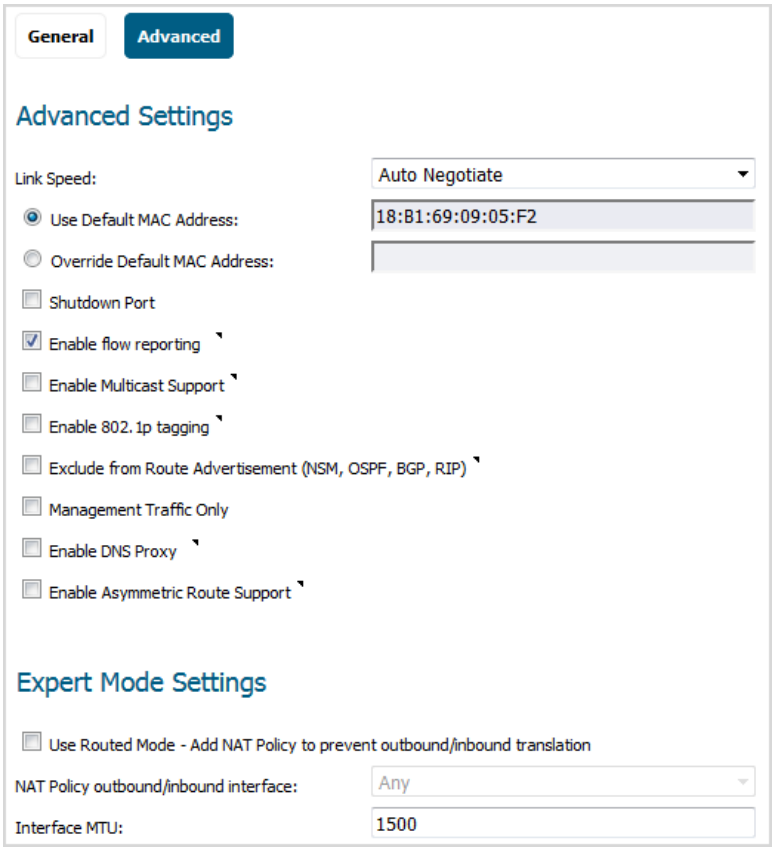
Link Aggregation Configuration

To configure Link Aggregation:

1 Navigate to **MANAGE | System Setup** | **Network > Interfaces**.

2 Click the **Configure** icon for the interface that is to be designated the master of the Link Aggregation Group. The **Edit Interface** dialog displays.

3 Click **Advanced**.



4 From **Redundant/Aggregate Ports**, select **Link Aggregation**. More options appear.

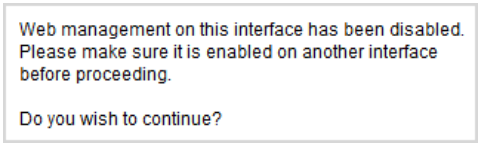


5 The **Aggregate Port** option displays with each of the currently unassigned interfaces on the Security Appliance. None of the ports are selected. Select up to three other interfaces to assign to the LAG.

> **NOTE:** After an interface is assigned to a Link Aggregation Group, its configuration is governed by the Link Aggregation master interface and it cannot be configured independently. In the **Interface Settings** table, the interface's zone is displayed as **Aggregate Port** and the **Configuration** icon is removed.

6 Set the **Link Speed** for the interface to **Auto-Negotiate**.

7 Click **OK**. If Web Management has not been configured for the interface, a message displays.



a Click **OK**.

b Enable Web Management on another interface.

IMPORTANT: Link Aggregation requires a matching configuration on the Switch. The switch's method of load balancing varies depending on the vendor. Consult the documentation for the switch for information on configuring Link Aggregation. Remember that it may be referred to as Port Channel, Ether Channel, Trunk, or Port Grouping.

Port Redundancy

Port Redundancy provides a simple method for configuring a redundant port for a physical Ethernet port. This is a valuable feature, particularly in high-end deployments, to protect against switch failures being a single point of failure.

When the primary interface is active, it processes all traffic to and from the interface. If the primary interface goes down, the secondary interface takes over all outgoing and incoming traffic. The secondary interface assumes the MAC address of the primary interface and sends the appropriate gratuitous ARP on a failover event. When the primary interface comes up again, it resumes responsibility for all traffic handling duties from the secondary interface.

In a typical Port Redundancy configuration, the primary and secondary interfaces are connected to different switches. This provides for a failover path in case the primary switch goes down. Both switches must be on the same Ethernet domain. Port Redundancy can also be configured with both interfaces connected to the same switch.

Port Redundancy Failover

SonicWall provides multiple methods for protecting against loss of connectivity in the case of a link failure, including High Availability (HA), Load Balancing Groups (LB Groups), and now Port Redundancy. If all three of these features are configured on a Security Appliance, the following order of precedence is followed in the case of a link failure:

1 Port Redundancy

2 HA

3 LB Group

When Port Redundancy is used with HA, Port Redundancy takes precedence. Typically an interface failover causes an HA failover to occur, but if a redundant port is available for that interface, then an interface failover occurs, but not an HA failover. If both the primary and secondary redundant ports go down, then an HA failover occurs (assuming the secondary Security Appliance has the corresponding port active).

When Port Redundancy is used with a LB Group, Port Redundancy again takes precedence. Any single port (primary or secondary) failures are handled by Port Redundancy just like with HA. When both the ports are down then LB kicks in and tries to find an alternate interface.

Port Redundancy Configuration

To configure Port Redundancy:

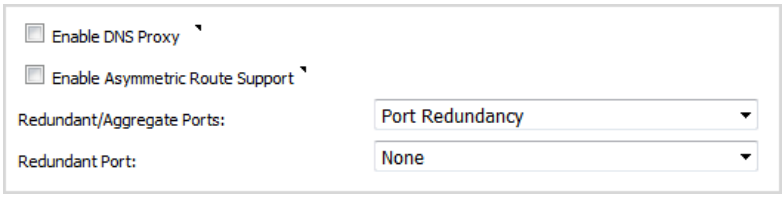
1 Navigate to **MANAGE | System Setup** | **Network > Interfaces**.

2 Click the **Configure** icon for the interface that is to be designated the master of the Link Aggregation Group. The **Edit Interface** dialog displays.

3 Click **Advanced**.

4 Set the **Link Speed** for the interface to **Auto-Negotiate**.

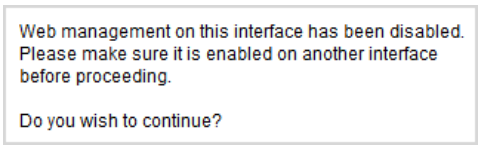
5 From **Redundant/Aggregate Ports**, select **Port Redundancy**. Another option displays.



6 The **Redundant Port** option displays all of the currently unassigned interfaces available. Select one of the interfaces; the default is **None**.

NOTE: After an interface is selected as a Redundant Port, its configuration is governed by the primary interface and it can not be configured independently. In the **Interface Settings** table, the interface's zone is displayed as **Redundant Port**, and the **Configuration** icon is removed.

7 Click **OK**. If Web Management has not been configured for the interface, a message displays.



a Click **OK**.

b Enable Web Management on another interface.

Configuring Virtual Interfaces (VLAN Subinterfaces)

When you add a VLAN subinterface, you need to assign it to a zone, assign it a VLAN Tag, and assign it to a physical interface. Based on your zone assignment, you configure the VLAN subinterface the same way you configure a physical interface for the same zone.

To add a virtual interface:

1 Navigate to **MANAGE | System Setup | Network > Interfaces**.

2 At the bottom of the **Interface Settings** table, select **Virtual Interface** from **Add Interface**. The **Add Interface** dialog displays.

3 Select a zone to assign to the interface. You can select LAN, WAN, DMZ, WLAN, or a custom zone. The zone assignment does not have to be the same as the parent (physical) interface. In fact, the parent interface can even remain **Unassigned**.

Your configuration choices for the network settings of the subinterface depend on the zone you select.

- **LAN**, **DMZ**, or a custom zone of Trusted type: **Static** or **Transparent**
- **WLAN** or a custom Wireless zone: static IP only (no IP Assignment list).

4 Assign a VLAN tag (ID) to the subinterface in the **VLAN Tag** field. Valid VLAN IDs are **0** (default) to 4094, although some switches reserve VLAN 1 for native VLAN designation, and VLAN 0 is reserved for QoS. You need to create a VLAN subinterface with a corresponding VLAN ID for each VLAN you wish to secure with your Security Appliance.

IMPORTANT: If X-Series switches are provisioned, VLAN IDs from 0 - 35 are internal VLAN IDs and cannot be used for VLAN subinterfaces.

5 Select the parent (physical) interface to which this subinterface belongs from **Parent Interface**. There is no per-interface limit to the number of subinterfaces you can assign – you may assign subinterfaces up to the system limit.

6 Configure the subinterface network settings based on the zone you selected. See the interface configuration instructions:

- Configuring a Static Interface on page 272
- Configuring Advanced Settings for a Static Interface on page 275
- Configuring Interfaces in Transparent IP Mode (Splice L3 Subnet) on page 282
- Configuring Wireless Interfaces on page 286
- Configuring a WAN Interface on page 292

7 Select the management and user-login methods for the subinterface.

8 Click **OK**.

Configuring IPS Sniffer Mode

To configure the Security Appliance for IPS Sniffer Mode, you use two interfaces in the same zone for the L2 Bridge-Pair. You can use any interfaces except the WAN interface. For this example, X2 and X3 are used for the Bridge-Pair and are configured in the LAN zone. The WAN interface (X1) is used by the Security Appliance for access to the Security Appliance Data Center as needed. The mirrored port on the switch connects to one of the interfaces in the Bridge-Pair.

Topics:

- Configuration Task List for IPS Sniffer Mode on page 306
- Configuring the Primary Bridge Interface on page 307
- Configuring the Secondary Bridge Interface on page 307
- Enabling and Configuring SNMP on page 308
- Configuring IPS Sniffer Mode on page 308

Configuration Task List for IPS Sniffer Mode

- Configure the Primary Bridge Interface
  - Select LAN as the Zone for the Primary Bridge Interface
  - Assign a static IP address
- Configure the Secondary Bridge Interface
  - Select LAN as the Zone for the Secondary Bridge Interface
  - Enable the L2 Bridge to the Primary Bridge interface
- Enable SNMP and configure the IP address of the SNMP manager system where traps can be sent
- Configure Security Services for LAN traffic
- Configure logging alert settings to “Alert” or below
- Connect the mirrored port on the switch to either one of the interfaces in the Bridge-Pair
- Connect and configure the WAN to allow access to dynamic signature data over the Internet

Configuring the Primary Bridge Interface

To configure the primary bridge interface:

1. Navigate to **MANAGE | System Setup | Network > Interfaces**.
2. Click the **Configure** icon in the right column of interface X2. The **Edit Interface** dialog displays.
3. Select **LAN** from the **Zone** drop-down menu. More options display.

   NOTE: You do not need to configure settings on the **Advanced** or **VLAN Filtering** tabs.

4. For **IP Assignment**, select **Static IP Mode.**
5. Configure the interface with a static IP Address (for example, `10.1.2.3`). The IP address you choose should not collide with any of the networks that are seen by the switch.

   NOTE: The Primary Bridge Interface must have a static IP assignment.

6. Configure the **Subnet Mask**.
7. Type in a descriptive comment.
8. Choose **Management** option(s) for the interface: **HTTPS**, **Ping**, **SNMP**, **SSH**.
9. Choose **User Login** options: **HTTP**, **HTTPS**.
10. To enable redirect to HTTPS from HTTP, select **Add rule to enable redirect from HTTP to HTTPS**. For more information about this option, see HTTP/HTTPS Redirection on page 267.
11. Click **OK**.

Configuring the Secondary Bridge Interface

Our example continues with X3 as the secondary bridge interface.

To configure the secondary bridge interface:

1. Navigate to **MANAGE | System Setup | Network > Interfaces**.
2. Click the **Configure** icon in the right column of interface X2. The **Edit Interface** dialog displays.
3. Select **LAN** from the **Zone** drop-down menu. More options display.

   NOTE: You do not need to configure settings on the Advanced or VLAN Filtering tabs.

4. From **IP Assignment**, select **Layer 2 Bridged Mode**.
5. From **Bridged to**, select the **X2** interface.
6. Do not enable the **Block all non-IPv4 traffic** setting if you want to monitor non-IPv4 traffic.
7. Select **Never route traffic on this bridge-pair** to ensure that the traffic from the mirrored switch port is not sent back out onto the network.
8. Select **Only sniff traffic on this bridge-pair** to enable sniffing or monitoring of packets that arrive on the L2 Bridge from the mirrored switch port.
9. Select **Disable stateful-inspection on this bridge-pair** to exempt these interfaces from stateful high availability inspection. If Deep Packet Inspection services are enabled for these interfaces, the DPI services continue to be applied.
10. Choose **Management** option(s) for the interface: **HTTPS**, **Ping**, **SNMP**, **SSH**.

11 Choose **User Login** options: **HTTP**, **HTTPS**.

12 To enable redirect to HTTPS from HTTP, select **Add rule to enable redirect from HTTP to HTTPS**. For more information about this option, see HTTP/HTTPS Redirection on page 267.

13 Click **OK**.

Enabling and Configuring SNMP

When SNMP is enabled, SNMP traps are automatically triggered for many events that are generated by SonicWall Security Services such as Intrusion Prevention and Gateway Anti-Virus (GAV).

More than 50 IPS and GAV events currently trigger SNMP traps. The *SonicOS Log Events Reference Guide* contains a list of events that are logged by SonicOS, and includes the SNMP trap number where applicable. This guide is available online at https://www.sonicwall.com/support/technical-documentation by selecting any SonicWall platform that runs SonicOS 6.5.

To determine the traps that are possible when using IPS Sniffer Mode with Intrusion Prevention enabled, search for **Intrusion** in the table found in the Index of Log Event Messages section in the *SonicOS Log Events Reference Guide*. The SNMP trap number, if available for that event, is printed in the **SNMP Trap Type** column of the table.

To determine the possible traps with Gateway Anti-Virus enabled, search the table for **Security Services**, and view the SNMP trap number in the **SNMP Trap Type** column.

To enable and configure SNMP:

1 Navigate to **MANAGE | System Setup | Appliance | SNMP**.

2 Select **Enable SNMP**.

3 Click **ACCEPT**. **CONFIGURE** becomes active and the **View**, **User/Group**, and **Access** sections are displayed.

4 Click **Configure**. The **SNMP Settings** dialog displays.

5 In the **System Name** field, type the name of the SNMP manager system that receives the traps sent from the Security Appliance.

6 Enter the name or email address of the contact person for the SNMP Contact in the **System Contact** field.

7 Enter a description of the system location, such as `3rd floor lab`, in the **System Location** field.

8 Enter the system's asset number in the **Asset Number** field.

9 In the **Get Community Name** field, type the community name that has permissions to retrieve SNMP information from the firewall, for example, `public`.

10 In the **Trap Community Name** field, type the community name that is used to send SNMP traps from the firewall to the SNMP manager, for example, `public`.

11 In the **Host 1/2/3/4** fields, type in the IP address(es) of the SNMP manager system(s) that receives the traps.

12 Click **OK**.

Configuring IPS Sniffer Mode

To configure IPS sniffer mode:

1 Navigate to **MANAGE | System Setup | Network > Interfaces**.

2 Click on the **Edit** icon for the **X2** interface. The **Edit Interface** dialog displays.

3 Set the **Mode / IP Assignment** to **Layer 2 Bridged Mode**. The options change.

4 Set the **Bridged To:** interface to **X0**.

5 Select **Only sniff traffic on the bridge-pair**.

6 Click **OK** to save and activate the change. The dialog closes, and the **Network > Interfaces** page redisplays.

7 Click the **Edit** icon for the **X1 WAN** interface. The **Edit Interface** dialog displays.

8 Assign the X1 WAN interface a unique IP address for the *internal* LAN segment of your network — this may sound wrong, but this is actually be the interface from which you manage the appliance, and it is also the interface from which the Security Appliance sends its SNMP traps as well as the interface from which it gets security services signature updates.

9 Click **OK**.

10 For traffic to pass successfully, you must also modify the firewall rules to allow traffic from the

- LAN to WAN,
- WAN to the LAN

11 Connect the:

- Span/mirror switch port to X0 on the Security Appliance, not to X2 (in fact, X2 is not plugged in at all)
- X1 to the internal network

IMPORTANT: Use care when programming ports spanned/mirrored to X0.

VIDEO: Informational videos with interface configuration examples are available online. For example, see How to configure the SonicWall WAN / X1 Interface with PPPoE Connection. Additional videos are available at: https://support.SonicWall.com/videos-product-select.

Configuring Security Services (Unified Threat Management)

The settings that you enable in this section control what type of malicious traffic you detect in IPS Sniffer Mode. Typically, you want to enable Intrusion Prevention, but you might also want to enable other Security Services, such as Gateway Anti-Virus or Anti-Spyware.

To enable Security Services, your SonicWall Security Appliance must be licensed for them and the signatures must be downloaded from the SonicWall Data Center. For complete instructions on enabling and configuring IPS, GAV, and Anti-Spyware, see *SonicOS Security Configuration*.

Topics:

- Configuring Logging on page 309
- Connecting the Mirrored Switch Port to a IPS Sniffer Mode Interface on page 310
- Connecting and Configuring the WAN Interface to the Data Center on page 310

Configuring Logging

You can configure logging on the **Log > Settings** page to record entries for attacks that are detected by the firewall. For how to enable logging, see *SonicOS Logs and Reporting*.

Connecting the Mirrored Switch Port to a IPS Sniffer Mode Interface

Use a standard Cat-5 Ethernet cable to connect the mirrored switch port to either interface in the Bridge-Pair. Network traffic is sent automatically from the switch to the Security Appliance where it can be inspected.

Consult the switch documentation for instructions on setting up the mirrored port.

Connecting and Configuring the WAN Interface to the Data Center

Connect the WAN port on the Security Appliance, typically port X1, to your gateway or to a device with access to the gateway. The Security Appliance communicates with the SonicWall Data Center automatically. For detailed instructions on configuring the WAN interface, see Configuring a WAN Interface on page 292.

Configuring Wire and Tap Mode

SonicOS supports Wire Mode and Tap Mode, which provide methods of non-disruptive, incremental insertion into networks. Wire and Tap mode settings describes the wire and tap modes.

NOTE: Wire mode is supported on NSA 2600 and higher appliances.

Wire and Tap mode settings

Wire mode setting	Description
Bypass Mode	Bypass Mode allows for the quick and relatively non-interruptive introduction of Security Appliance hardware into a network. Upon selecting a point of insertion into a network (for example, between a core switch and a perimeter Security Appliance, in front of a VM server farm, at a transition point between data classification domains), the Security Appliance is inserted into the physical data path, requiring a very short maintenance window. One or more pairs of switch ports on the Security Appliance are used to forward all packets across segments at full line rates, with all the packets remaining on the Security Appliance's 240Gbps switch fabric rather than getting passed up to the multi-core inspection and enforcement path. While Bypass Mode does not offer any inspection or firewalling, this mode allows you to physically introduce the Security Appliance into the network with a minimum of downtime and risk, and to obtain a level of comfort with the newly inserted component of the networking and security infrastructure. You can then transition from Bypass Mode to Inspect or Secure Mode instantaneously through a simple user-interface driven reconfiguration.
Inspect Mode	Inspect Mode extends Bypass Mode without functionally altering the low-risk, zero-latency packet path. Packets continue to pass through the Security Appliance's switch fabric, but they are also mirrored to the multi-core RF-DPI engine for the purposes of passive inspection, classification, and flow reporting. This reveals the Security Appliance's Application Intelligence and threat detection capabilities without any actual intermediate processing.

Wire and Tap mode settings (Continued)

Wire mode setting	Description
Secure Mode	Secure Mode is the progression of Inspect Mode, actively interposing the Security Appliance's multi-core processors into the packet processing path. This unleashes the inspection and policy engines' full-set of capabilities, including Application Intelligence and Control, Intrusion Prevention Services, Gateway and Cloud-based Anti-Virus, Anti-Spyware, and Content Filtering. Secure Mode affords the same level of visibility and enforcement as conventional NAT or L2 Bridged Mode deployments, but without any L3/L4 transformations, and with no alterations of ARP or routing behavior. Secure Mode thus provides an incrementally attainable NGFW deployment requiring no logical and only minimal physical changes to existing network designs. Secure mode should be used when creating wire-mode pairs for VLAN translation.
Tap Mode	Tap Mode provides the same visibility as Inspect Mode, but differs from the latter in that it ingests a mirrored packet stream through a single switch port on the Security Appliance, eliminating the need for physically intermediated insertion. Tap Mode is designed for use in environments employing network taps, smart taps, port mirrors, or SPAN ports to deliver packets to external devices for inspection or collection. Like all other forms of Wire Mode, Tap Mode can operate on multiple concurrent port instances, supporting discrete streams from multiple taps.

Wire modes: Functional differences summarizes the key functional differences between modes of interface configuration:

Wire modes: Functional differences

Interface configuration	Bypass mode	Inspect mode	Secure mode	Tap mode	L2 Bridge, Transparent, NAT, Route modes
Active/Active Clustering [a]	No	No	No	No	Yes
Application Control	No	No	Yes	No	Yes
Application Visibility	No	Yes	Yes	Yes	Yes
ARP/Routing/NAT [a]	No	No	No	No	Yes
Comprehensive Anti-Spam Service [a]	No	No	No	No	Yes
Content Filtering	No	No	Yes	No	Yes
DHCP Server [a]	No	No	No	No	Yes [b]
DPI Detection	No	Yes	Yes	Yes	Yes
DPI Prevention	No	No	Yes	No	Yes
DPI-SSL[a]	No	No	Yes	No	Yes
High-Availability	Yes	Yes	Yes	Yes	Yes
Link-State Propagation [c]	Yes	Yes	Yes	No	No
Stateful Packet Inspection	No	Yes	Yes	Yes	Yes
TCP Handshake Enforcement [d]	No	No	No	No	Yes
Virtual Groups [a]	No	No	No	No	Yes
VLAN Translation [e]	No	No	Yes	No	No

a. These functions or services are unavailable on interfaces configured in Wire Mode, but remain available on a system-wide level for any interfaces configured in other compatible modes of operation.

b. Not available in L2 Bridged Mode.

c. **Link State Propagation** is a feature whereby interfaces in a Wire Mode pair mirror the link-state triggered by transitions of their partners. This is essential to proper operations in redundant path networks. Link State Propagation is not supported in Wire Mode over VLAN interfaces.

d. Disabled by design in Wire Mode to allow for failover events occurring elsewhere on the network to be supported when multiple Wire Mode paths, or when multiple Security Appliance units are in use along redundant or asymmetric paths.

e. VLAN Translation is not supported in Wire Mode over VLAN interfaces.

NOTE: When operating in Wire Mode, the firewall's dedicated Management interface is used for local management. To enable remote management and dynamic security services and application intelligence updates, a WAN interface (separate from the Wire Mode interfaces) must be configured for Internet connectivity. This is easily done given that SonicOS supports interfaces in mixed-modes of almost any combination.

Configuring an Interface for Wire Mode

Wire Mode can be configured on WAN, LAN, DMZ, and custom zones (except wireless zones). Wire Mode is a simplified form of Layer 2 Bridged Mode, and is configured as a pair of interfaces. In Wire Mode, the destination zone is the **Paired Interface Zone**. Access rules are applied to the Wire Mode pair based on the direction of traffic between the source **Zone** and its **Paired Interface Zone**. For example, if the source **Zone** is **WAN** and the **Paired Interface Zone** is **LAN**, then WAN to LAN and LAN to WAN rules are applied, depending on the direction of the traffic.

In Wire Mode, you can enable **Link State Propagation**, which propagates the link status of an interface to its paired interface. If an interface goes down, its paired interface is forced down to mirror the link status of the first interface. Both interfaces in a Wire Mode pair always have the same link status.

In Wire Mode, you can **Disable Stateful Inspection**. When **Disable Stateful Inspection** is selected, Stateful Packet Inspection is turned off. When **Disable Stateful Inspection** is *not* selected, new connections can be established without enforcing a 3-way TCP handshake. **Disable Stateful Inspection** must be selected if asymmetrical routes are deployed.

To configure an interface for Wire Mode:

1 Navigate to **MANAGE | System Setup | Network > Interfaces**,

2 Click the **Configure** icon for the interface you want to configure for Wire Mode. The **Edit Interface** dialog displays.

3 From **Zone**, select any zone type except WLAN.

4 From **Mode / IP Assignment**, to configure the Interface for:

- Tap mode, select **Tap Mode (1-Port Tap)**
- Wire Mode, select **Wire Mode (2-Port Wire)**.

5 From **Wire Mode Type**, select the appropriate mode:

- **Bypass (via Internal Switch/Relay)**
- **Inspect (Passive DPI of Mirrored Traffic)**
- **Secure (Active DPI of Inline Traffic)**

6 From **Paired Interface**, select the interface that connects to the upstream Security Appliance. The paired interfaces must be of the same type (two 1 GB interfaces or two 10 GB interfaces).

NOTE: Only unassigned interfaces are available from **Paired Interface**. To make an interface unassigned, click its **Configure**, and from **Zone**, select **Unassigned**.

7 Click **OK**.

Configuring Wire Mode for a WAN/LAN Zone Pair

The following configuration is an example of how Wire Mode can be configured. This example is for a WAN zone paired with a LAN zone. Wire Mode can also be configured for DMZ and custom zones.

To configure Wire Mode for a WAN/LAN Zone Pair:

1 Navigate to **MANAGE | System Setup | Network > Interfaces**,

2 Click one of these:

- **Add Interface**.
- **Configure** icon for the interface you want to configure.

The **Add/Edit Interface** dialog displays.

3 From **IP Assignment**, select **Wire Mode (2-Port Wire)**.

4 From **Zone**, select **WAN**.

5 From **Paired Interface Zone**, select **LAN**.

6 Select the **Disable Stateful Inspection** option.

7 Select the **Enable Link State Propagation** option.

8 Click **OK**. The **Interface Settings** table is updated:

Wire Mode with Link Aggregation

NOTE: Wire Mode over VLAN interfaces does not support Link Aggregation.

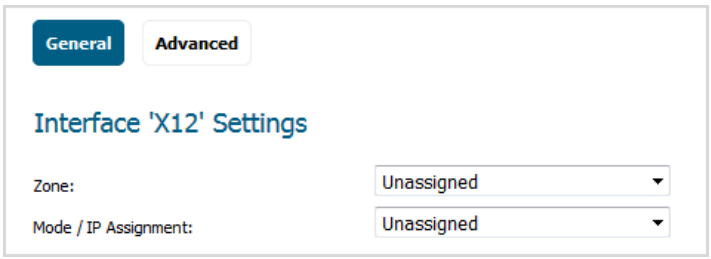
Link Aggregation (LAG) is used to bundle multiple links into a single interface to increase bandwidth. To inspect traffic over a LAG interface, a SonicWall Security Appliance can be connected inline, allowing packets sent on one link to be bridged across to the destination transparently. Existing Wire Mode features such as link state propagation are supported. Up to 8 members per LAG are supported.

Wire Mode and Link Aggregation are configured from **Network > Interfaces**. When **Link Aggregation** is selected on the **Edit Interface > Advanced** dialog, it also lists unassigned interfaces. You can select member interfaces for each side of the Wire Mode connection. The number of members on each side must be equal. It is recommended that the type and bandwidth size of the member interfaces also match.

To configure Wire Mode with LAG:

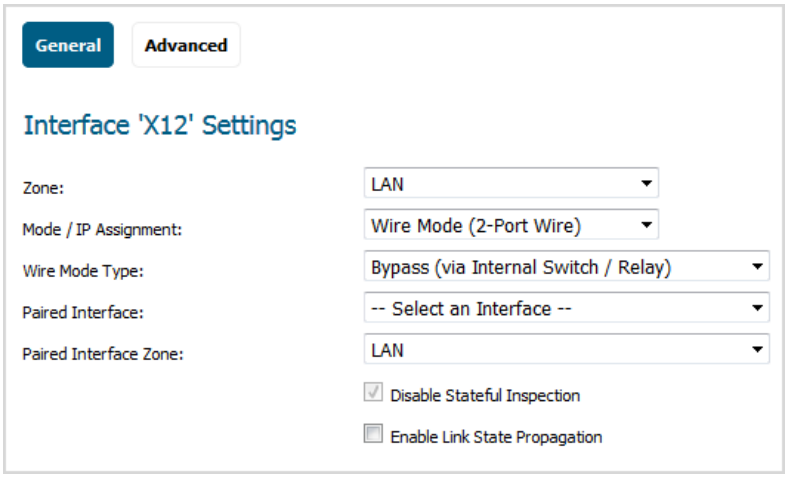
1 Navigate to **MANAGE | System Setup | Network > Interfaces**.

2 Click the **Configure** icon for the interface you want to configure. The **Edit Interface** dialog displays.



3 From **Zone**, select the zone you want. The options change.

4 From **Mode / IP Assignment**, select **Wire Mode (2-Port Wire)**. The options change again.



5 From **Wire Mode Type**, select **Secure (Active DPI of Inline Traffic)**.

6 From **Paired Interface**, select the interface to be paired.

7 From **Paired Interface Zone**, select the zone of the interface to be paired.

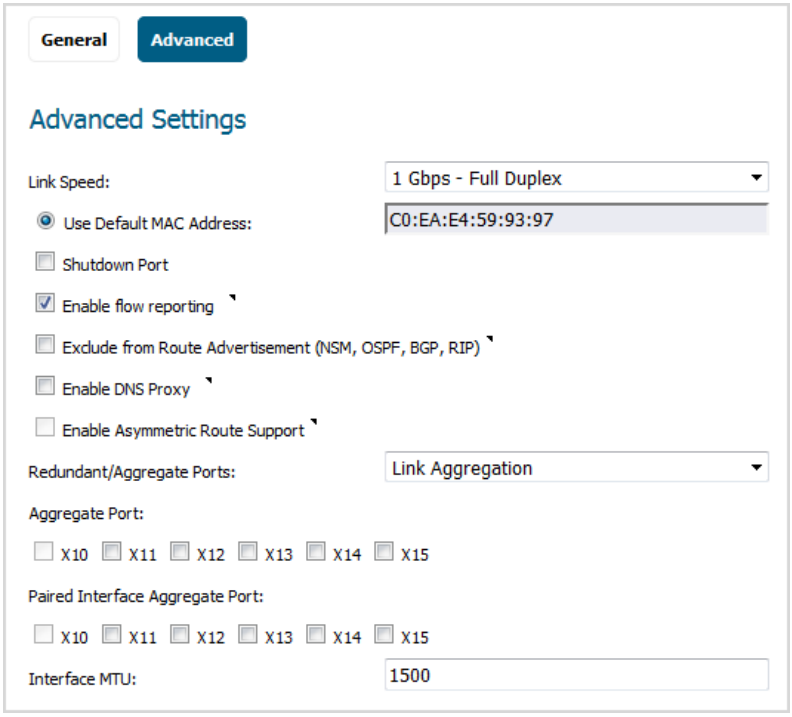
8 Select the **Disable Stateful Inspection** option. This option is selected by default.

9 Optionally, select the **Enable Link State Propagation** option if you want it. This option is not selected by default.

10 Click **Advanced**.

To continue on Advanced:

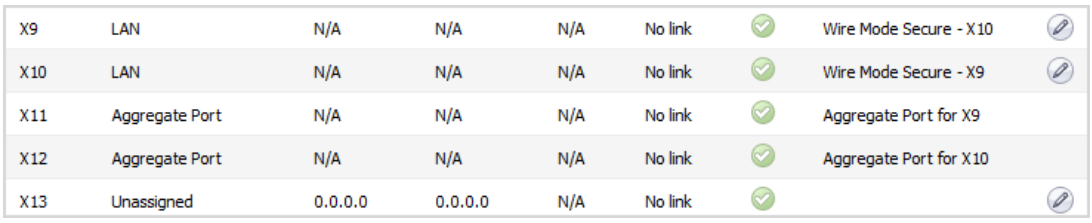
1 From **Redundant/Aggregate Ports**, select **Link Aggregation**. The options change.



2 From **Aggregate Port**, select the port for aggregation.

3 From **Paired Interface Aggregate Port**, select the paired port for aggregation.

4 Click **OK**. The configuration is displayed in the **Interface Settings** table on **Network > Interfaces**.



X9	LAN	N/A	N/A	N/A	No link		Wire Mode Secure - X10	
X10	LAN	N/A	N/A	N/A	No link		Wire Mode Secure - X9	
X11	Aggregate Port	N/A	N/A	N/A	No link		Aggregate Port for X9	
X12	Aggregate Port	N/A	N/A	N/A	No link		Aggregate Port for X10	
X13	Unassigned	0.0.0.0	0.0.0.0	N/A	No link			

Layer 2 Bridged Mode

SonicOS includes **L2 (Layer 2) Bridged Mode**, a method of unobtrusively integrating a Security Appliance into any Ethernet network. L2 Bridged Mode is ostensibly similar to SonicOS's Transparent Mode in that it enables a Security Appliance to share a common subnet across two interfaces, and to perform stateful and deep-packet inspection on all traversing IP traffic, but it is functionally more versatile.

In particular, L2 Bridged Mode employs a secure learning bridge architecture, enabling it to pass and inspect traffic types that cannot be handled by many other methods of transparent Security Appliance integration. Using L2 Bridged Mode, a SonicWall Security Appliance can be non-disruptively added to any Ethernet network to provide in-line deep-packet inspection for all traversing IPv4 TCP and UDP traffic. In this scenario, the Security Appliance is not used for security enforcement, but instead for bidirectional scanning, blocking viruses and spyware, and stopping intrusion attempts.

Unlike other transparent solutions, L2 Bridged Mode can pass all traffic types, including IEEE 802.1Q VLANs, Spanning Tree Protocol, multicast, broadcast, and IPv6, ensuring that all network communications continues uninterrupted.

Another aspect of the versatility of L2 Bridged Mode is that you can use it to configure IPS Sniffer Mode. Supported on SonicWall Security Appliances, IPS Sniffer Mode uses a single interface of a Bridge-Pair to monitor network traffic from a mirrored port on a switch. IPS Sniffer Mode provides intrusion detection, but cannot block malicious traffic because the Security Appliance is not connected inline with the traffic flow. For more information about IPS Sniffer Mode, see IPS Sniffer Mode on page 264.

L2 Bridged Mode provides an ideal solution for networks that already have an existing Security Appliance, and do not have immediate plans to replace their existing Security Appliance, but wish to add the security of SonicWall deep-packet inspection and security services, such as Intrusion Prevention Services, Gateway Anti Virus, and Gateway Anti Spyware. If you do not have SonicWall security services subscriptions, you may sign up for free trials from MySonicWall.

You can also use L2 Bridged Mode in a High Availability deployment. This scenario is explained in the Layer 2 Bridged Mode with High Availability on page 331.

NOTE: Link Aggregation is not supported in Layer 2 Bridged Mode.

Topics:

- Key Features of SonicOS Layer 2 Bridged Mode on page 316
- Key Concepts to Configuring L2 Bridged Mode and Transparent Mode on page 316
- Comparing L2 Bridged Mode to Transparent Mode on page 318
- L2 Bridge Path Determination on page 325
- L2 Bridge Interface Zone Selection on page 326
- Sample Topologies on page 328

Key Features of SonicOS Layer 2 Bridged Mode

SonicOS Layer 2 Bridged Mode: Key features and benefits outlines the benefits of each key feature of layer 2 bridged mode.

SonicOS Layer 2 Bridged Mode: Key features and benefits

Feature	Benefit
L2 Bridging with Deep Packet Inspection	This method of transparent operation means that a SonicWall Security Appliance can be added to any network without the need for readdressing or reconfiguration, enabling the addition of deep-packet inspection security services with no disruption to existing network designs. Developed with connectivity in mind as much as security, L2 Bridged Mode can pass all Ethernet frame types, ensuring seamless integration.
Secure Learning Bridge Architecture	True L2 behavior means that all allowed traffic flows natively through the L2 Bridge. Whereas other methods of transparent operation rely on ARP and route manipulation to achieve transparency, which frequently proves problematic, L2 Bridged Mode dynamically learns the topology of the network to determine optimal traffic paths.
Universal Ethernet Frame-Type Support	All Ethernet traffic can be passed across an L2 Bridge, meaning that all network communications continue uninterrupted. While many other methods of transparent operation only support IPv4 traffic, L2 Bridged Mode inspects all IPv4 traffic and passes (or blocks, if desired) all other traffic, including LLC, all Ethertypes, and even proprietary frame formats.
Mixed-Mode Operation	L2 Bridged Mode can concurrently provide L2 Bridging and conventional Security Appliance services, such as routing, NAT, VPN, and wireless operations. This means it can be used as an L2 Bridge for one segment of the network, while providing a complete set of security services to the remainder of the network. This also allows for the introduction of the SonicWall Security Appliance as a pure L2 bridge, with a smooth migration path to full security services operation.
Wireless Layer 2 Bridging	Use a single IP subnet across multiple zone types, including LAN, WLAN, DMZ, or custom zones. This feature allows wireless and wired clients to seamlessly share the same network resources, including DHCP addresses. The Layer 2 protocol can run between paired interfaces, allowing multiple traffic types to traverse the bridge, including broadcast and non-ip packets.

Key Concepts to Configuring L2 Bridged Mode and Transparent Mode

The following terms are used when referring to the operation and configuration of L2 Bridged Mode:

- **L2 Bridged Mode** – A method of configuring a SonicWall Security Appliance, which enables it to be inserted inline into an existing network with absolute transparency, beyond even that provided by Transparent Mode. Layer 2 Bridged Mode also refers to the *IP Assignment* configuration that is selected for *Secondary Bridge Interfaces* that are placed into a *Bridge-Pair*.
- **Transparent Mode** – A method of configuring a SonicWall Security Appliance that allows it to be inserted into an existing network without the need for IP reconfiguration by spanning a single IP subnet across two or more interfaces through the use of automatically applied ARP and routing logic.
- **IP Assignment** – When configuring a Trusted (LAN) or Public (DMZ) interface, the IP Assignment for the interface can either be:
    - **Static** – The IP address for the interface is manually entered.

    - **Transparent Mode** – The IP address(es) for the interface is assigned using an Address Object (Host, Range, or Group) that falls within the WAN Primary IP subnet, effectively spanning the subnet from the WAN interface to the assigned interface.
    - **Layer 2 Bridged Mode** – An interface placed in this mode becomes the *Secondary Bridge Interface* to the *Primary Bridge Interface* to which it is paired. The resulting Bridge-Pair then behaves like a two-port learning bridge with full L2 transparency, and all IP traffic that passes through is subjected to full stateful failover and deep packet inspection.

- **Bridge-Pair** – The logical interface set composed of a *Primary Bridge Interface* and a *Secondary Bridge Interface*. The terms primary and secondary do not imply any inherent level of operational dominance or subordination; both interfaces continue to be treated according to their zone type, and to pass IP traffic according to their configured Access Rules. Non-IPv4 traffic across the Bridge-Pair is controlled by the *Block all non-IPv4 traffic* setting on the *Secondary Bridge Interface*. A system may support as many Bridge Pairs as it has interface pairs available. In other words, the maximum number of Bridge-Pairs is equal to ½ the number of physical interfaces on the platform. Membership in a Bridge-Pair does not preclude an interface from conventional behavior; for example, if X1 is configured as a *Primary Bridge Interface* paired to X3 as a *Secondary Bridge Interface*, X1 can simultaneously operate in its traditional role as the Primary WAN, performing NAT for Internet-bound traffic through the *Auto-added X1 Default NAT Policy*.
- **Primary Bridge Interface** – A designation that is assigned to an interface after a *Secondary Bridge Interface* has been paired to it. A Primary Bridge Interface can belong to an Untrusted (WAN), Trusted (LAN), or Public (DMZ) zone.
- **Secondary Bridge Interface** – A designation that is assigned to an interface whose *IP Assignment* has been configured for *Layer 2 Bridged Mode*. A Secondary Bridge Interface can belong to a Trusted (LAN), or Public (DMZ) zone.
- **Bridge Management Address** – The address of the Primary Bridge Interface is shared by both interfaces of the *Bridge-Pair*. If the Primary Bridge Interface also happens to be the Primary WAN interface, it is this address that is used for outbound communications by the Security Appliance, such as NTP, and License Manager updates. Hosts that are connected to either segment of the Bridge-Pair may also use the Bridge Management Address as their gateway, as is common in *Mixed-Mode* deployments.
- **Bridge-Partner** – The term used to refer to the other member of a *Bridge-Pair*.
- **Non-IPv4 Traffic** - SonicOS supports the following IP protocol types: ICMP (1), IGMP (2), TCP (6), UDP (17), GRE (47), ESP (50), AH (51), EIGRP (88), OSPF (89), PIM-SM (103), L2TP (115). More esoteric IP types, such as Combat Radio Transport Protocol (126), are not natively handled by the Security Appliance, nor are non-IPv4 traffic types such as IPX or (currently) IPv6. L2 Bridged Mode can be configured to either pass or drop Non-IPv4 traffic.
- **Captive-Bridged Mode** – This optional mode of L2 Bridge operation prevents traffic that has entered an L2 bridge from being forwarded to a non-Bridge-Pair interface. By default, L2 Bridge logic forwards traffic that has entered the L2 Bridge to its destination along the most optimal path as determined by ARP and routing tables. In some cases, the most optimal path might involve routing or NATing to a non-Bridge-Pair interface. Activating Captive-Bridged Mode ensures that traffic that enters an L2 Bridge exits the L2 Bridge rather than taking its most logically optimal path. In general, this mode of operation is only required in complex networks with redundant paths, where strict path adherence is required.
- **Pure L2 Bridge Topology** – Refers to deployments where the Security Appliance is used strictly in *L2 Bridged Mode* for the purposes of providing in-line security to a network. This means that all traffic entering one side of the *Bridge-Pair* is bound for the other side, and is not routed/NATed through a different interface. This is common in cases where there is an existing perimeter Security Appliance, or where in-line security is desired along some path (for example, inter-departmentally, or on a trunked link between two switches) of an existing network. Pure L2 Bridge Topology is not a functional limitation, but rather a topological description of a common deployment in heterogeneous environments.
- **Mixed-Mode Topology** – Refers to deployments where the *Bridge-Pair* are not the only point of ingress/egress through the Security Appliance. This means that traffic entering one side of the *Bridge-Pair* may be destined to be routed/NATed through a different interface. This is common when the

Security Appliance is simultaneously used to provide security to one or more Bridge-Pair while also providing:

- Perimeter security, such as WAN connectivity, to hosts on the Bridge-Pair or on other interfaces.
- Firewall and Security services to additional segments, such as Trusted (LAN) or Public (DMZ) interface, where communications occur between hosts on those segments and hosts on the Bridge-Pair.
- Wireless services with SonicPoints, where communications occur between wireless clients and hosts on the Bridge-Pair.

Comparing L2 Bridged Mode to Transparent Mode

While Transparent Mode allows a Security Appliance running SonicOS to be introduced into an existing network without the need for re-addressing, it presents a certain level of disruptiveness, particularly with regard to ARP, VLAN support, multiple subnets, and non-IPv4 traffic types. Consider a scenario where a Transparent Mode SonicWall Security Appliance has just been added to the network with a goal of minimally disruptive integration, particularly:

- Negligible or no unscheduled downtime
- No need to re-address any portion of the network
- No need to reconfigure or otherwise modify the gateway router (as is common when the router is owned by the ISP)

Topics:

- ARP in Transparent Mode on page 318
- VLAN Support in Transparent Mode on page 319
- Multiple Subnets in Transparent Mode on page 319
- Non-IPv4 Traffic in Transparent Mode on page 319
- ARP in L2 Bridged Mode on page 319
- VLAN Support in L2 Bridged Mode on page 320
- L2 Bridge IP Packet Path on page 321
- Multiple Subnets in L2 Bridged Mode on page 322
- Non-IPv4 Traffic in L2 Bridged Mode on page 322
- Comparison of L2 Bridged Mode to Transparent Mode on page 323
- Benefits of Transparent Mode over L2 Bridged Mode on page 325

ARP in Transparent Mode

ARP (Address Resolution Protocol: the mechanism by which unique hardware addresses on network interface cards are associated to IP addresses) is *proxied* in Transparent Mode. If the Workstation on Server on the left had previously resolved the Router (`192.168.0.1`) to its MAC address `00:99:10:10:10:10`, this cached ARP entry would have to be cleared before these hosts could communicate through the Security Appliance. This is because the Security Appliance proxies (or answers on behalf of) the gateway's IP (`192.168.0.1`) for hosts connected to interfaces operating in Transparent Mode. So when the Workstation at the left attempts to resolve `192.168.0.1`, the ARP request it sends is responded to by the Security Appliance with its own X0 MAC address (`00:06:B1:10:10:10`).

The Security Appliance also proxy ARPs the IP addresses specified in the Transparent Range (`192.168.0.100` to `192.168.0.250`) assigned to an interface in Transparent Mode for ARP requests received on the X1 (Primary

WAN) interface. If the Router had previously resolved the Server (192.168.0.100) to its MAC address 00:AA:BB:CC:DD:EE, this cached ARP entry would have to be cleared before the router could communicate with the host through the Security Appliance. This typically requires a flushing of the router's ARP cache either from its management interface or through a reboot. When the router's ARP cache is cleared, the router can then send a new ARP request for 192.168.0.100, to which the Security Appliance responds with its X1 MAC 00:06:B1:10:10:11.

VLAN Support in Transparent Mode

While the network depicted in the above diagram is simple, it is not uncommon for larger networks to use VLANs for segmentation of traffic. If this was such a network, where the link between the switch and the router was a VLAN trunk, a Transparent Mode SonicWall Security Appliance would have been able to terminate the VLANs to subinterfaces on either side of the link, but it would have required unique addressing; that is, non-Transparent Mode operation requiring re-addressing on at least one side. This is because only the Primary WAN interface can be used as the *source* for Transparent Mode address space.

Multiple Subnets in Transparent Mode

It is also common for larger networks to employ multiple subnets, be they on a single wire, on separate VLANs, multiple wires, or some combination. Transparent Mode is capable of supporting multiple subnets through the use of Static ARP and Route entries.

Non-IPv4 Traffic in Transparent Mode

Transparent Mode drops (and generally logs) all non-IPv4 traffic, precluding it from passing other traffic types, such as IPX, or unhandled IP types.

L2 Bridged Mode addresses these common Transparent Mode deployment issues and is described in these sections:

- ARP in L2 Bridged Mode on page 319
- VLAN Support in L2 Bridged Mode on page 320
- L2 Bridge IP Packet Path on page 321
- Multiple Subnets in L2 Bridged Mode on page 322
- Non-IPv4 Traffic in L2 Bridged Mode on page 322
- Comparison of L2 Bridged Mode to Transparent Mode on page 323
- Benefits of Transparent Mode over L2 Bridged Mode on page 325

ARP in L2 Bridged Mode

L2 Bridged Mode employs a learning bridge design where it dynamically determines which hosts are on which interface of an L2 Bridge (referred to as a Bridge-Pair). ARP is passed through natively, meaning that a host communicating across an L2 Bridge sees the actual host MAC addresses of their peers. For example, the Workstation communicating with the Router (192.168.0.1) sees the router as 00:99:10:10:10:10, and the Router sees the Workstation (192.168.0.100) as 00:AA:BB:CC:DD:EE.

This behavior allows for a SonicWall Security Appliance operating in L2 Bridged Mode to be introduced into an existing network with no disruption to most network communications other than that caused by the momentary discontinuity of the physical insertion.

NOTE: Stream-based TCP protocols communications (for example, an FTP session between a client and a server) needs to be re-established upon the insertion of an L2 Bridged Mode Security Appliance. This is by design so as to maintain the security afforded by stateful packet inspection. As the stateful packet inspection engine can not have knowledge of the TCP connections which pre-existed it, it drops these *established* packets with a log event such as *TCP packet received on non-existent/closed connection; TCP packet dropped*.

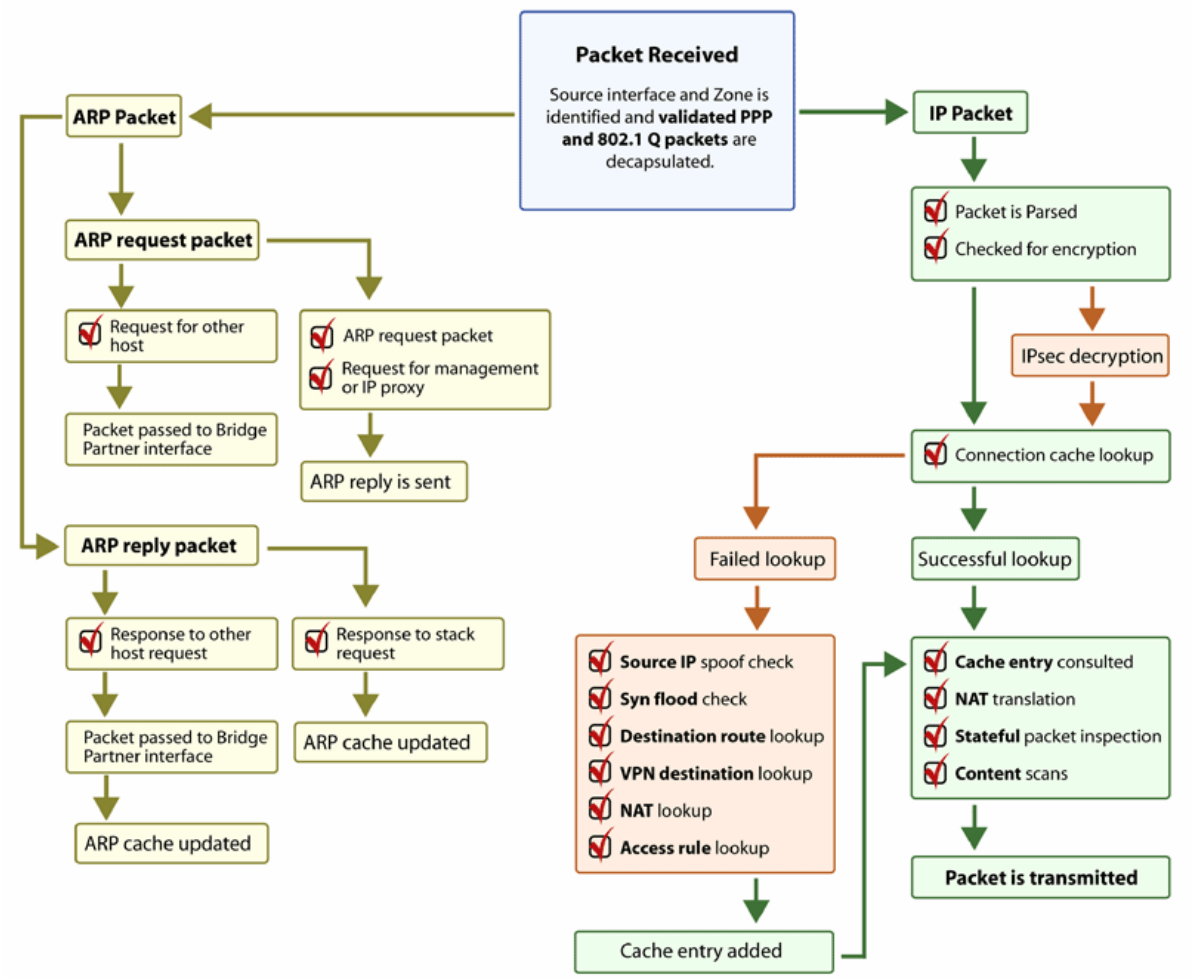
VLAN Support in L2 Bridged Mode

On SonicWall Security Appliances, L2 Bridged Mode provides fine control over 802.1Q VLAN traffic traversing an L2 Bridge. The default handling of VLANs is to allow and preserve all 802.1Q VLAN tags as they pass through an L2 Bridge, while still applying all firewall rules, and stateful and deep-packet inspection to the encapsulated traffic. It is further possible to specify white/black lists for allowed/disallowed VLAN IDs through the L2 Bridge.

This allows a Security Appliance operating in L2 Bridged Mode to be inserted, for example, inline into a VLAN trunk carrying any number of VLANs, and to provide full security services to all IPv4 traffic traversing the VLAN without the need for explicit configuration of any of the VLAN IDs or subnets. Access Rules can also, optionally, be applied to all VLAN traffic passing through the L2 Bridged Mode because of the method of handling VLAN traffic.

L2 Bridge IP Packet Path

L2 Bridge IP packet flow



The following sequence of events describes the flow in L2 Bridge IP packet flow:

1 802.1Q encapsulated frame enters an L2 Bridge interface (this first step, Step 2, and Step 12 apply only to 802.1Q VLAN traffic).

2 The 802.1Q VLAN ID is checked against the VLAN ID white/black list. If the VLAN ID is:

- Disallowed, the packet is dropped and logged.
- Allowed, the packet is de-capsulated, the VLAN ID is stored, and the inner packet (including the IP header) is passed through the full packet handler.

3 As any number of subnets is supported by L2 Bridging, no source IP spoof checking is performed on the source IP of the packet. It is possible to configure L2 Bridges to only support a certain subnet or subnets using Access Rules.

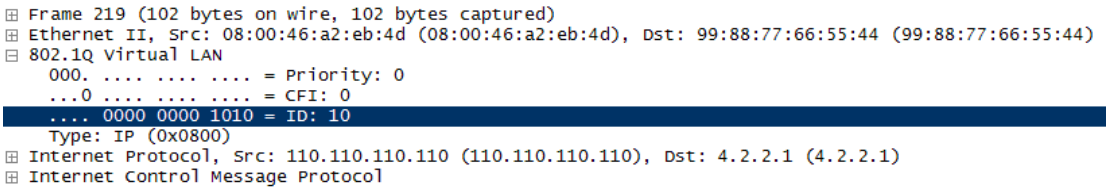
4 SYN Flood checking is performed.

5 A destination route lookup is performed to the destination zone, so that the appropriate Access rule can be applied. Any zone is a valid destination, including the same zone as the source zone (for example, LAN to LAN), the Untrusted zone (WAN), the Encrypted (VPN), Wireless (WLAN), Multicast, or custom zones of any type.

6 A NAT lookup is performed and applied, as needed:

- In general, the destination for packets entering an L2 Bridge is the *Bridge-Partner* interface (that is, the other side of the bridge). In these cases, no translation is performed.
- In cases where the L2 Bridge Management Address is the gateway, as is sometimes the case in *Mixed-Mode topologies*, then NAT is applied as needed (for more details, see L2 Bridge Path Determination on page 325).

7 Access Rules are applied to the packet. For example, on SonicWall Security Appliances, the following packet decode shows an ICMP packet bearing VLAN ID 10, source IP address `110.110.110.110` destined for IP address `4.2.2.1`.



It is possible to construct an Access Rule to control any IP packet, independent of its VLAN membership, by any of its IP elements, such as source IP, destination IP, or service type. If the packet is disallowed, it is dropped and logged. If the packet is allowed, it continues.

8 A connection cache entry is made for the packet, and required NAT translations (if any) are performed.

9 Stateful packet inspection and transformations are performed for TCP, VoIP, FTP, MSN, Oracle, RTSP and other media streams, PPTP and L2TP. If the packet is disallowed, it is dropped and logged. If the packet is allowed, it continues.

10 Deep packet inspection, including GAV, IPS, Anti-Spyware, CFS and email-filtering is performed. If the packet is disallowed, it is dropped and logged. If the packet is allowed, it continues. Client notification is performed as configured.

11 If the packet is destined for the Encrypted zone (VPN), the Untrusted zone (WAN), or some other connected interface (the last two of which might be the case in Mixed-Mode Topologies) the packet is sent through the appropriate path.

12 If the packet is not destined for the VPN/WAN/Connected interface, the stored VLAN tag is restored, and the packet (again bearing the original VLAN tag) is sent out the *Bridge-Partner* interface.

Multiple Subnets in L2 Bridged Mode

L2 Bridged Mode is capable of handling any number of subnets across the bridge, as described in L2 Bridge IP Packet Path on page 321. The default behavior is to allow all subnets, but Access Rules can be applied to control traffic as needed.

Non-IPv4 Traffic in L2 Bridged Mode

Unsupported traffic is, by default, passed from one L2 Bridge interface to the Bridge-Partner interface. This allows the Security Appliance to pass other traffic types, including LLC packets such as Spanning Tree, other EtherTypes, such as MPLS label switched packets (EtherType 0x8847), Appletalk (EtherType 0x809b), and the ever-popular Banyan Vines (EtherType 0xbad). These non-IPv4 packets are only passed across the Bridge, they are not inspected or controlled by the packet handler. If these traffic types are not needed or desired, the bridging behavior can be changed by enabling the **Block all non-IPv4 traffic** option on the **Secondary Bridge Interface** configuration dialog.

Comparison of L2 Bridged Mode to Transparent Mode

Comparison of L2 Bridged Mode to Transparent Mode

Attribute	Layer 2 Bridged Mode	Transparent Mode
Layer of Operation	Layer 2 (MAC)	Layer 3 (IP)
ARP behavior	ARP (Address Resolution Protocol) information is unaltered. MAC addresses natively traverse the L2 bridge. Packets that are destined for SonicWall Security Appliance's MAC addresses are processed, others are passed, and the source and destinations are learned and cached.	ARP is proxied by the interfaces operating in Transparent Mode.
Path determination	Hosts on either side of a Bridge-Pair are dynamically learned. There is no need to declare interface affinities.	The Primary WAN interface is always the master ingress/egress point for Transparent mode traffic, and for subnet space determination. Hosts transparently sharing this subnet space must be explicitly declared through the use of Address Object assignments.
Maximum interfaces	Two interfaces, a Primary Bridge Interface and a Secondary Bridge Interface.	Two or more interfaces. The master interface is always the Primary WAN. There can be as many transparent subordinate interfaces as there are interfaces available.
Maximum pairings	The maximum number of Bridge-Pairs allowed is limited only by available physical interfaces. This can be described as "many One-to-One pairings."	Transparent Mode only allows the Primary WAN subnet to be spanned to other interfaces, although it allows for multiple interfaces to simultaneously operate as transparent partners to the Primary WAN. This can be described as "a single One-to-One" or "a single One-to-Many pairing."
Zone restrictions	The Primary Bridge Interface can be Untrusted, Trusted, or Public. The Secondary Bridge Interface can be Trusted or Public.	Interfaces in a Transparent Mode pair must consist of one Untrusted interface (the Primary WAN, as the master of the pair's subnet) and one or more Trusted/Public interface (such as, LAN or DMZ).
Subnets supported	Any number of subnets is supported. Access Rules can be written to control traffic to/from any of the subnets as needed.	In its default configuration, Transparent Mode only supports a single subnet (that which is assigned to, and spanned from the Primary WAN). It is possible to manually add support for additional subnets through the use of ARP entries and routes.
Non-IPv4 Traffic	All non-IPv4 traffic, by default, is bridged from one Bridge-Pair interface to the Bridge-Partner interface, unless disabled on the Secondary Bridge Interface configuration page. This includes IPv6 traffic, STP (Spanning Tree Protocol), and unrecognized IP types.	Non IPv4 traffic is not handled by Transparent Mode, and is dropped and logged.

Comparison of L2 Bridged Mode to Transparent Mode

Attribute	Layer 2 Bridged Mode	Transparent Mode
VLAN traffic	VLAN traffic is passed through the L2 Bridge, and is fully inspected by the Stateful and Deep Packet Inspection engines.	VLAN subinterfaces can be created and can be given Transparent Mode Address Object assignments, but the VLANs are terminated by the Security Appliance rather than passed.
VLAN subinterfaces	VLAN subinterfaces can be configured on Bridge-Pair interfaces, but they are passed through the bridge to the Bridge-Partner unless the destination IP address in the VLAN frame matches the IP address of the VLAN subinterface on the Security Appliance, in which case it is processed (for example, as management traffic).	VLAN subinterfaces can be assigned to physical interfaces operating in Transparent Mode, but their mode of operation is independent of their parent. These VLAN subinterfaces can also be given Transparent Mode Address Object assignments, but in any event VLAN subinterfaces are terminated rather than passed.
Dynamic addressing	Although a Primary Bridge Interface may be assigned to the WAN zone, only static addressing is allowable for Primary Bridge Interfaces.	Although Transparent Mode employs the Primary WAN as a master interface, only static addressing is allowable for Transparent Mode.
VPN support	VPN operation is supported with one additional route configured. See VPN Integration with Layer 2 Bridged Mode on page 340 for details.	VPN operation is supported with no special configuration requirements.
DHCP support	DHCP can be passed through a Bridge-Pair.	Interfaces operating in Transparent Mode can provide DHCP services, or they can pass DHCP using IP Helper.
Routing and NAT	Traffic is intelligently routed in/out of the L2 Bridge-Pair from/to other paths. By default, traffic is NATed from one Bridge-Pair interface to the Bridge-Partner, but it can be NATed to other paths, as needed. Custom routes and NAT policies can be added as needed.	Traffic is intelligently routed from/to other paths. By default, traffic is not NATed from/to the WAN to/from Transparent Mode interface, but it can be NATed to other paths, as needed. Custom routes and NAT policies can be added as needed.
Stateful Packet Inspection	Full stateful packet inspection is applied to all IPv4 traffic traversing the L2 Bridge for all subnets, including VLAN traffic on firewalls.	Full stateful packet inspection is applied to traffic from/to the subnets defined by Transparent Mode Address Object assignment.
Security services	All security services (GAV, IPS, Anti-Spy, CFS) are fully supported. All regular IP traffic, as well as all 802.1Q encapsulated VLAN traffic.	All security services (GAV, IPS, Anti-Spy, CFS) are fully supported from/to the subnets defined by Transparent Mode Address Object assignment.

Comparison of L2 Bridged Mode to Transparent Mode

Attribute	Layer 2 Bridged Mode	Transparent Mode
Broadcast traffic	Broadcast traffic is passed from the receiving Bridge-Pair interface to the Bridge-Partner interface.	Broadcast traffic is dropped and logged, with the possible exception of NetBIOS which can be handled by IP Helper.
Multicast traffic	Multicast traffic is inspected and passed across L2 Bridge-Pairs providing Multicast has been activated on **MANAGE | Security Configuration |Firewall Settings > Multicast**. It is not dependent upon IGMP messaging, nor is it necessary to enable multicast support on the individual interfaces.	Multicast traffic, with IGMP dependency, is inspected and passed by Transparent Mode providing Multicast has been activated on **MANAGE | Security Configuration |Firewall Settings > Multicast**, and multicast support has been enabled on the relevant interfaces.

Benefits of Transparent Mode over L2 Bridged Mode

Two interfaces are the maximum allowed in an L2 Bridge Pair. If more than two interfaces are required to operate on the same subnet, Transparent Mode should be considered.

L2 Bridge Path Determination

Packets received by the Security Appliance on Bridge-Pair interfaces must be forwarded along to the appropriate and optimal path toward their destination, whether that path is the Bridge-Partner, some other physical or sub interface, or a VPN tunnel. Similarly, packets arriving from other paths (physical, virtual or VPN) bound for a host on a Bridge-Pair must be sent out over the correct Bridge-Pair interface.

The following summary describes, in order, the logic applied to path determinations for these cases:

1. If present, the most specific *non-default* route to the destination is chosen. This would cover, for example:
    a. A packet arriving on X3 (non-L2 Bridge LAN) destined for host `15.1.1.100` subnet, where a route to the `15.1.1.0/24` subnet exists through `192.168.0.254` through the X0 (Secondary Bridge Interface, LAN) interface. The packet would be forwarded through X0 to the destination MAC address of `192.168.0.254`, with the destination IP address `15.1.1.100`.
    b. A packet arriving on X4 (Primary Bridge Interface, LAN) destined for host `10.0.1.100`, where a route to the `10.0.1.0/24` exists through `192.168.10.50` through the X5 (DMZ) interface. The packet would be forwarded through X5 to the destination MAC address of `192.168.10.50`, with the destination IP address `10.0.1.100`.
2. If no specific route to the destination exists, an ARP cache lookup is performed for the destination IP address. A match indicates the appropriate destination interface. This would cover, for example:
    a. A packet arriving on X3 (non-L2 Bridge LAN) destined for host `192.168.0.100` (residing on L2 Primary Bridge Interface X2). The packet would be forwarded through X2 to the known destination MAC and IP address of `192.168.0.100`, as derived from the ARP cache.
    b. A packet arriving on X4 (Primary Bridge Interface, LAN) destined for host `10.0.1.10` (residing on X5 – DMZ). The packet would be forwarded through X5 to the known destination MAC and IP address of `10.0.1.10`, as derived from the ARP cache.
3. If no ARP entry is found:
    a. If the packet arrives on a Bridge-Pair interface, it is sent to the Bridge-Partner interface.
    b. If the packet arrives from some other path, the Security Appliance sends an ARP request out both interfaces of the Bridge-Pair to determine on which segment the destination IP resides.

In this last case, as the destination is unknown until after an ARP response is received, the destination zone also remains unknown until that time. This precludes the Security Appliance from being able to apply the appropriate Access Rule until after path determination is completed. Upon completion, the correct Access Rule is applied to subsequent related traffic.

With regard to address translation (NAT) of traffic arriving on an L2 Bridge-Pair interface, if it is determined to be bound for:

1 The Bridge-Partner interface, no IP translation (NAT) is performed.

2 A different path, appropriate NAT policies applies; if the path is:

   a Another connected (local) interface, there is likely no translation. That is, it is effectively routed as a result of hitting the *last-resort* Any->Original NAT Policy.

   b Determined to be through the WAN, then the default Auto-*added [interface] outbound NAT Policy for X1 WAN* applies, and the packet's source is translated for delivery to the Internet. This is common in the case of Mixed-Mode topologies as described in Internal Security on page 330.

L2 Bridge Interface Zone Selection

Bridge-Pair interface zone assignment should be done according to your network's traffic flow requirements. Unlike Transparent Mode, which imposes a system of "more trusted to less trusted" by requiring that the source interface be the Primary WAN, and the transparent interface be Trusted or Public, L2 Bridged Mode allows for greater control of operational levels of trust. Specifically, L2 Bridged Mode allows for the *Primary* and *Secondary Bridge Interfaces* to be assigned to the same or different zones (for example, LAN+LAN, LAN+DMZ, WAN+CustomLAN) This affects not only the default Access Rules that are applied to the traffic, but also the manner in which Deep Packet Inspection security services are applied to the traffic traversing the bridge. Important areas to consider when choosing and configuring interfaces to use in a Bridge-Pair are Security Services, Access Rules, and WAN connectivity:

Security Services Directionality

As it is one of the primary employments of L2 Bridged Mode, understanding the application of security services is important to the proper zone selection for Bridge-Pair interfaces. Security services applicability is based on the following criteria:

1 **The direction of the service**:

- GAV is primarily an Inbound service, inspecting inbound HTTP, FTP, IMAP, SMTP, POP3, and TCP Streams. It also has an additional Outbound element for SMTP.
- Anti Spyware is primarily Inbound, inspecting inbound HTTP, FTP, IMAP, SMTP, POP3 for the delivery (that is, retrieval) of Spyware components as generally recognized by their class IDs. It also has an additional Outbound component, where Outbound is used relative to the directionality (namely, Outgoing) ascribed to it by the IPS signatures that trigger the recognition of these Spyware components. The Outgoing classifier (described in IPS: Direction of traffic) is used because these components are generally retrieved by the client (for example, LAN host) through HTTP from a Web-server on the Internet (WAN host). Referring to IPS: Direction of traffic, that would be an *Outgoing* connection, and requires a signature with an Outgoing directional classification.
- IPS has three directions: Incoming, Outgoing, and Bidirectional. Incoming and Outgoing are described in IPS: Direction of traffic, and Bidirectional refers to all points of intersection on the table.
- For additional accuracy, other elements are also considered, such as the state of the connection (for example, SYN or Established), and the source of the packet relative to the flow (for example, initiator or responder).

2 **The direction of the traffic**. The direction of the traffic as it pertains to IPS is primarily determined by the Source and Destination zone of the traffic flow. When a packet is received by the Security Appliance, its source zone is generally immediately known, and its destination zone is quickly determined by doing a route (or VPN) lookup.

Based on the source and destination, the packet's directionality is categorized as either *Incoming* or *Outgoing,* (not to be confused with Inbound and Outbound) where the criteria shown in IPS: Direction of traffic is used to make the determination.

IPS: Direction of traffic [a]

Dest/Src	Untrusted	Public	Wireless	Encrypted	Trusted	Multicast
Untrusted	Incoming	Incoming	Incoming	Incoming	Incoming	Incoming
Public	Outgoing	Outgoing	Outgoing	Incoming	Incoming	Incoming
Wireless	Outgoing	Outgoing	Trust	Trust	Trust	Incoming
Encrypted	Outgoing	Outgoing	Trust	Trust	Trust	Outgoing
Trusted	Outgoing	Outgoing	Trust	Trust	Trust	Outgoing

a. Table data is subject to change.

In addition to this categorization, packets traveling to/from zones with levels of additional trust, which are inherently afforded heightened levels of security (LAN|Wireless|Encrypted <--> LAN|Wireless|Encrypted) are given the special *Trust* classification. Traffic with the Trust classification has all signatures applied (Incoming, Outgoing, and Bidirectional).

3 **The direction of the signature**. This pertains primarily to IPS, where each signature is assigned a direction by SonicWall's signature development team. This is done as an optimization to minimize false positives. Signature directions are:

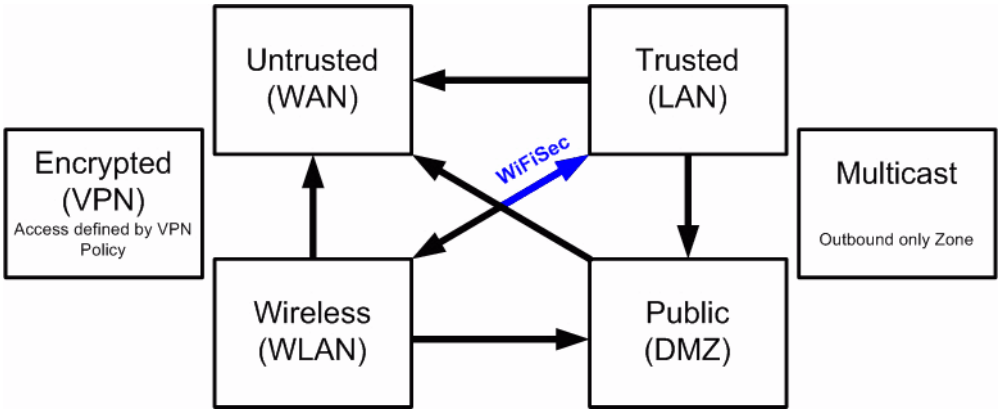
- **Incoming** – Applies to *Incoming* and *Trust*. The majority of signatures are Incoming, and they include all forms of application exploits and all enumeration and footprinting attempts. Approximately 85% of signatures are Incoming.
- **Outgoing** – Applies to *Outgoing* and *Trust*. Examples of Outgoing signatures would include IM and P2P login attempts, and responses to successfully launched exploits (for example, Attack Responses). Approximately 10% of signatures are Outgoing.
- **Bidirectional** – Applies to all. Examples of Bidirectional signatures would include IM file transfers, various NetBIOS attacks (for example, Sasser communications) and a variety of DoS attacks (for example, UDP/TCP traffic destined to port 0). Approximately 5% of signatures are Bidirectional.

4 **Zone application**. For a signature to be triggered, the desired security service *must be active on at least one of the zones it traverses*. For example, a host on the Internet (X1, WAN) accessing a Microsoft Terminal Server (on X3, Secondary Bridge Interface, LAN) triggers the *Incoming* signature "IPS Detection Alert: MISC MS Terminal server request, SID: 436, Priority: Low" if IPS is active on the *WAN*, the *LAN*, or both.

Access Rule Defaults

Default, zone-to-zone Access Rules. The default Access Rules should be considered, although they can be modified as needed. The defaults are shown in Access rule defaults:

Access rule defaults



WAN Connectivity

Internet (WAN) connectivity is required for *stack* communications, such as licensing, security services signature downloads, NTP (time synchronization), and CFS (Content Filtering Services). At present, these communications can only occur through the Primary WAN interface. If you require these types of communication, the Primary WAN should have a path to the Internet. Whether or not the Primary WAN is employed as part of a Bridge-Pair does not affect its ability to provide these stack communications.

NOTE: If Internet connectivity is not available, licensing can be performed manually and signature updates can also be performed manually (http://www.MySonicWall.com/).

Sample Topologies

The following are sample topologies depicting common deployments:

- **Inline Layer 2 Bridged Mode** represents the addition of a SonicWall Security Appliance to provide security services in a network where an existing Security Appliance is in place.
- **Perimeter Security** represents the addition of a SonicWall Security Appliance in *pure L2 Bridged Mode* to an existing network, where the Security Appliance is placed near the perimeter of the network.
- **Internal Security** represents the full integration of a SonicWall Security Appliance in *mixed-mode*, where it provides simultaneous L2 bridging, WLAN services, and NATed WAN access.
- **Layer 2 Bridged Mode with High Availability** represents the mixed-mode scenario where the Security Appliance HA pair provide high availability along with L2 bridging.
- **Layer 2 Bridged Mode with SSL VPN** represents the scenario where a SonicWall SMA SSL VPN or SonicWall SSL VPN Series appliance is deployed in conjunction with L2 Bridged Mode.

Topics:

- Wireless Layer 2 Bridge on page 329
- Inline Layer 2 Bridged Mode on page 329
- Perimeter Security on page 330
- Internal Security on page 330
- Layer 2 Bridged Mode with High Availability on page 331
- Layer 2 Bridged Mode with SSL VPN on page 332

Wireless Layer 2 Bridge

NOTE: Wireless Layer 2 Bridge does not apply to the SuperMassive 9800.

In wireless mode, after bridging the wireless (WLAN) interface to a LAN or DMZ zone, the WLAN zone becomes the secondary bridged interface, allowing wireless clients to share the same subnet and DHCP pool as their wired counterparts.

To configure a WLAN to LAN Layer 2 interface bridge:

1. Navigate to **MANAGE | System Setup | Network > Interfaces**.
2. Click the **Configure** icon for the wireless interface you wish to bridge. The **Edit Interface** dialog displays.

   TIP: If you have a Virtual Access Point configured, then you already have a VLAN interface under an interface, such as X4, in the WLAN zone, and your Virtual Access Point is configured to use that VLAN ID.

3. From **Layer 2 Bridged Mode**, select **Mode / IP Assignment**.

   NOTE: Although a general rule is automatically created to allow traffic between the WLAN zone and your chosen bridged interface, WLAN zone type security properties still apply. Any specific rules must be manually added.

4. Select the Interface to which the WLAN should be bridged from **Bridged To**. In this instance, the X0 (default LAN zone) is chosen.
5. Configure the remaining options normally. For more information on configuring WLAN interfaces, see Configuring Wireless Interfaces on page 286.

Inline Layer 2 Bridged Mode

This method is useful in networks where there is an existing Security Appliance that remains in place, but you wish to utilize the Security Appliance's security services without making major changes to the network. By placing the Security Appliance in Layer 2 Bridged Mode, the X0 and X1 interfaces become part of the same broadcast domain/network (that of the X1 WAN interface).

This example refers to a SonicWall Security Appliance installed in a Hewlett Packard ProCurve switching environment. SonicWall is a member of HP's ProCurve Alliance – more details can be found at the following location: http://www.procurve.com/alliance/members/SonicWall.htm.

HP's ProCurve Manager Plus (PCM+) and HP Network Immunity Manager (NIM) server software packages can be used to manage the switches as well as some aspects of the SonicWall Security Appliance.

To configure inline Layer 2 bridged mode:

1. Navigate to **MANAGE | System Setup | Network > Interfaces**.
2. Click the **Configure** icon for the **X0 LAN** interface.
3. On the **Edit Interface** dialog, set the IP Assignment to **Layer 2 Bridged Mode (IP Route Option)**. The options change.
4. Set the **Bridged To:** interface to **X1**.
5. To block all non-IP traffic on the bridged pair, select **Block all non-IP traffic**. This option is not selected by default.
6. To prevent traffic from being routed on the bridged pair, select **Never route traffic on this bridge-pair**. This option is not selected by default.
7. To only sniff traffic on the bridged pair, select **Only sniff traffic on this bridge-pair**. This option is not selected by default.

8 To prevent stateful inspection on the bridged pair, select **Disable stateful-inspection on this bridge-pair**. This option is not selected by default.

9 Ensure the interface is configured for **HTTPS** and **SNMP** so it can be managed from the DMZ by **PCM+/NIM**.

10 Configure the remaining options normally.

11 Click **OK** to save and activate the change.

You also need to make sure to modify the Access Rules to allow traffic from the LAN to WAN, and from the WAN to the LAN, otherwise traffic cannot pass successfully. You might also need to modify routing information on your firewall if your PCM+/NIM server is placed on the DMZ.

Perimeter Security

Perimeter Security is a network scenario where the Security Appliance is added to the perimeter to provide security services (the network may or may not have an existing Security Appliance between the Security Appliance and the router). In this scenario, everything below the Security Appliance (the *Primary Bridge Interface* segment) is generally considered as having a lower level of trust than everything to the left of the Security Appliance (the *Secondary Bridge Interface* segment). For that reason, it would be appropriate to use X1 (Primary WAN) as the *Primary Bridge Interface*.

Traffic from hosts connected to the *Secondary Bridge Interface* (LAN) would be permitted outbound through the firewall to their gateways (VLAN interfaces on the L3 switch and then through the router), while traffic from the *Primary Bridge Interface* (WAN) would, by default, not be permitted inbound.

If there are public servers, for example, a mail and Web server, on the *Secondary Bridge Interface* (LAN) segment, an Access Rule allowing WAN -> LAN traffic for the appropriate IP addresses and services could be added to allow inbound traffic to those servers.

Internal Security

A network scenario where the Security Appliance acts as the perimeter security device and secure wireless platform. Simultaneously, it provides L2 Bridge security between the workstation and server segments of the network *without having to readdress any of the workstation or servers*.

This typical inter-departmental Mixed Mode topology deployment demonstrates how the Security Appliance can simultaneously Bridge and route/NAT. Traffic to/from the *Primary Bridge Interface* (Server) segment from/to the *Secondary Bridge Interface* (Workstation) segment pass through the L2 Bridge.

As both interfaces of the Bridge-Pair are assigned to a Trusted (LAN) zone, the following apply:

- All traffic is allowed by default, but Access Rules could be constructed as needed.

  Consider, for the point of contrast, what would occur if the X2 (Primary Bridge Interface) was instead assigned to a Public (DMZ) zone: All the Workstations would be able to reach the Servers, but the Servers would not be able to initiate communications to the Workstations. While this would probably support the traffic flow requirements (that is, Workstations initiating sessions to Servers), it would have two undesirable effects:

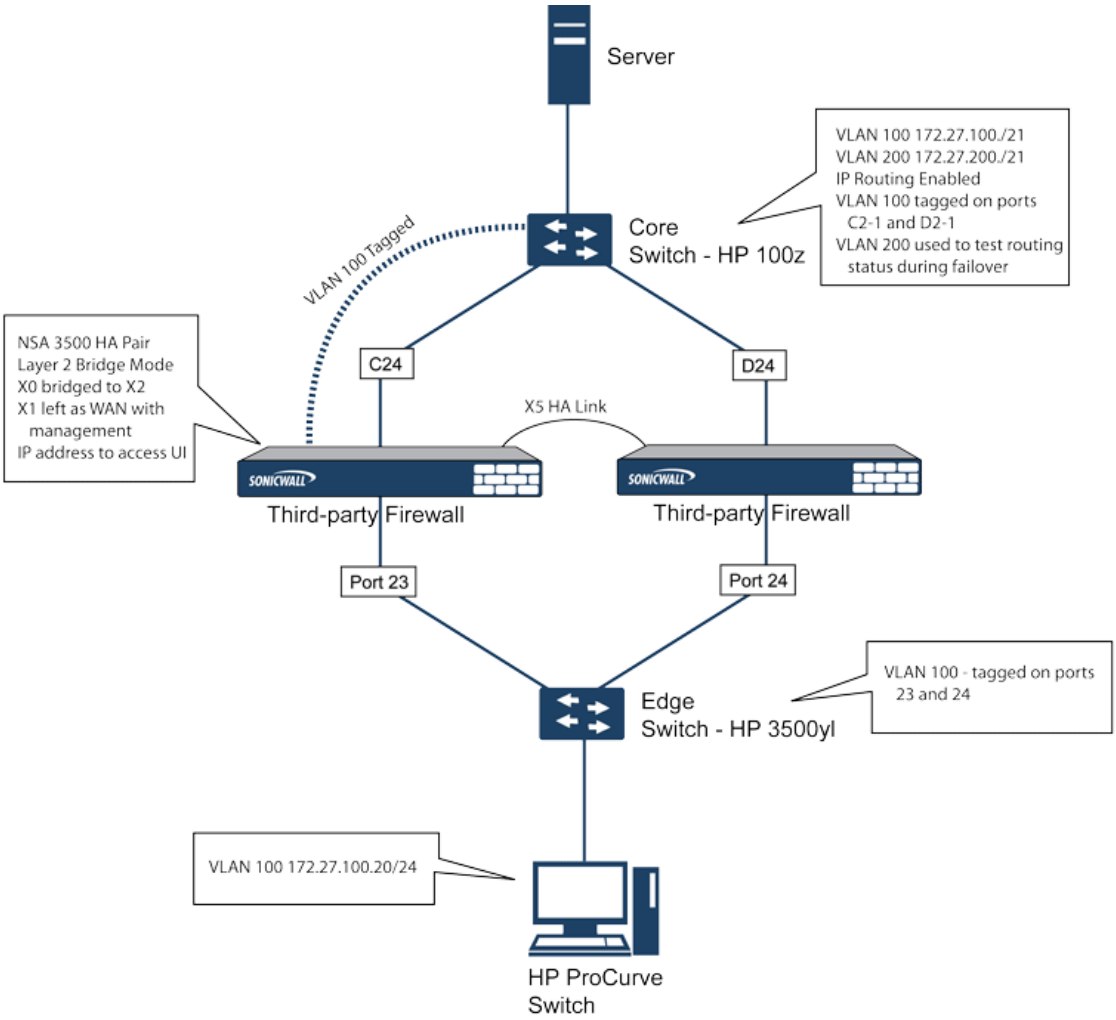
- The DHCP server would be in the DMZ. DHCP requests from the Workstations would pass through the L2 Bridge to the DHCP server (192.168.0.100), but the DHCP offers from the server would be dropped by the default DMZ->LAN Deny Access Rule. An Access Rule would have to be added, or the default modified, to allow this traffic from the DMZ to the LAN.
- Security services directionality would be classified as *Outgoing* for traffic from the Workstations to the Server because the traffic would have a Trusted source zone and a Public destination zone. This might be sub-optimal because it would provide less scrutiny than the *Incoming* or (ideally) *Trust* classifications.
- Security services directionality would be classified as *Trust*, and all signatures (*Incoming*, *Outgoing*, and *Bidirectional*) are applied, providing the highest level of security to/from both segments.

For detailed instructions on configuring interfaces in Layer 2 Bridged Mode, see Configuring Layer 2 Bridged Mode on page 334

Layer 2 Bridged Mode with High Availability

This method is appropriate in networks where both High Availability (HA) and Layer 2 Bridged Mode are desired. This example is for SonicWall Security Appliances, and assumes the use of switches with VLANs configured. See Internal security example: Both High Availability and Layer 2 Bridged Mode are desired.

Internal security example: Both High Availability and Layer 2 Bridged Mode are desired



The Security Appliance HA pair consists of two Security Appliances, connected together on port X5, the designated HA port. Port X1 on each appliance is configured for normal WAN connectivity and is used for access to the management interface of that device. Layer 2 Bridged Mode is implemented with port X0 bridged to port X2.

When setting up this scenario, there are several things to take note of on both the Security Appliances and the switches.

On the Security Appliances:

- Do not enable the Virtual MAC option when configuring High Availability. In a Layer 2 Bridged Mode configuration, this function is not useful.

- Enabling Preempt Mode is not recommended in an inline environment such as this. If Preempt Mode is required, follow the recommendations in the documentation for your switches, as the trigger and failover time values play a key role here.
- Consider reserving an interface for the management network (this example uses X1). If it is necessary to assign IP addresses to the bridge interfaces for probe purposes or other reasons, SonicWall recommends using the management VLAN network assigned to the switches for security and administrative purposes.

  NOTE: The IP addresses assigned for HA purposes do not directly interact with the actual traffic flow.

On the switches:

- Using multiple tag ports: As shown in Internal security example: Both High Availability and Layer 2 Bridged Mode are desired, two tag (802.1q) ports were created for VLAN 100 on both the Edge switch (ports 23 and 24) and Core switch (C24 - D24). The Security Appliances are connected inline between these two switches. In a high-performance environment, it is usually recommended to have Link Aggregation/ Port Trunking, Dynamic LACP, or even a completely separate link designated for such a deployment (using OSPF), and the fault tolerance of each of the switches must be considered. Consult your switch documentation for more information.
- On HP ProCurve switches, when two ports are tagged in the same VLAN, the port group is automatically placed into a failover configuration. In this case, as soon as one port fails, the other one becomes active.

Layer 2 Bridged Mode with SSL VPN

This sample topology covers the proper installation of a SonicWall Security Appliance into your existing SonicWall EX-Series SSL VPN or SonicWall SSL VPN networking environment. By placing the Security Appliance into Layer 2 Bridged Mode, with an internal, private connection to the SSL VPN appliance, you can scan for viruses, spyware, and intrusions in both directions. In this scenario the Security Appliance is not used for security enforcement, but instead for bidirectional scanning, blocking viruses and spyware, and stopping intrusion attempts. When programmed correctly, the Security Appliance does not interrupt network traffic, unless the behavior or content of the traffic is determined to be undesirable. Both one- and two-port deployments of the SonicWall Security Appliance are covered in this section.

WAN to LAN Access Rules

Because the Security Appliance is used in this deployment scenario only as an enforcement point for anti-virus, anti-spyware, and intrusion prevention, its existing security policy must be modified to allow traffic to pass in both directions between the WAN and LAN. For information about allowing traffic to pass in both directions between WAN and LAN, see *SonicOS Policies*.

Configure the Network Interfaces and Activate L2B Mode

In this scenario, the WAN interface is used for:

- Access to the management interface for the administrator
- Subscription service updates on MySonicWall
- The default route for the device and subsequently the “next hop” for the internal traffic of the SSL VPN appliance (this is why the WAN interface must be on the same IP segment as the internal interface of the SSL VPN appliance)

The LAN interface on the Security Appliance is used to monitor the unencrypted client traffic coming from the external interface of the SSL VPN appliance. This is the reason for running in Layer 2 Bridged Mode (instead of reconfiguring the external interface of the SSL VPN appliance to see the LAN interface as the default route).

To activate L2B mode on an interface:

1. Navigate to **MANAGE | System Setup > Network > Interfaces**.
2. Click the **Configure** icon for the **WAN** interface. The **Edit Interface** dialog displays.
3. Assign the interface an address that can access the Internet so that the Security Appliance can obtain signature updates and communicate with NTP. The gateway and internal/external DNS address settings must match those of your SSL VPN appliance:
   - **IP address**: This must match the address for the internal interface on the SSL VPN appliance.
   - **Subnet Mask**, **Default Gateway**, and **DNS Server(s)**: Make these addresses match your SSL VPN appliance settings.
4. For the **Management** setting, choose **HTTPS** and **Ping**.
5. Click **OK** to save and activate the changes.

To configure the LAN interface settings:

1. Navigate to **MANAGE | System Setup | Network > Interfaces**.
2. Click the **Configure** icon for the **LAN** interface.
3. For the IP Assignment setting, select **Layer 2 Bridged Mode**.
4. For the **Bridged to** setting, select **X1**.
5. If you also need to pass VLAN tagged traffic, supported on Security Appliance, click **VLAN Filtering**.
6. Add all of the VLANs that need to be passed.
7. Click **OK** to save and activate the change.

You may be automatically disconnected from the Security Appliance's management interface. You can now disconnect your management laptop or desktop from the Security Appliance's X0 interface, and power the Security Appliance off before physically connecting it to your network.

Install the Security Appliance between the Network and SSL VPN Appliance

Regardless of your deployment method (single- or dual-homed), the Security Appliance should be placed between the X0/LAN interface of the SSL VPN appliance and the connection to your internal network. This allows the device to connect out to SonicWall's licensing and signature update servers, and to scan the decrypted traffic from external clients requesting access to internal network resources.

If your SSL VPN appliance is in two-port mode behind a third-party firewall, it is dual-homed.

To connect a dual-homed SSL VPN appliance:

1. Cable the X0/LAN port on the Security Appliance to the X0/LAN port on the SSL VPN appliance.
2. Cable the X1/WAN port on the Security Appliance to the port where the SSL VPN was previously connected.
3. Power on the Security Appliance.

If your SSL VPN appliance is in one-port mode in the DMZ of a third-party firewall, it is single-homed.

To connect a single-homed SSL VPN appliance:

1. Cable the X0/LAN port on the Security Appliance to the X0/LAN port of the SSL VPN appliance.
2. Cable the X1/WAN port on the Security Appliance to the port where the SSL VPN was previously connected.

3 Power on the Security Appliance.

Configure or Verify Settings

From a management station inside your network, you should now be able to access the management interface on the Security Appliance using its WAN IP address.

To configure or verify settings:

1 Ensure that all security services for the SonicWall Security Appliance are enabled. See Licensing Services on page 335 and Activating Security Services on Each Zone on page 336.

2 SonicWall Content Filtering Service must be disabled before the device is deployed in conjunction with a SonicWall SMA SSL VPN appliance.

   a Navigate to **MANAGE | System Setup | Network > Zones** page.

   b Click **Configure** next to the **LAN (X0)** zone.

   c Clear **Enforce Content Filtering Service**.

   d Click **OK**.

3 If you have not yet changed the administrative password on the SonicWall Security Appliance, you can do so on **MANAGE | System Setup | Appliance > Base Settings**.

4 To test access to your network from an external client, connect to the SSL VPN appliance and log in.

5 When connected, attempt to access to your internal network resources. If there are any problems, review your configuration and see Configuring the Common Settings for L2 Bridged Mode Deployments on page 335.

Configuring Layer 2 Bridged Mode

Topics:

- Configuration Task List for Layer 2 Bridged Mode on page 334
- Configuring Layer 2 Bridged Mode Procedure on page 337
- VLAN Integration with Layer 2 Bridged Mode on page 339
- VPN Integration with Layer 2 Bridged Mode on page 340

Configuration Task List for Layer 2 Bridged Mode

- Choose a topology that suits your network
- Configuring the Common Settings for L2 Bridged Mode Deployments on page 335
  - License security services
  - Disable DHCP server
  - Configure and enable SNMP and HTTP/HTTPS management
  - Enable syslog
  - Activate security services on affected zones
  - Create Access Rules
  - Configure log settings

    - Configure wireless zone settings
- Configuring the Primary Bridge Interface on page 337
    - Select the zone for the Primary Bridge Interface
    - Activate management
    - Activate security services
- Configuring the Secondary Bridge Interface on page 338
    - Select the zone for the Secondary Bridge Interface
    - Activate management
    - Activate security services
- Apply security services to the appropriate zones

Configuring the Common Settings for L2 Bridged Mode Deployments

The following settings need to be configured on your SonicWall Security Appliance before using it in most of the Layer 2 Bridged Mode topologies:

- Licensing Services on page 335
- Disabling DHCP Server on page 335
- Configuring SNMP Settings on page 336
- Enabling SNMP and HTTPS on the Interfaces on page 336
- Enabling Syslog on page 336
- Activating Security Services on Each Zone on page 336
- Creating Access Rules on page 336
- Configuring Log Settings on page 337
- Configuring Wireless Zone Settings on page 337

Licensing Services

When the Security Appliance is successfully registered:

1 Navigate to **MANAGE | Updates | Licenses**.

2 Click **SYNCHRONIZE** under **Manage Security Services Online**.

This contacts the Security Appliance licensing server and ensures the Security Appliance is properly licensed.

To check licensing status, go to **MONITOR | Current Status | System Status** page and view the license status of all the security services (Gateway Anti-Virus, Anti-Spyware, and Intrusion Prevention).

Disabling DHCP Server

When using a SonicWall Security Appliance in Layer 2 Bridged Mode in a network configuration where another device is acting as the DHCP server, you must first disable the Security Appliance's internal DHCP engine, which is configured and running by default.

To disable the DHCP server:

1. Navigate to **MANAGE | System Setup | Network > DHCP Server**.
2. Clear **Enable DHCP Server**.
3. Click **ACCEPT**.

Configuring SNMP Settings

To configure SNMP settings:

1. Navigate to **MANAGE | System Setup | Appliance | SNMP**.
2. Select **Enable SNMP**.
3. Click **ACCEPT**. **Configure** becomes active and the SNMP information is populated.
4. Click **CONFIGURE**. The **Configure SNMP** dialog displays. For how to configure SNMP, see Setting Up SNMP Access on page 51.

Enabling SNMP and HTTPS on the Interfaces

To enable SNMP and HTTPS on the interfaces:

1. Navigate to **MANAGE | System Setup | Network > Interfaces**.
2. Click the **Edit** icon for the interface through which you manage the appliance. The **Edit Interface** dialog displays.
3. For the **Management** option, enable **HTTPS** and **SNMP**.
4. Click **OK**.

Enabling Syslog

You enable Syslog on the L**og > Syslog** page. For how to enable Syslog, see *SonicOS Logs and Reporting*.

Activating Security Services on Each Zone

On **MANAGE | System Setup | Network > Zones**, for each zone you are using, make sure that the security services are activated.

Then, for each service on **MANAGE | Security Configuration | Security Services**, activate and configure the settings that are most appropriate for your environment. For information about activating and configuring security services, see *SonicOS Security Configuration*.

Creating Access Rules

If you plan to manage the Security Appliance from a different zone, or if you are using a third-party server for management, SNMP, or syslog services, create Access Rules for traffic between the zones. On **MANAGE | Policies | Rules > Access Rules**, click on the icon for the intersection of the zone of the server and the zone that has users and servers (your environment could have more than one of these intersections). Create a new rule to

allow the server to communicate with all devices in that zone. For information about Access Rules, see *SonicOS Policies*.

Configuring Log Settings

On **MANAGE | Logs & Reporting | Log Settings |Name Resolution**, set the **Name Resolution Method** to **DNS then NetBios**. For information about configuring log settings, see *SonicOS Logs and Reporting*.

Configuring Wireless Zone Settings

When you are using a HP PCM+/NIM system, if it is managing a HP ProCurve switch on an interface assigned to a WLAN/Wireless zone, you need to deactivate two features; otherwise, you cannot manage the switch.

To configure wireless zone settings:

1. Navigate to **MANAGE | System Setup | Network > Zones**.
2. Select your Wireless zone.
3. On **Wireless**, clear the **Only allow traffic generated by a SonicPoint and WiFiSec Enforcement** option.
4. Click **OK**.

Configuring Layer 2 Bridged Mode Procedure

Refer to the L2 Bridge Interface Zone Selection on page 326 for choosing a topology that best suits your network. This example uses a topology that most closely resembles the Simple L2 Bridge Topology.

Choose an interface to act as the Primary Bridge Interface. Refer to the L2 Bridge Interface Zone Selection on page 326 for information in making this selection. This example uses X1 (automatically assigned to the Primary WAN):

Topics:

- Configuring the Primary Bridge Interface on page 337
- Configuring the Secondary Bridge Interface on page 338
- Configuring an L2 Bypass for Hardware Failures on page 338

Configuring the Primary Bridge Interface

To configure the primary bridge interface:

1. Navigate to **MANAGE | System Setup | Network > Interfaces**.
2. Click the **Configure** icon in the right column of the X1 (WAN) interface.
3. Configure the interface with a Static IP address (for example, `192.168.0.12`).

   NOTE: The Primary Bridge Interface must have a Static IP assignment.

4. For WAN interfaces only:
   a. Configure the default gateway. This is required for the Security Appliance itself to reach the Internet.
   b. Configure the DNS server.

5 Choose one or more **Management** options for the interface: **HTTPS**, **Ping** (selected by default), **SNMP**, **SSH**.

> **NOTE:** Selecting **HTTPS** activates and selects **Add rule to enable redirect from HTTP to HTTPS automatically**. For more information about HTTP/HTTPS redirection, see HTTP/HTTPS Redirection on page 267.

6 Choose **User Login** options: **HTTP**, **HTTPS**.

7 To enable redirect to HTTPS from HTTP, select **Add rule to enable redirect from HTTP to HTTPS**. For more information about this option, see HTTP/HTTPS Redirection on page 267.

8 Click **OK**.

Choose an interface to act as the Secondary Bridge Interface. Refer to the L2 Bridge Interface Zone Selection on page 326 for information in making this selection.

Configuring the Secondary Bridge Interface

This example uses X0 (automatically assigned to the LAN):

1 Navigate to **MANAGE | System Setup | Network > Interfaces**.

2 Click the **Configure** icon in the right column of the X0 (LAN) interface.

3 From **IP Assignment**, select **Layer 2 Bridged Mode**.

4 From **Bridged to**, select the **X1** interface.

5 Choose one or more **Management** options for the interface: **HTTPS**, **Ping** (selected by default), **SNMP**, **SSH**.

> **NOTE:** Selecting **HTTPS** activates and selects **Add rule to enable redirect from HTTP to HTTPS automatically**. For more information about HTTP/HTTPS redirection, see HTTP/HTTPS Redirection on page 267.

6 Choose **User Login** options: **HTTP**, **HTTPS**.

7 To enable redirect to HTTPS from HTTP, select **Add rule to enable redirect from HTTP to HTTPS**. For more information about this option, see HTTP/HTTPS Redirection on page 267.

8 You may optionally enable **Block all non-IPv4 traffic** to prevent the L2 bridge from passing non-IPv4 traffic.

9 To control VLAN traffic through the L2 bridge, click **VLAN Filtering**. By default, all VLANs are allowed:

- Select **Block listed VLANs (blacklist)** from the drop-down list and add the VLANs you wish to block from the left pane to the right pane. All VLANs added to the right pane is blocked, and all VLANs remaining in the left pane are allowed.
- Select **Allow listed VLANs (whitelist)** from the drop-down list and add the VLANs you wish to explicitly allow from the left pane to the right pane. All VLANs added to the right pane are allowed, and all VLANs remaining in the left pane are blocked.

10 Click **OK**. The **Interface Settings** table displays the updated configuration:

You may now apply security services to the appropriate zones, as desired. In this example, they should be applied to the LAN, WAN, or both zones.

Configuring an L2 Bypass for Hardware Failures

An L2 bypass enables you to perform a physical bypass of the Security Appliance when an interface is bridged to another interface with LAN bypass capability. This allows network traffic to continue flowing if an unrecoverable firewall error occurs.

When the L2 bypass relay is closed, the network cables attached to the bypassed interfaces (X0 and X1) are physically connected as if they were a single continuous network cable. The **Engage physical bypass on malfunction** option provides the user the choice of avoiding disruption of network traffic by bypassing the firewall in the event of a malfunction.

L2 bypass is only applicable to interfaces in **Layer 2 Bridged Mode**. The **Engage physical bypass on malfunction** option only appears when the **Layer 2 Bridged Mode** option is selected from **Mode / IP Assignment**. This option does not appear unless a physical bypass relay exists between the two interfaces of the bridge-pair.

When the **Engage physical bypass on malfunction** option is enabled, the other **Layer 2 Bridged Mode** options are automatically set

- **Block all non-IPv4 traffic** – disabled. When enabled, this option blocks all non-IPv4 Ethernet frames. So, this option is disabled.
- **Never route traffic on this bridge-pair** – enabled. When enabled, this option prevents packets from being routed to a network other than the peer network of the bridged pair. So, this option is enabled.
- **Only sniff traffic on this bridge-pair** – disabled. When enabled, traffic received on the bridge-pair interface is never forwarded. So, this option is disabled.
- **Disable stateful-inspection on this bridge-pair** – unchanged. This option is not affected.

To configure an L2 bypass:

1 Navigate to **MANAGE | System Setup | Network > Interfaces**.

2 Click on the **Edit** icon in the **Configure** column for the interface you want to configure. The **Edit Interface** dialog displays.

3 Select **Engage physical bypass on malfunction.**

NOTE: The **Engage physical bypass on malfunction** option is available only when the X0 and X1 interfaces are bridged together on an NSA-6600 or above.

4 Click **OK**.

VLAN Integration with Layer 2 Bridged Mode

VLANs are supported on SonicWall Security Appliances. When a packet with a VLAN tag arrives on a physical interface, the VLAN ID is evaluated to determine if it is supported. The VLAN tag is stripped, and packet processing continues as it would for any other traffic. A simplified view of the inbound and outbound packet path includes these potentially reiterative steps:

- IP validation and reassembly
- Decapsulation (802.1q, PPP)
- Decryption
- Connection cache lookup and management
- Route policy lookup
- NAT Policy lookup
- Access Rule (policy) lookup
- Bandwidth management
- NAT translation
- Advanced Packet Handling (as applicable)
    - TCP validation

- Management traffic handling
- Content Filtering
- Transformations and flow analysis (on SonicWall Security Appliances): H.323, SIP, RTSP, ILS/LDAP, FTP, Oracle, NetBIOS, Real Audio, TFTP
- IPS and GAV

At this point, if the packet has been validated as acceptable traffic, it is forwarded to its destination. The packet egress path includes:

- Encryption
- Encapsulation
- IP fragmentation

On egress, if the route policy lookup determines that the gateway interface is a VLAN subinterface, the packet is tagged (encapsulated) with the appropriate VLAN ID header. The creation of VLAN subinterfaces automatically updates the firewall's routing policy table:

The auto-creation of NAT policies, Access Rules with regard to VLAN subinterfaces behave exactly the same as with physical interfaces. Customization of the rules and policies that govern the traffic between VLANs can be performed with customary SonicOS ease and efficiency.

When creating a zone (either as part of general administration, or as a step in creating a subinterface), a checkbox is presented on the zone creation page to control the auto-creation of a GroupVPN for that zone. By default, only newly created Wireless type zones have **Create GroupVPN for this zone** enabled, although the option can be enabled for other zone types by selecting the checkbox during creation.

Management of security services between VLAN subinterfaces is accomplished at the zone level. All security services are configurable and applicable to zones comprising physical interfaces, VLAN subinterfaces, or combinations of physical and VLAN subinterfaces.

Gateway Anti-Virus and Intrusion Prevention Services between the different workgroups can easily be employed with the use of VLAN segmentation, obviating the need for dedicated physical interfaces for each protected segment.

VLAN support enables organizations to offer meaningful internal security (as opposed to simple packet filtering) between various workgroups, and between workgroups and server farms without having to use dedicated physical interfaces on the firewall.

Here the ability to assign VLAN subinterfaces to the WAN zone, and to use the WAN client mode (only Static addressing is supported on VLAN subinterfaces assigned to the WAN zone) is illustrated, along with the ability to support WAN Load Balancing and failover. Also demonstrated is the distribution of SonicPoints throughout the network by means of connecting them to access mode VLAN ports on workgroup switches. These switches are then backhauled to the core switch, which then connects all the VLANs to the appliance through a trunk link.

VPN Integration with Layer 2 Bridged Mode

When configuring a VPN on an interface that is also configured for Layer 2 Bridged Mode, you must configure an additional route to ensure that incoming VPN traffic properly traverses the Security Appliance.

To configure VPN integration with Layer 2 bridged mode:

1. Navigate to **MANAGE | System Setup | Network > Interfaces**.
2. Click the **Add** icon. The **Add Route Policy** dialog displays.
3. Configure the route as follows:
    - Source: **ANY**

- Destination: *custom-VPN-address-object* (This is the address object for the local VPN tunnel IP address range.)
- Service: **ANY**
- Gateway: **0 . 0 . 0 . 0**
- Interface: **X0**

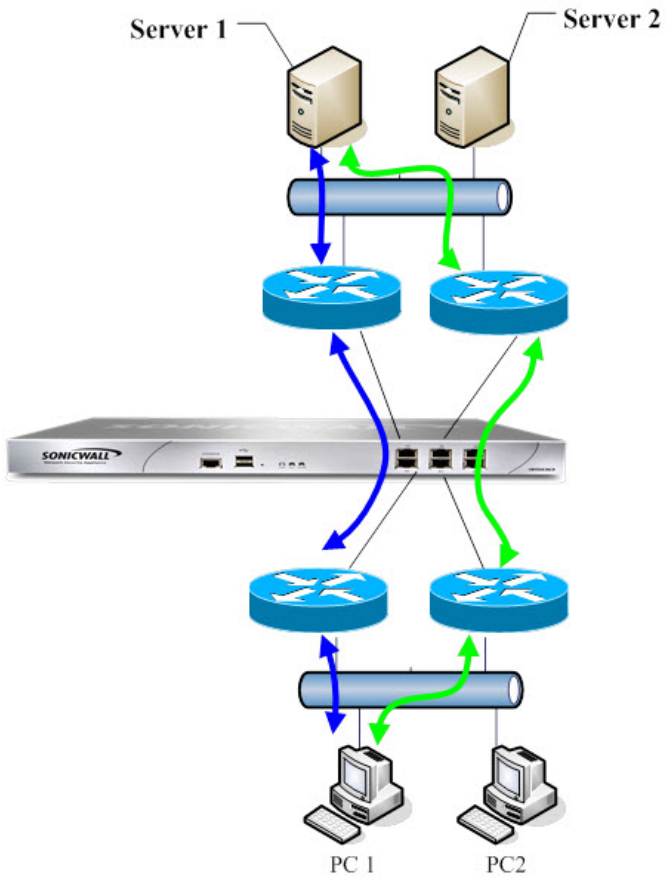
4 Click **OK**.

Asymmetric Routing

SonicOS supports asymmetric routing. Asymmetric routing is when the flow of packets in one direction passes through a different interface than that used for the return path. This can occur when traffic flows across different layer 2 bridged pair interfaces on the Security Appliance or when it flows across different Security Appliances in a high availability cluster.

Any Security Appliance that performs deep packet inspection or stateful firewall activity must “see” all packets associated with a packet flow. This is in contrast to traditional IP routing in which each packet in a flow may technically be forwarded along a different path as long as it arrives at its intended destination — the intervening routers do not have to see every packet. Today’s routers do attempt to forward packets with a consistent next-hop for each packet flow, but this applies only to packets forwarded in one direction. Routers make no attempt to direct return traffic to the originating router. This IP routing behavior presents problems for a Security Appliance cluster that does not support asymmetric routing because the set of Cluster Nodes all provide a path to the same networks. Routers forwarding packets to networks through the cluster may choose any of the Cluster Nodes as the next-hop. The result is asymmetric routing, in which the flow of packets in one direction go through a node different than that used for the return path. This difference in flow causes traffic to be dropped by one or both Cluster Nodes as neither is “seeing” all of the traffic from the flow. See Asymmetric routing.

Asymmetric routing



Asymmetric Routing Traffic

In Asymmetric routing, PC1 communicates with Server1, two-way traffic passes through different routers, that is, some packets of same connection go through blue path, some go through green path. On such deployment, the routers may run some redundancy route protocol or load balancing protocol. for example.Cisco HSRP protocol.

SonicOS uses stateful inspection. All connections passing through the Security Appliance are bound to interfaces. With support for asymmetric routing, however, SonicOS tracks ingress and egress traffic, even when the flows go across different interfaces, and provides stateful, deep packet inspection.

NOTE: Asymmetric routing is not the same as one-way connections without reply, that is, TCP State Bypass.

Configuring Interfaces for IPv6

For a complete description of configuring IPv6 interfaces, see IPv6 Interface Configuration on page 938.

31-Bit Network

SonicOS 6.2.7 introduces support for RFC 3021, which defines the use of a 31-bit subnet mask. This mask allows only two host addresses in the subnet, with no network or gateway address and no broadcast address. Such a configuration can be used within a larger network to connect two hosts with a point-to-point link. The savings in

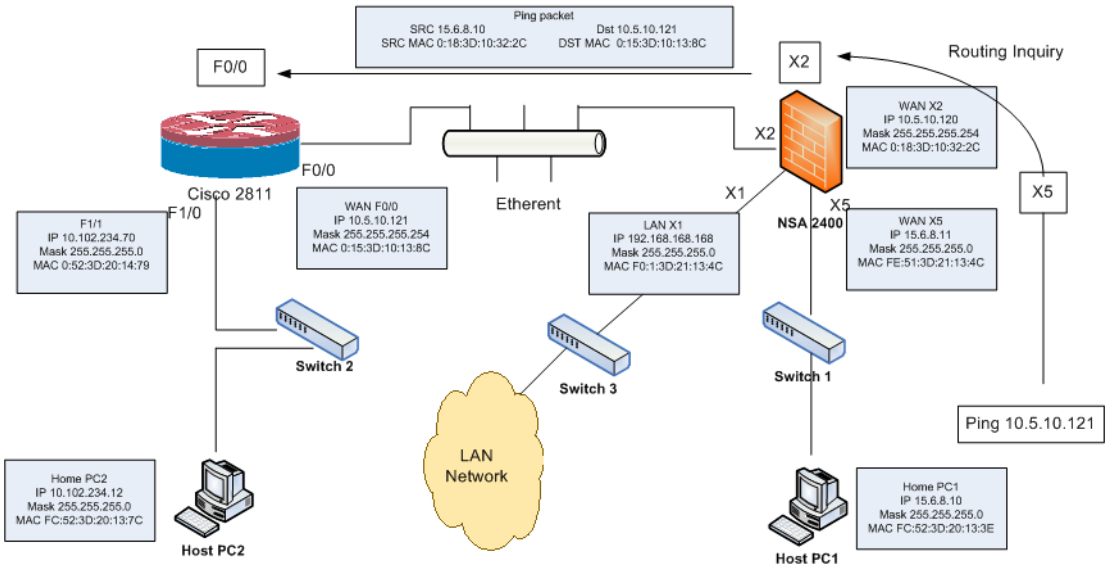
address space resulting from this change is easily seen as each point-to-point link in a large network would consume two addresses instead of four.

In this context, the point-to-point link is not equivalent to PPP (point to point protocol). A point-to-point link using a 31-bit mask can use or not use the PPP protocol. 31-bit prefixed IPv4 addresses on a point-to-point link can also be used in the Ethernet network.

Topics:

- Example Network Environment on page 343
- Configuring SonicOS on page 344

Example Network Environment



In this network environment, Host PC1 and Host PC2 can visit each other, while hosts in the LAN network can visit Host PC2.

To configure settings for this environment:

1 For Host PC1, add two route entries:

- `Route add 10.5.10.0 mask 255.255.255.0  15.6.8.10`
- `Route add 10.102.234.0 mask 255.255.255.0  15.6.8.10`

2 For Host PC2, add two route entries:

- `Route add 10.5.10.0 mask 255.255.255.0  10.102.234.70`
- `Route add 15.6.8.0 mask 255.255.255.0  10.102.234.70`

3 On the Cisco router (F0/0):

- `interface fastEthernet 0/0`
- `ip address 10.5.10.120 255.255.255.254`

4 On the Cisco 2811, add one route entry:

```
!
ip route 15.6.8.0 255.255.255.0 10.5.10.120
!
```

5 On the firewall, add one route entry to enable the WAN zone data flow from X2 to X5, and X5 to X2:

```
Any    10.102.234.0    Any    X2 Default Gateway    X2
```

Configuring SonicOS

To configure an interface for a 31-bit subnet:

1 Navigate to **MANAGE | System Setup | Network > Interfaces**.

2 Edit the desired interface.

3 Set the **Subnet Mask** to `255.255.255.254`.

4 Enter one host IP address into the **IP Address** field.

5 Enter the other host IP address into the **Default Gateway** field.

6 Set the other fields according to your network, as needed.

7 Click **OK**.

PPPoE Unnumbered Interface Support

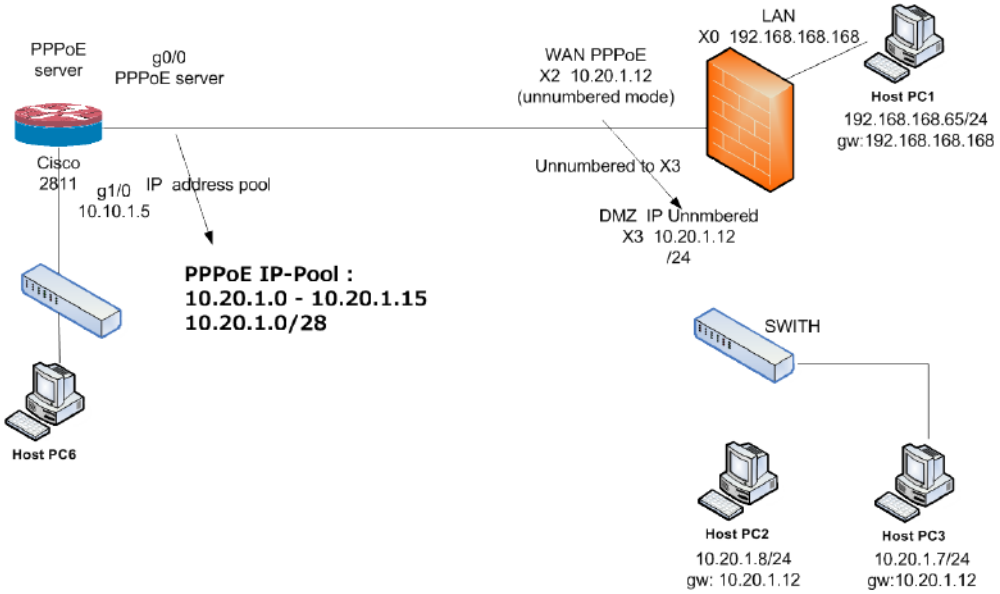
A PPPoE unnumbered interface allows you to manage a range of IP addresses with only a single PPPoE connection. The Internet Service Provider (ISP) provides multiple static IP addresses that can be allocated within a subnet. The first address is designated as the network address, and the last one as the broadcast address.

The default MTU of PPPoE is **1492**.

Topics:

- Sample Network Topography on page 345
- Caveats on page 345
- Configuring a PPPoE Unnumbered Interface on page 345
- Configuring HA with PPPoE Unnumbered on page 346

Sample Network Topography



In this topology, X2 is the PPPoE unnumbered interface, and X3 is an unnumbered interface.

SonicOS adds two policies to the **Network > Routing > Route Policies** table.

SonicOS also adds two NAT policies.

Caveats

To change X3 to another mode when X2 unnumbered to X3 is configured, first terminate the relationship with X2 by changing X2 to another mode. Otherwise, if you change the IP address or mask of interface X3, it causes X3 to reconnect to the PPPoE server.

If X3 is set as unnumbered interface, other interfaces cannot connect to X3 using an L2 Bridge.

Configuring a PPPoE Unnumbered Interface

To configure a PPPoE unnumbered interface:

1. Navigate to **MANAGE | System Setup | Network > Interfaces**.
2. Configure the PPPoE client settings on a WAN interface by clicking its **Edit** icon. The **Edit Interface** dialog displays.
3. Select **Unnumbered interface**. The drop-down menu activates.
4. Select **Create new unnumbered Interface**. The **Add Unnumbered Interface** dialog displays.
5. For **Zone**, select **LAN**, **DMZ**, or create a new zone.

   NOTE: **Mode / IP Assignment** is set to **IP Unnumbered** and dimmed.

6. For **IP Address**, enter the address provided by your ISP. Usually it is the second IP address assigned by the provider.
7. Enter the subnet mask assigned by the ISP in the **Subnet Mask** field.
8. Finish configuring this interface.

9 Click **OK**.

10 Finish configuring the first interface.

11 Click **OK**.

Configuring HA with PPPoE Unnumbered

For how to configure HA with PPPoE Unnumbered, see Configuring Active/Standby High Availability Settings on page 704.

14

Configuring PortShield Interfaces

NOTE: The NSA 2600 Security Appliance does not support PortShield, and the SOHO W Security Appliance does not support the X-Series/N-Series Solution.

- **Network > PortShield Groups** on page 347
  - About PortShield on page 347
  - SonicOS Support of X-Series/N-Series Switches on page 348
  - SonicOS Support of N-Series Switches on page 356
  - Managing Ports on page 361
  - Configuring PortShield Groups on page 370+

Network > PortShield Groups

Topics:

- About PortShield on page 347
- SonicOS Support of X-Series/N-Series Switches on page 348
- SonicOS Support of N-Series Switches on page 356
- Managing Ports on page 361
- Configuring PortShield Groups on page 370

About PortShield

A PortShield interface is a virtual interface with a set of ports, including ports on Dell X-Series/N-Series, or extended, switches, assigned to it. PortShield architecture enables you to configure some or all of the LAN ports into separate security contexts, providing protection not only from the WAN and DMZ, but between devices inside your network as well. In effect, each context has its own wire-speed PortShield that enjoys the protection of a dedicated, deep packet inspection Security Appliance.

TIP: Zones can always be applied to multiple interfaces in **MANAGE | System Setup | Network > Interfaces**, even without the use of PortShield groupings. These interfaces, however, do not share the same network subnet unless they are grouped using PortShield.

You can assign any combination of ports to a PortShield interface. All ports not assigned to a PortShield interface are assigned to the LAN interface.

Static Mode and Transparent Mode

There are two IP assignment methods you can deploy to create PortShield interfaces:

- Static mode
- Transparent mode

Working in Static Mode

When you create a PortShield interface in Static Mode, you manually create an explicit address to be applied to the PortShield interface. All ports mapped to the interface are identified by this address. Static mode is available on interfaces assigned to Trusted, Public, or Wireless zones.

NOTE: When you create a PortShield interface in Static Mode, make sure the IP address you assign to the interface is not already in use by another PortShield interface.

Working in Transparent Mode

Transparent Mode addressing allows for the WAN subnetwork to be shared by the current interface through Address Object assignments. The interface's IP address is the same as the WAN interface IP address. Transparent mode is available on interfaces assigned to Trusted and Public Zones.

NOTE: Make sure the IP address you assign to the PortShield interface is within the WAN subnetwork.

When you create a PortShield interface in Transparent Mode, you create a range of addresses to be applied to the PortShield interface. You include these addresses in one entity called an Address Object. Address Objects allow for entities to be defined one time and to be re-used in multiple referential instances throughout the SonicOS interface. When you create a PortShield interface using an address object, all ports mapped to the interface are identified by any of the addresses specified in the address range.

NOTE: Each statically addressed PortShield interface must be on a unique subnetwork. You can not overlap PortShield interfaces across multiple subnetworks.

SonicOS Support of X-Series/N-Series Switches

Topics:

- About the X-Series/N-Series Solution on page 348
- Supported Topologies on page 356

About the X-Series/N-Series Solution

NOTE: The X-Series/N-Series Solution is not supported on the NSA 2600 or SOHO W Security Appliance.

Critical network elements, such as a firewall and switch, need to be managed, usually individually. SonicOS allows unified management of the firewall and the switches using the Security Appliance Management Interface and GMS.

The maximum number of interfaces available on the SonicWall Security Appliances varies depending on the model. The feature is supported on all SonicWall firewalls running SonicOS 6.5 or SonicOS 6.2.7 and higher, except for the following:

- NSsp 12000 series
- SuperMassive 9800
- NSA 2600

In certain deployments, the number of ports required might easily exceed the maximum number of interfaces available on the firewall. With the X-Series /N-Series Solution, ports on switches are viewed as extended interfaces of the Security Appliance, thereby increasing the number of interfaces available for use up to 192, depending on the switch. These extended ports can be portshielded and/or configured for High Availability (HA) and treated as any other interface on the Security Appliance.

For complete information about X-Series/N-Series switches, refer to the following documents:

- Dell™ Networking™ X1000 and X4000 Series Switches Getting Started Guide
- Dell™ Networking™ X1000 and X4000 Series Switches User Guide
- Dell EMC Networking N-Series N1100-ON, N1500, N2000, N2100-ON, N3000, N3000-ON, N3100-ON, and N4000 Switches User's Configuration Guide
- SonicWall TZ Series and SonicWall X-Series solution managing SonicWall access points

Topics:

- Terminology on page 349
- Performance Requirements on page 350
- Key Features Supported with X-Series/N-Series Switches on page 350
- PortShield Functionality and X-Series/N-Series Switches on page 351
- Dell X-Series/N-Series Daisy-Chaining Support on page 352
- PoE/PoE+ and SFP/SFP+ Support on page 353
- X-Series/N-Series Solution and SonicPoints on page 354
- Managing Extended Switches using GMS on page 354
- Extended Switch Global Parameters on page 354
- About Links on page 355
- Logging and Syslog Support on page 355

Terminology

HA	High Availability
Extended switch	Same as X-Series/N-Series switch.
External switch	Same as X-Series/N-Series switch.
IDV	Interface Disambiguation through VLAN – The reconfiguring of ports, portshielded to Security Appliance interfaces, on the extended switch as access ports of the VLAN corresponding to the PortShield VLAN.

PoE	Power over Ethernet – A system than passes electrical power along with data on Ethernet cabling, which allows a single cable to provide both data connection and electrical power to devices.
PoE+	Power over Ethernet Plus – An enhanced version of PoE (standard 802.3at) that provides more power than PoE.
SFP	Small form-factor pluggable – A compact, hot-pluggable transceiver used for both telecommunication and data communications applications and supports 1Gb fiber modules.
SFP+	Enhanced small form-factor pluggable – An enhanced version of SFP that supports 10 Gb fiber modules.
SPM	Single Point Management
STP	Spanning Tree Protocol – A network protocol that ensures a loop-free topology for Ethernet networks and allows redundant (spare) links to provide backup paths if an active link fails.

Performance Requirements

A SonicWall Security Appliance can now:

- Be provisioned for a maximum of four switches.
- Manage an increased number of ports.

Key Features Supported with X-Series/N-Series Switches

NOTE: For more information on these features, refer to the *SonicWall SonicOS 6.5 X-Series/N-Series Solution Deployment Guide* located on the Support portal at https://www.sonicwall.com/support/technical-documentation/ and choose **NS*a* Series**, **NSA Series**, **SuperMassive 9000 Series**, and **TZ Series** in the **Select A Product** field.

- Provisioning an X-Series/N-Series switch as an extended switch
- PortShield functionality
- Configuring extended switch Interface settings
- Managing basic extended switch global parameters
- Managing the extended switch using GMS
- High Availability (HA) with PortShield functionality

  Support for PortShield functionality in HA mode is available using Common Uplink. In this configuration, a link between the active/standby Security Appliance and the switch serves as a common uplink to carry all the PortShield traffic. In this configuration, Security Appliance interfaces that serve as PortShield hosts should be connected to a separate switch and not the same switch connected to the active and standby units. This avoids looping of packets for the same PortShield VLAN. The PortShield members can be connected to ports on the switch that is controlled by the active/standby Security Appliance.
- Diagnostics support for extended switch
- Support for VLANs in a common uplink with SPM configuration
- Support for VLANs in a dedicated uplink configuration
- Single Point of Management over Common Uplink for VLAN Traffic

  VLANs are also supported with Common Uplink. This allows a single link between the Security Appliance and the switch to carry management traffic of the Security Appliance managing the switch plus

PortShield traffic for the *Interface Disambiguation via VLAN* (IDV) VLANs corresponding to the Security Appliance interfaces plus traffic for the VLAN sub-interfaces present under the Common Uplink interface.

NOTE: Overlapping VLANs cannot exist under Security Appliance interfaces configured as dedicated uplinks or common uplinks to the same switch. This is because the VLAN space is global on the switch.

NOTE: PortShield of Extended Switch Interfaces to Common Uplink Interfaces without selecting any VLANs for access/trunk configuration is not supported.

- PoE/PoE+ and SFP/SFP+ functionality for SonicWall Security Appliances by certain X-Series/N-Series switches.
- Batching configuration messages – To facilitate support of the X-Series /N-Series switches, configuration messages can be batched before being sent to a switch.

PortShield Functionality and X-Series/N-Series Switches

PortShield architecture allows configuration of Security Appliance ports into separate security zones, thereby allowing protection of a deep-packet inspection Security Appliance for traffic between devices across zones. For more information about PortShield functionality, see Configuring PortShield Interfaces on page 347.

The SonicWall X-Series/N-Series Solution allows support for portshielding interfaces on the extended switch to Security Appliance interfaces. X-Series/N-Series switches are L2 switches, and by default, all ports on the extended switch are configured as access ports of the default VLAN 1. When ports of the extended switch are portshielded to Security Appliance interfaces, the ports are reconfigured as access ports of the VLAN corresponding to the PortShield VLAN, also known as the IDV VLAN of the PortShield host interface.

Topics:

- Different Traffic Scenarios with PortShield on page 351
- Prerequisites for PortShielding X-Series/N-Series Switches on page 352
- Dell X-Series/N-Series Daisy-Chaining Support on page 352

Different Traffic Scenarios with PortShield

- Traffic between network devices connected to the ports on the extended switch that are part of the same PortShield group are switched automatically by the extended switch.
- Traffic between network devices connected to the ports on the extended switch and devices connected to ports on the Security Appliance that are part of the same PortShield group are switched by the internal switch on the Security Appliance.
- Traffic between network devices connected to the ports on the extended switch destined to Security Appliance interfaces are handled by the data path in software. Such traffic may be subjected to Security Appliance security services such as access rules, deep packet inspection, and intrusion prevention.
- Traffic between network devices connected to the ports on the extended switch and devices connected to ports on the Security Appliance that are part of a different zone or part of a different PortShield group are forwarded by the data path in software. Such traffic is subjected to Security Appliance security services in software.

Prerequisites for PortShielding X-Series/N-Series Switches

IMPORTANT: If the topology has two or more switches, the switches can be cascaded or daisy chained, that is, one switch can be connected to another one that is connected to the Security Appliance.

- X-Series switches (excluding X1052/X1052P models) are delivered from the factory in unmanaged mode to avoid unauthorized access to the switch. You need to put the switch into Managed mode by pressing **Mode**, near the power plug, for at least seven seconds.

  X1052/X1052P models delivered from the factory are by default in Managed mode.

  During the initial set up of the switch, to ensure the X-Series switch's IP does not change dynamically when the DHCP server is enabled on the Security Appliance interfaces, choose **Static IP** instead of **Dynamic IP**.

  For more information on these features, refer to the *SonicWall SonicOS 6.5 X-Series/N-Series Solution Deployment Guide* located on the Support portal at https://www.sonicwall.com/support/technical-documentation/ and select **NS*a* Series**, **NSA Series**, **SuperMassive 9000 Series**, and **TZ Series** in the **Select A Product** field.

- Apart from the initial IP address, username/password configuration, which can be found on the switch, no other configuration is recommended to be performed on the X-Series switch directly through the switch's GUI/console. To do so results in the Security Appliance being out-of-sync with the configuration state of the X-Series switch.

- To manage the X-Series switch from the Security Appliance, one of the interfaces of the Security Appliance must be in the same subnet as the X-Series switch. For example, to manage an X-Series switch with a default IP `192.168.2.1`, an interface of the Security Appliance needs to be configured in the `192.168.2.0/24` subnet and connected to the X-Series switch.

- Ensure the Security Appliance can reach the X-Series switch by pinging the X-Series switch from the Security Appliance before provisioning/managing the switch from the Security Appliance.

- VLAN support:

  - Support for VLANs is available on shared and common uplinks. For example, VLANs can be configured under the Security Appliance interface, which is provisioned as the shared uplink for the X-Series switch.

  - For details on VLAN support, refer to the *SonicWall SonicOS 6.5 X-Series/N-Series Solution Deployment Guide* located on the Support portal at https://www.sonicwall.com/support/technical-documentation/ and select **NS*a* Series**, **NSA Series**, **SuperMassive 9000 Series**, and **TZ Series** in the **Select A Product** field. Overlapping VLANs cannot exist under Security Appliance interfaces configured as dedicated uplinks. For example, if X3 and X5 are configured for dedicated uplinks, VLAN 100 cannot be present under both X3 and X5. Such a configuration is rejected.

Dell X-Series/N-Series Daisy-Chaining Support

NOTE: This feature is not supported on NSA 2600 platforms.

The SonicOS X-Switch/N-Switch Daisy Chaining solution enables integration of a SonicWall Security Appliance with switches connected in daisy-chained mode. Integration with all Dell switch models is supported in daisy-chain mode.

Daisy Chaining allows those with large facilities, such as warehouses, to deploy two switches more than 1000 ft apart on a given site, to be connected to each other through fiber, to have the first switch—the parent switch—connected to the Security Appliance, and to manage both the switches from the Security Appliance. This deployment also allows you access to an increased number of interfaces on the switch by using a single interface on the Security Appliance. All the interfaces of the parent switch and the child switch are available to be managed from the Security Appliance.

Topics:

- Assumptions and Dependencies on page 353
- Daisy Chaining Support on page 353

Assumptions and Dependencies

- SonicOS switch daisy-chaining solution allows support for single level of chaining only. Multi-level chaining, where more than two switches are connected in series, is not supported. For example, the parent switch can be connected to a child switch, but the child switch cannot be connected to another child switch.
- There is a maximum limitation of 4 extended switches that can be provisioned. For example, a parent switch can have up to three child switches.
- In daisy-chaining mode, the only supported topology for the child switch is Common Uplink in which the child switch is connected to the parent switch through a single uplink. Other variations, such as dedicated uplinks, isolated links, etc. are not supported for the child switch.

Daisy Chaining Support

Both switches connected in daisy chained mode must have the IP address in the same subnet, and the Security Appliance must be able to reach this subnet. Provisioning the switches in daisy-chained mode is a two-step process:

1 Provision the parent switch as a standalone switch.

2 Provision the child switch as a daisy-chained switch.

PoE/PoE+ and SFP/SFP+ Support

Some SonicWall Security Appliances do not support PoE/PoE+, but this functionality can be added with certain switches, as shown in X-Series switch PoE/PoE+ and SFP/SFP+ support. This additional functionality enhances SonicWave/SonicPoint usage by SonicWall Security Appliances, especially for wireless access points supporting 802.11ac (supports up to 30W maximum power; 802.11a/b/g/h supports up to 15.4 W maximum power).

Some X-Series switches also support SFP/SFP+, as shown in X-Series switch PoE/PoE+ and SFP/SFP+ support.

NOTE: Configuration of the PoE/PoE+ ports on the X-Series switch is managed from the UI of the X-Series switch and not **MANAGE | System Setup** | **Network > PortShield Groups** on the SonicWall Security Appliance.

X-Series switch PoE/PoE+ and SFP/SFP+ support

This X-Series switch	Supports
X1008	1 PoE PD port; by default, port 8 is the PD port
X1008P	8 PoE ports, up to 123W total; by default, ports 1 through 8 support PoE
X1018	2 1GbE SFP ports; by default, ports 17 and 18 support SFP
X1018P	16 PoE ports, up to 246W total; by default, ports 1 through 16 support PoE 2 1GbE SFP ports; by default, ports 17 and 18 support SFP
X1026	2 1GbE SFP ports; by default, ports 25 and 26 support SFP
X1026P	24 PoE/12 PoE+ ports, up to 369W total; by default: • Ports 1 through 12 support PoE+ • Ports 13 through 24 support PoE 2 1GbE SFP ports; by default, ports 25 and 26 support SFP

X-Series switch PoE/PoE+ and SFP/SFP+ support (Continued)

This X-Series switch	Supports
X1052	4 10GbE SFP+ ports; by default, ports 49 through 52 support SFP+
X1052P	24 PoE/12 PoE+ ports, up to 369W total; by default: • Ports 1 through 12 support PoE+ • Ports 13 through 24 support PoE • Ports 25 through 48 support neither PoE nor PoE+ 4 10GbE SFP+ ports; by default, ports 49 through 52 support SFP+
X4012	12 10GbE SFP+ ports; by default, ports 1 through 12 support SFP+

IMPORTANT: A SonicWave or SonicPoint AC without an external power source must be portshielded through ports 1 through 12 on an X1026P or X1052P X-Series switch.

Any SonicPoint non-AC model without an external power source can be portshielded through ports 1 through 8 (X1008P), 1 through 16 (X1018P), or 1 through 24 (X1026P and X1052P).

Any SonicPoint with an external power source can be portshielded to any ethernet port.

X-Series/N-Series Solution and SonicPoints

Ports on an extended switch can be portshielded to the WLAN zone of the Security Appliance, and SonicWall wireless access points can be connected to these ports.

When connecting access points to an X-Series/N-Series switch, it is important to consider the access point's power requirements. A SonicPointACe/ACi/N2 requires a minimum of 25.5 watts. If your switch model does not support PoE+, you must use a SonicPoint power injector. For which switches support PoE+, see PoE/PoE+ and SFP/SFP+ Support on page 353. For more information about managing SonicWall access points, see the Knowledge Base article, *SonicWall TZ Series and SonicWall X-Series Solution Managing SonicWall Access Points*.

Managing Extended Switches using GMS

The switch integration feature allows unified management of both the Security Appliance and the switch using the SonicOS management interface and SonicWall GMS version 8.1 SP1 or higher. GMS supports all configuration operations, such as provisioning of an extended switch, configuration of extended switch interface settings, and manageability of extended switch global parameters.

For more information, refer to the GMS administration documentation located on the **Support** portal. Go to https://www.sonicwall.com/support/technical-documentation/ and select **GMS** in the **Select A Product** field.

Extended Switch Global Parameters

Extended switch global parameters shows the extended switch global parameters that can be configured through the SonicOS Management Interface.

NOTE: For more information, refer to the *SonicWall SonicOS 6.5 X-Series/N-Series Solution Deployment Guide* located on the Support portal at https://www.sonicwall.com/support/technical-documentation/ and select **NS*a* Series**, **NSA Series**, **SuperMassive 9000 Series**, and **TZ Series** in the **Select A Product** field.

Extended switch global parameters

All Switches	Only X1026P and X1052P switches
STP Mode	PoE Alert Usage Threshold

Extended switch global parameters (Continued)

All Switches	Only X1026P and X1052P switches
STP State	PoE Traps
	PoE Power Limit Mode

About Links

Management (MGMT) links carry only management traffic and cannot be portshielded.

Data links carry all PortShield traffic. If all they carry are data, the links are called common links. In a few topologies, data links also carry management traffic, in which case they are called shared links.

Shared or common links can carry all the portshielded groups.

Dedicated links can carry only one portshielded group, and that group must be portshielded to the dedicated port on the Security Appliance.

About Uplink Interfaces

Uplink interfaces can be viewed as “trunk” ports set up to carry tagged/untagged traffic. When an extended switch is added with Security Appliance uplink and X-Switch uplink options, the port on the Security Appliance configured as the SuperMassive uplink and the port on the extended switch configured as the switch uplink are set up automatically to receive/send tagged traffic for all IDV VLANs. The IDV VLAN of the tagged traffic allows the firmware to derive the PortShield host interface for the traffic.

Criteria for Configuring an Uplink Interface

- The interface must be a physical interface; virtual interfaces are not allowed.
- The interface must be a switch interface. (On some platforms, some Security Appliance interfaces are not connected to the switch. Such interfaces are not allowed.)
- The interface cannot be a PortShield host (some other Security Appliance interface cannot be portshielded to it) or a PortShield group member (cannot be portshielded to another Security Appliance interface).
- The interface cannot be a bridge primary or bridge secondary interface.
- The interface cannot have any children (it cannot be a parent interface for other child interfaces).

Logging and Syslog Support

Support for logging critical configuration events such as addition/deletion of a switch, configuration of PortShield on an extended switch port, and network events such as port coming up/going down is available.

Supported Topologies

IMPORTANT: Before setting up the interface between the Security Appliance and the switch, learn about provisioning, configuring, and setting up these topologies in the *SonicWall SonicOS 6.5 X-Series/N-Series Solution Deployment Guide* found on the Support portal at https://www.sonicwall.com/support/technical-documentation/ by selecting **NS*a* Series**, **NSA Series**, **SuperMassive 9000 Series**, and **TZ Series** in the **Select A Product** field.

NOTE: For basic details on configuring PortShield interfaces with X-Series/N-Series switches, see Managing Ports on page 361.

The key supported topologies for X-Series/N-Series switch support are:

- Common uplink configuration
- Dedicated uplink configuration

  IMPORTANT: SonicPoints must be portshielded through the port that is part of the dedicated link.

- Hybrid configuration with common and dedicated uplink(s)
- Shared link configuration for both management and data traffic
- Isolated links for management and data uplinks
- HA and PortShield configurations with dedicated uplink(s)
- HA and PortShield configurations with a common uplink
- VLAN(s) with common uplinks through SPM configuration
- VLAN(s) with dedicated uplink(s) configuration
- Dedicated link for SonicPoint access

SonicOS Support of N-Series Switches

Topics:

- About N-Series Switches
- Configuring the N-Series Switch
- Provisioning an N-Series Switch as an Extended Switch
- Significance of Uplink Interfaces
- Provisioning an N-Series Switch
- Configuring an Extended Switch in PortShield

About N-Series Switches

All platforms, except the NSA 2600 and SOHO W platforms, support these Dell® N-Series switches and the N-Series Solution.

One of the major pain-points for customers has been the need to manage critical network elements such as the Security Appliance and switch individually. Dell N-Series switches and Dell N-Series Solution allow unified management of both the Security Appliance and the N-Series switch using the Security Appliance's management interface and GMS. The maximum number of interfaces available on the TZ Series appliances range from 5 (TZ300) to 10 (TZ600). In certain deployments, the number of ports required might easily exceed the maximum number of interfaces available on the TZ Series appliance. With Dell N-Series Solution, ports on

the N-Series switch can be viewed as extended interfaces of the Security Appliance, thereby increasing the number of available interfaces.

One fundamental difference between N-Series switches and N-Series Solution, is the method by which the switch is programmed by the Security Appliance;

- N-Series Solution uses XML API as the mechanism for the Security Appliance to push/retrieve configuration to/from the switch.
- N-Series Solution uses CLI as the mechanism to configure the switch and retrieve config from the switch.

The functionality of the N-Series Solution supported on TZ Series, NSA Series, and SM series platforms is on a par with the feature set of the N-Series Solution solution supported on these platforms and is similar to the Dell X-Series Switch Solution. A maximum of four N-Series switches is supported. In the scenario where both N-series and N-Series switches are integrated with the same firewall, a maximum of four N-series + N-Series switches combined can be supported. Daisy-chaining of N-series switches is also supported.

Below are the key set of features supported with the initial phase of Dell Switch Integration solution:

- Provisioning of N-Series Switch as Extended Switch
- PortShield Functionality
- Configuration of N-Series Switch Interface Settings
- Manageability of basic N-Series Switch Global Parameters
- Manageability of Extended Switch using GMS
- High Availability plus PortShield
- Diagnostics Support for N-Series switch
- Daisy-Chaining of N-Series switches

Configuring the N-Series Switch

N-Series switches delivered from the factory by default have no IP address configured and DHCP enabled. For example, the configuration from a N1524 switch booted from default settings:

```
console#show running-config

!Current Configuration:
!System Description "Dell Networking N1524, 6.2.5.3, Linux 3.6.5"
!System Software Version 6.2.5.3
!
configure
stack
member 1 1    ! N1524
exit
interface vlan 1
ip address dhcp
exit
snmp-server engineid local 800002a203f48e3807701e
exit

console#
```

After the N-Series switch reboots, you can configure the initial setup using the Easy Setup Wizard. The wizard guides you through the initial switch configuration and gets you up and running as quickly as possible.

NOTE: You can exit the setup wizard at any point by entering `Ctrl Z`.

TIP: You can skip the setup wizard, and enter CLI mode to manually configure the switch.

IMPORTANT: You must respond to this question to run the setup wizard within 60 seconds:

```
Would you like to run the setup wizard (you must answer this question within
60 seconds)? (y/n) y
```

Otherwise, the system continues with normal operation using the default system configuration. Reset the switch with an empty startup configuration to rerun the Dell Easy Setup Wizard.

For how to configure the N-Series switch with the Dell Easy Setup Wizard, see the *Dell Getting Started Guide* for your switch.

Provisioning an N-Series Switch as an Extended Switch

NOTE: PoE-related fields are configurable only on PoE models of N-series switches.

On a standalone TZ Series system, provisioning of an Extended Switch specify these following Extended Switch parameters:

Mandatory Parameters		Optional Parameters	
• ID	• Username	• Firewall Uplink	• PoE Usage Threshold
• Switch Model	• Password	• Switch Uplink	• PoE Management Mode
• IP Address	• Switch Management	• STP Mode	• PoE Detection Type
		• STP State	

On a TZ Series system with High Availability enabled, an Extended Switch can be added by specifying these Extended Switch Parameters:

Mandatory Parameters		Optional Parameters	
• ID	• Username	• STP Mode	• PoE Usage Threshold
• Switch Model	• Password	• STP State	• PoE Management Mode
• IP Address	• Switch Management		• PoE Detection Type

NOTE: The Firewall Uplink and Switch Uplink parameters are not relevant for a Security Appliance operating in High Availability mode. Currently, the mandatory parameters cannot be modified after configuration.

Significance of Uplink Interfaces

Uplink Interfaces can be viewed as trunk ports setup to carry tagged/untagged traffic. When an Extended Switch is added with firewall uplink and switch uplink parameters, the port on the Security Appliance configured as the firewall uplink and the port on the extended switch configured as the switch uplink are setup to receive/send tagged traffic for all IDV VLANs automatically. The IDV VLAN of the tagged traffic allows SonicOS to derive the ingress interface, that is, the portshield host interface, for the traffic.

When configuring an interface as a firewall uplink, the interface:

- Should be a physical interface; virtual interfaces are not allowed.
- Should be a switch interface (on some platforms, some Security Appliance interfaces are not connected to the switch; such interfaces are excluded).
- Cannot be a PortShield host (some other Security Appliance interfaces cannot be portshielded to it) or cannot be a PortShield member (cannot be portshielded to another Security Appliance interface).

- Cannot be a bridge primary or bridge secondary interface.
- Cannot have any children (it cannot be a parent interface for other child interfaces).

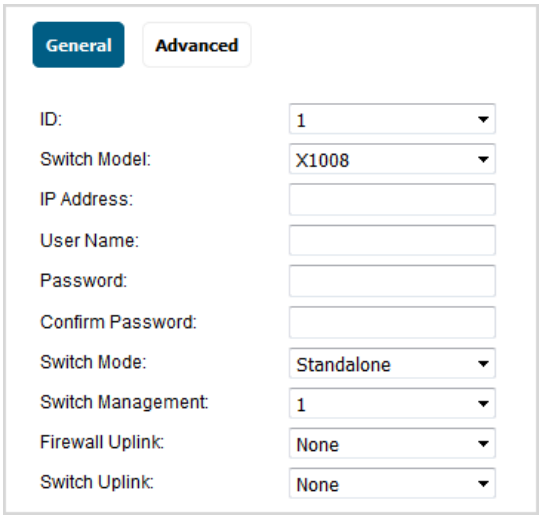
Provisioning an N-Series Switch

To provision an N-Series switch:

1 Navigate to **MANAGE | System Setup > Network -> PortShield Groups**.

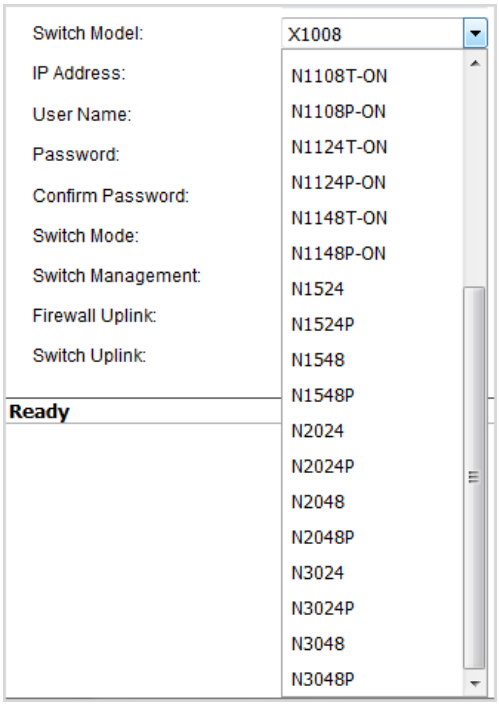
2 Click **External Switch Configuration**.

3 Click **ADD SWITCH**. The **Add External Switch** dialog displays.



4 Select an ID for this switch from **ID**. The default is **1**.

5 Select the type of switch from Switch Model.



6 Finish configuring the General options.

7 Click **Advanced**.

8 Finish configuring the Advanced options.

NOTE: PoE options display only for PoE N-Series switches.

9 Click **ADD**. The switch is added to the **External Switch Configuration** table.

Configuring an Extended Switch in PortShield

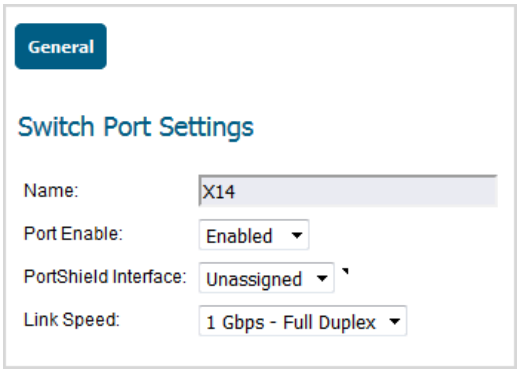
To configure an extended switch:

1 Navigate to **MANAGE | System Setup > Network > PortShield Groups**.

2 Click **Port Graphics**.

3 Click the port to be configured.

4 Click **CONFIGURE**. The **Edit Switch Port** dialog displays.

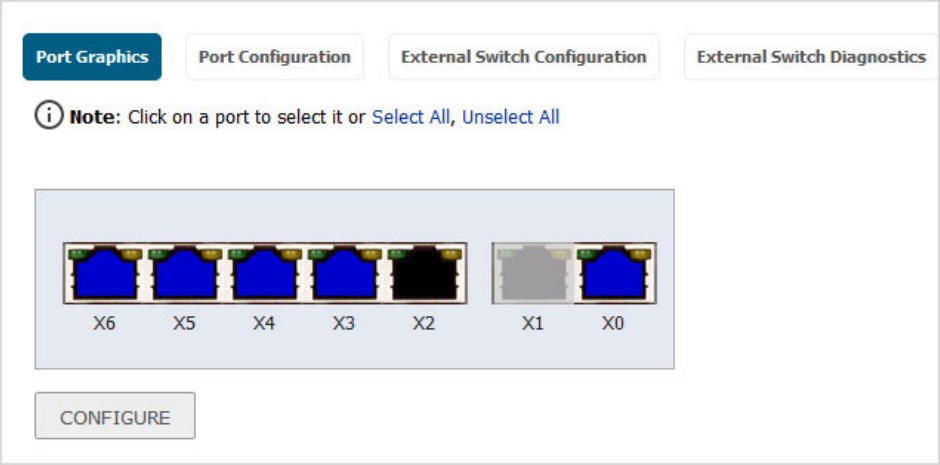


5 Configure the options.

6 Click **OK**.

Managing Ports

IMPORTANT: SOHO W Security Appliances do not support the X-Series Solution. Although all Security Appliance ports are managed the same, **MANAGE | System Setup | Network > PortShield Groups** is different for these Security Appliances; see Managing Ports on the SOHO W Firewall on page 369.



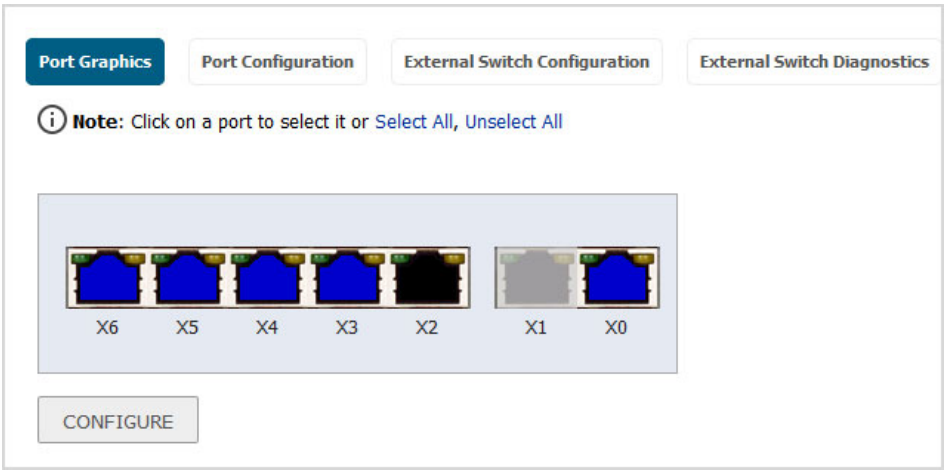
MANAGE | System Setup | Network > PortShield Groups allows you to manage the assignments of ports to PortShield interfaces through:

- **Port Graphics**
- **Port Configuration**
- **External Switch Configuration**
- **External Switch Diagnostics**

Topics:

- Viewing Interfaces (Ports) on Port Graphics on page 362
- Viewing Status of and Editing PortShield Interfaces on the Port Configuration Tab on page 364
- Viewing and Managing the External Switch Configuration on page 366
- Monitoring External Switch Diagnostics and Managing Firmware on page 367
- Managing Ports on the SOHO W Firewall on page 369

Viewing Interfaces (Ports) on Port Graphics



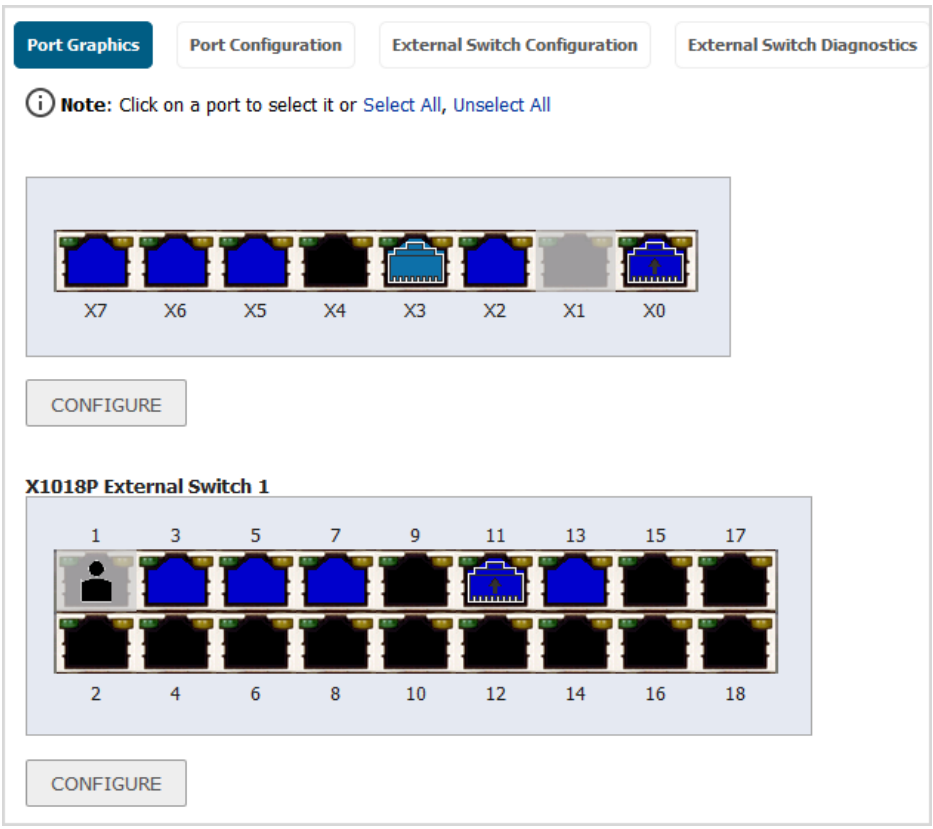
Port Graphics displays the PortShield interfaces (ports) for the Security Appliance. The large graphic represents the Security Appliance's interfaces. The interfaces are color coded to reflect their configuration:

Color code for interface configuration

This color	Designates this type of interface
Black	Unassigned, that is, not part of a PortShield group
Yellow	Selected to be configured
Same color (other than black, yellow, or grey)	Part of a PortShield group, with the master interface having a white outline around the color
Greyed out	Cannot be assigned, that is, added to a PortShield group
Grey interfaces with a person graphic	Switch MGMT
Any (other than black, yellow, or grey) with an up arrow	Uplink

Each port graphic is labeled with its associated port name: X0 - X*n*. When you select an interface or interfaces, you can configure them as described in Configuring PortShield Groups on page 370.

When an Extended Switch is Configured



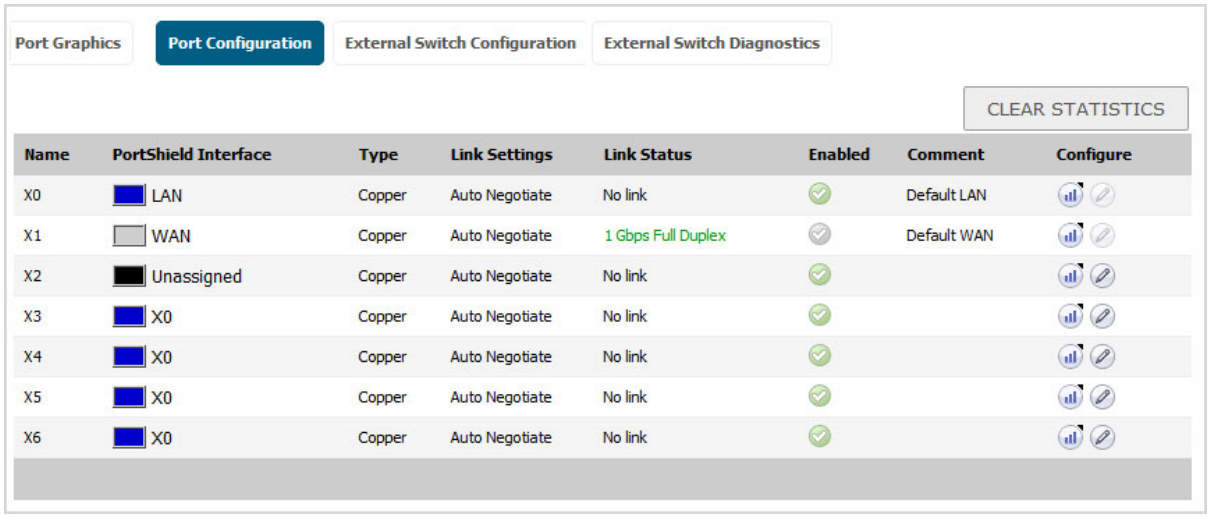
When one or more extended switches are provisioned, **Port Graphics** displays the PortShield interfaces (ports) for both the Security Appliance and the switch(es):

- The first graphic displays the Security Appliance's ports and is not labeled.
- The next graphic displays the ports for the first external switch, External Switch 1, which is labeled ***SwitchModel*** **External Switch 1**, for example, X1018P External Switch 1.
- If more external switches are provisioned, subsequent graphics display the ports for the other external switches in order of their ID, that is, External Switch 2, External Switch 3, and External Switch 4.

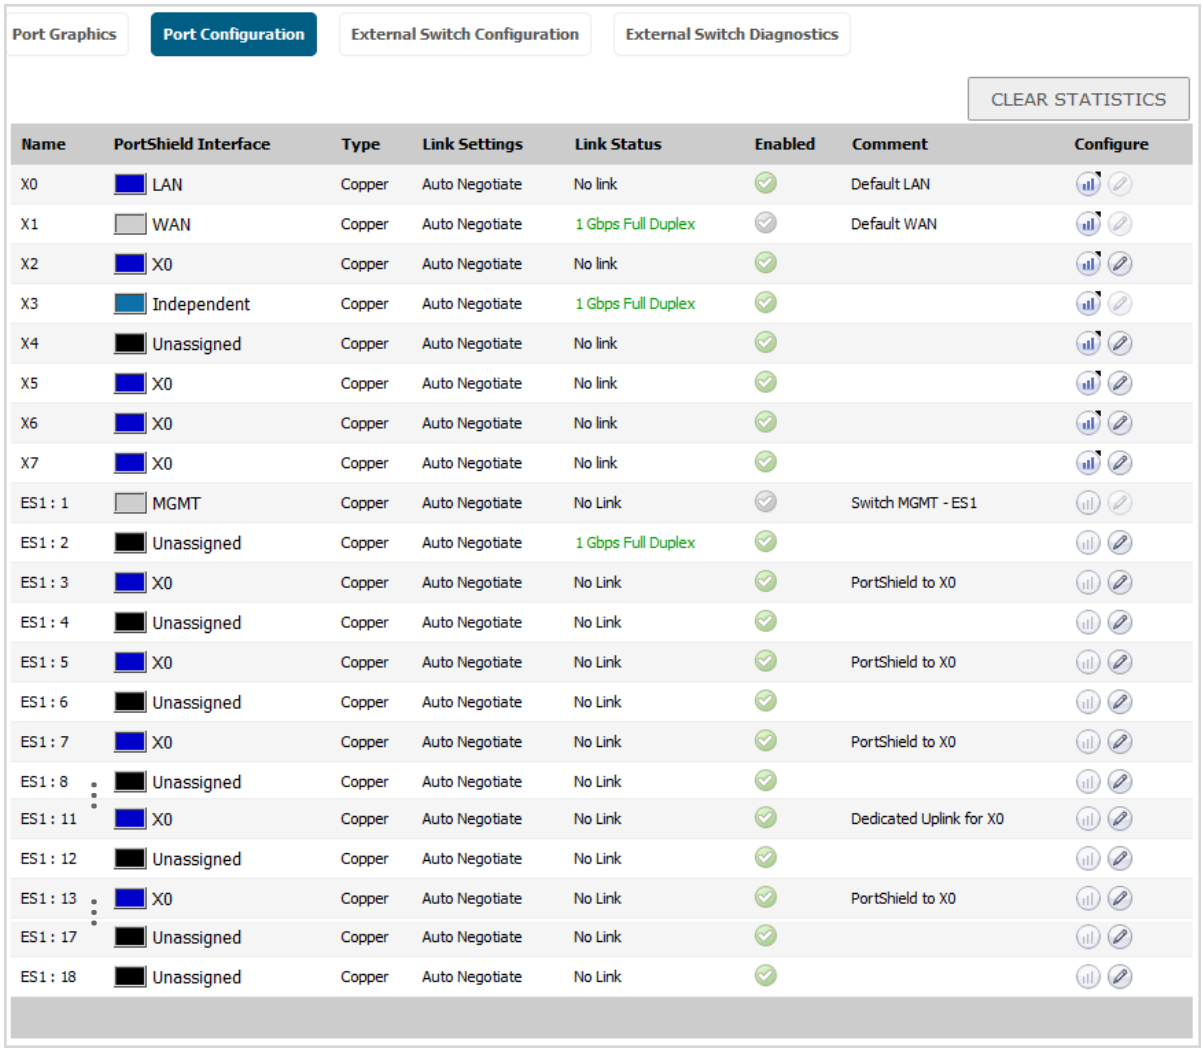
The color coding for external interfaces is the same as for the Security Appliance; see Color code for interface configuration.

Viewing Status of and Editing PortShield Interfaces on the Port Configuration Tab

Without an extended switch



With extended switches



Name	PortShield Interface	Type	Link Settings	Link Status	Enabled	Comment	Configure
X0	LAN	Copper	Auto Negotiate	No link		Default LAN	
X1	WAN	Copper	Auto Negotiate	1 Gbps Full Duplex		Default WAN	
X2	X0	Copper	Auto Negotiate	No link			
X3	Independent	Copper	Auto Negotiate	1 Gbps Full Duplex			
X4	Unassigned	Copper	Auto Negotiate	No link			
X5	X0	Copper	Auto Negotiate	No link			
X6	X0	Copper	Auto Negotiate	No link			
X7	X0	Copper	Auto Negotiate	No link			
ES1 : 1	MGMT	Copper	Auto Negotiate	No Link		Switch MGMT - ES1	
ES1 : 2	Unassigned	Copper	Auto Negotiate	1 Gbps Full Duplex			
ES1 : 3	X0	Copper	Auto Negotiate	No Link		PortShield to X0	
ES1 : 4	Unassigned	Copper	Auto Negotiate	No Link			
ES1 : 5	X0	Copper	Auto Negotiate	No Link		PortShield to X0	
ES1 : 6	Unassigned	Copper	Auto Negotiate	No Link			
ES1 : 7	X0	Copper	Auto Negotiate	No Link		PortShield to X0	
ES1 : 8	Unassigned	Copper	Auto Negotiate	No Link			
ES1 : 11	X0	Copper	Auto Negotiate	No Link		Dedicated Uplink for X0	
ES1 : 12	Unassigned	Copper	Auto Negotiate	No Link			
ES1 : 13	X0	Copper	Auto Negotiate	No Link		PortShield to X0	
ES1 : 17	Unassigned	Copper	Auto Negotiate	No Link			
ES1 : 18	Unassigned	Copper	Auto Negotiate	No Link			

Port Configuration consists of a table that lists information about the PortShield interfaces:

Name — Port name associated with the PortShield interface, such as X0 or X15. Ports for any external switches are shown in the format **ES*s*:*n***, where ***s*** is the switch ID and ***n*** is the port number, as appropriate.

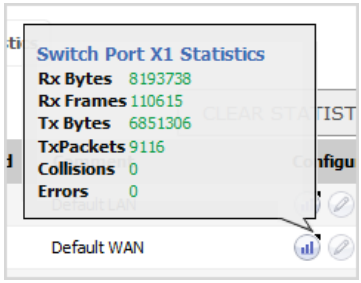
PortShield Interface — Color-coded graphic reflecting the PortShield interface's assignment and to which PortShield group it belongs. This graphic is a smaller version of the larger graphic(s) on **Port Graphics**.

Type — Type of port:

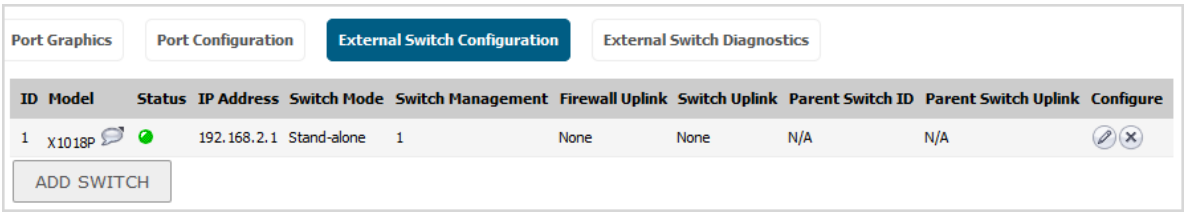
- **Copper**
- **Wireless**

Link Settings — Link speed:

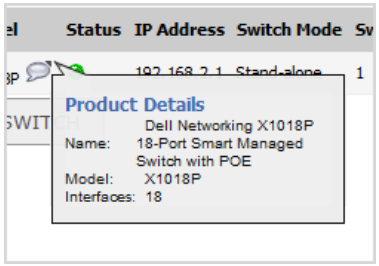
- **Auto Negotiate**
- **1000 Mbps – Full Duplex**
- **100 Mbps – Full Duplex**
- **100 Mbps – Half Duplex**
- **10 Mbps – Full Duplex**
- **10 Mbps – Half Duplex**

Link Status	Displays either: • The current link speed, in green, for example, **1000 Mbps – Full Duplex**. • **No link**.
Enabled	**Enable** icon that is: • Green if the interface is enabled. • Dimmed grey if the interface is disabled.
Comment	Any comment entered when the interface was configured.
Configure	Contains two icons: • **Statistics** – When clicked, displays a pop-up summary containing statistics about the interface:  **NOTE:** To clear all statistics, click **CLEAR STATISTICS** at the top of **Network > PortShield Groups > Port Configuration**. • **Edit** – When clicked, displays the **Edit Switch Port** dialog. For more information about this dialog, see the procedure in Configuring PortShield Interfaces on **Network > PortShield Groups** on page 371.

Viewing and Managing the External Switch Configuration



NOTE: This table displays **No Entries** if external switches have not been provisioned.

ID	ID number of the external switch: **1**, **2**, **3**, or **4**.
Model	Model number of the extended switch. For each switch, this column also contains a **Comment** icon that displays a popup summary with product details. 

Status	Status of the switch: A green **Enabled** icon indicates the switch is up and available. **NOTE:** When an extended switch has been powered off and then the Security Appliance is restarted (rebooted), it might take up to five minutes before the Security Appliance discovers the extended switch and reports the **Status** of the switch as up and available.
IP Address	IP address of the extended switch.
Switch Mode	Mode of the switch, such as **Stand-alone**.
Switch Management	Switch port used for management traffic.
Firewall Uplink	Port on the Security Appliance configured as the Security Appliance uplink. If no Security Appliance port has been configured as the Security Appliance uplink, the column displays **None**.
Switch Uplink	Port on the extended switch configured as the switch uplink. If no switch port has been configured as the switch uplink, the column displays **None**.
Parent Switch ID	For daisy-chained switches, the ID of the parent switch. If no switch port has been configured as the parent switch, the column displays **N/A**.
Parent Switch Uplink	Port on a daisy-chained parent switch configured as the switch uplink. If no switch port has been configured as the parent switch uplink, the column displays **N/A**.
Configure	Contains the: • **Edit** icon – Click to display the **Edit External Switch** dialog. • **Delete** icon – Click to delete the switch entry.

External Switch Configuration provides information about the external switches provisioned on the Security Appliance and allows you to manage the switch. You can also configure or delete an extended switch. To configure an extended switch, see Configuring PortShield Groups on page 370; to delete an extended switch, see the *SonicWall X-Series/N-Series Solution Deployment Guide*.

Monitoring External Switch Diagnostics and Managing Firmware

NOTE: The tables display **No Entries** if external switches have not been provisioned.

External Switch Diagnostics allows you to:

- Monitor statistics for the extended switch(es)
- Upload the firmware image and/or the boot image
- Restart the extended switch(es)

Topics:

- Changing the Display on page 367
- Monitoring Statistics on page 368
- Restarting the External Switch(es) on page 368
- Managing External Switch Firmware on page 369

Changing the Display

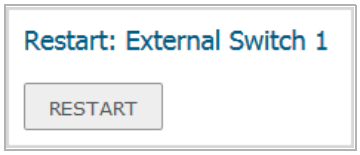
External Switch Diagnostics displays statistics and other information about only one switch at a time. By default, the data for External Switch 1, **ES1**, is displayed. If you have two or more external switches, to display data about a different external switch, choose **ES2**, **ES3**, or **ES4** from **Switch Name**:

Monitoring Statistics

The **Statistics** table displays a running tally of all statistics. To restart statistics collection, click **Clear** to reset the counters.

Name	Port name, 1 – *n*.
Status	Whether the port is **Up** or **Down**.
Rx Unicast Packets	Number of Unicast packets received on the port.
Rx Multicast Packets	Number of Multicast packets received on the port.
Rx Broadcast Packets	Number of Broadcast packets received on the port.
Rx Bytes	Number of bytes received on the port.
Rx Errors	Number of packets with errors received on the port.
Tx Unicast Packets	Number of Unicast packets transmitted on the port.
Tx Multicast Packets	Number of Multicast packets transmitted on the port.
Tx Broadcast Packets	Number of Broadcast packets transmitted on the port.
Tx Bytes	Number of bytes transmitted on the port.
FCS Errors	Number of packets with FCS (frame check sequence) errors received on the port.
Single Collision Frames	Number of frame collisions detected on the port.
Late Collisions	Number of frame collisions detected after the last frame bit was sent on the port.
Excessive Collisions	Number of frame collisions detected that exceeded the number of retries on the port.
Internal MAC Transmit Errors	Number of non-collision transmission errors detected on the port.
Oversized packets	Number of received packets larger than the port was expecting.
Rx Pause Frames	Number of pause frames received by the port.
Tx Pause Frames	Number of pause frames sent by the port.

Restarting the External Switch(es)

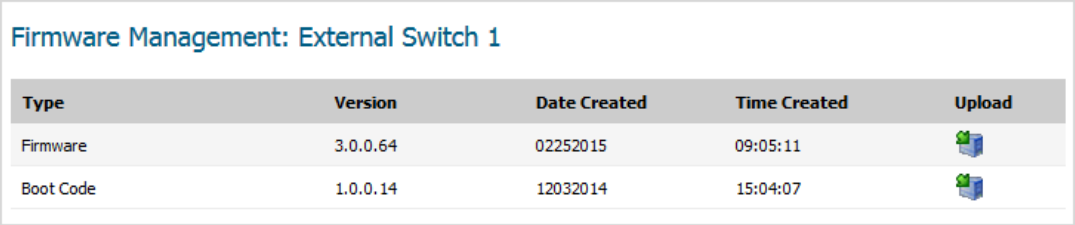


IMPORTANT: When an extended switch has been powered off and then the Security Appliance is restarted (rebooted), it might take up to five minutes before the Security Appliance discovers the extended switch and reports the **Status** of the switch as **Connected**.

To restart an external switch:

1. Navigate to **MANAGE | System Setup | Network > PortShield Groups**.
2. Click **External Switch Diagnostics**.
3. Select which external switch to restart from **Switch Name**.
4. Scroll to the **Restart: External Switch** section.
5. Click **RESTART**.

Managing External Switch Firmware

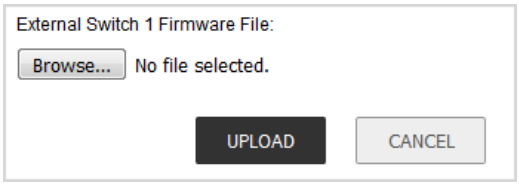


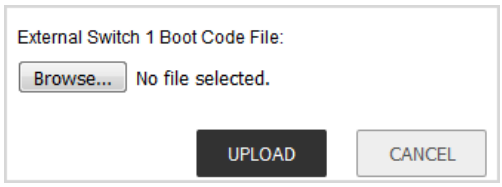
The **Firmware Management: External Switch** table displays information about the external switch's firmware and boot code:

Type	Either **Firmware** or **Boot Code**.
Version	Version of firmware or boot code on the external switch.
Date Created	Date the firmware or boot code was created.
Time Created	Time the firmware or boot code was created.
Upload	**Upload** icon; for • **Firmware**, displays the **Upload External Switch Firmware** dialog: • **Boot Code**, displays the **Upload External Switch Boot Code** dialog:

To upload firmware or boot code:

1 Click **Upload** for either **Firmware** or **Boot Code**. The **Upload External Switch Firmware** or **Upload External Switch Boot Code** dialog displays.





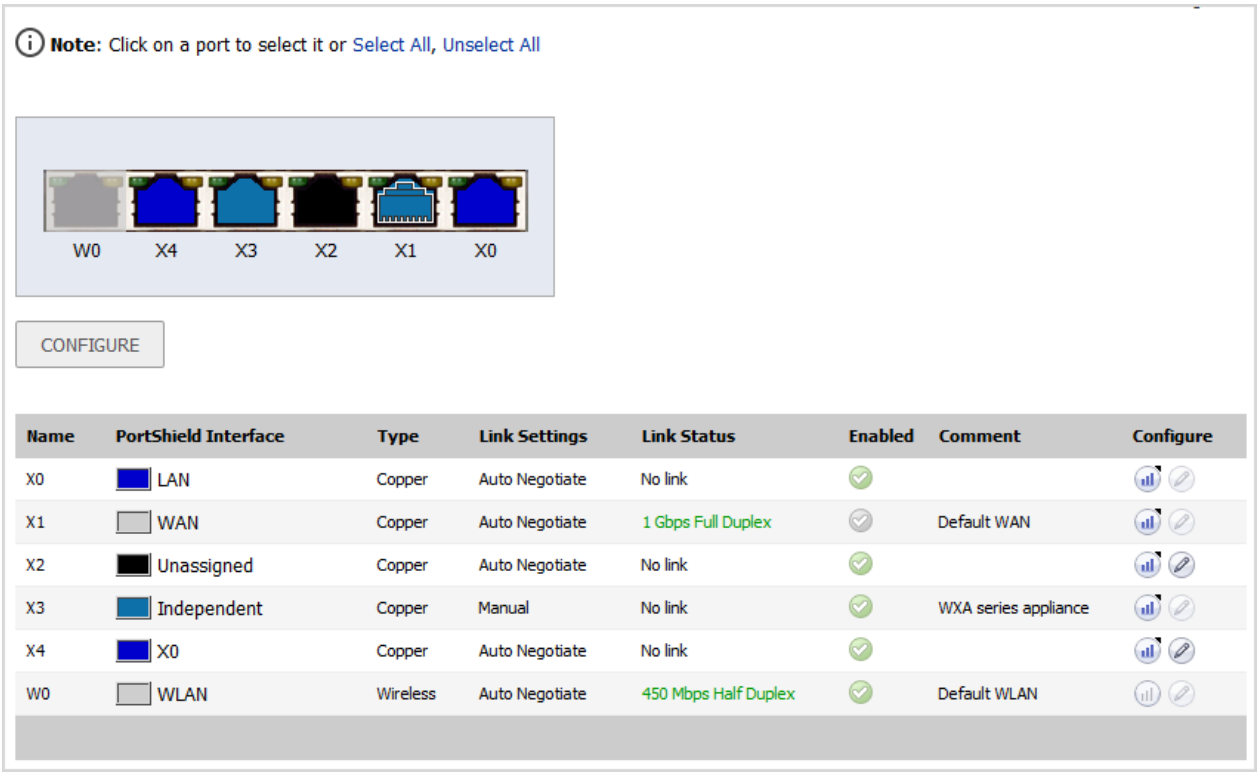
2 Click **Browse**. The **File Upload** dialog displays.

3 Select the file.

4 Click **Upload**.

Managing Ports on the SOHO W Firewall

The **Network > PortShield Groups** page for the SOHO W Security Appliance has a different look. The information on this page combines the information on **Port Graphics** (see Viewing Interfaces (Ports) on Port Graphics on page 362) and **Port Configuration** (Viewing Status of and Editing PortShield Interfaces on the Port Configuration Tab on page 364).



Configure Security Appliance interfaces as described in Configuring PortShield Groups on page 370.

Configuring PortShield Groups

PortShield groups can be configured on several different pages in the SonicOS Management Interface:

- Configuring PortShield Interfaces on **Network > Interfaces** on page 370
- Configuring PortShield Interfaces with the PortShield Interface Guide (TZ Series and SOHO W Firewalls Only) on page 371
- Configuring PortShield Interfaces on **Network > PortShield Groups** on page 371
- Configuring External Switch PortShield Groups from Port Graphics on page 373

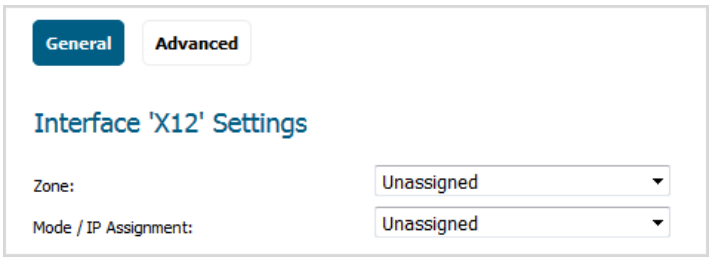
Configuring PortShield Interfaces on Network > Interfaces

IMPORTANT: For a port to be an interface, it must be configured with an IP address. Otherwise, the port is not listed in **PortShield Interface**.

To configure a PortShield interface:

1 Navigate to **MANAGE | System Setup | Network > Interfaces**.

2 In the **Interface Settings** table, click the **Configure** icon for the interface you want to configure. The **Edit Interface** dialog displays.



3 From **Zone**, select on a zone type option to which you want to map the interface. More options display.

NOTE: You can add PortShield interfaces only to **Trusted**, **Public**, and **Wireless** zones.

4 In the **Mode / IP Assignment** drop-down menu, select **PortShield Switch Mode**. The options change again.

5 From **PortShield to**, select the interface you want to map this port to. Only ports that match the zone you have selected are displayed.

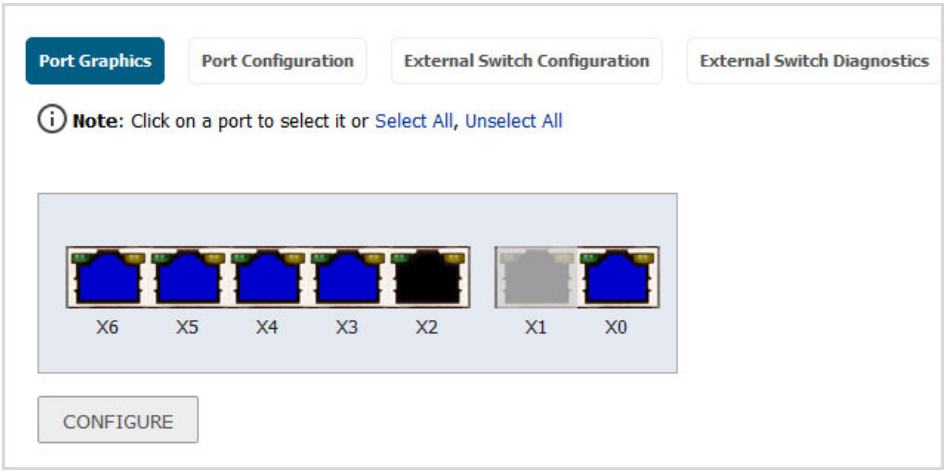
6 Click **OK**.

Configuring PortShield Interfaces with the PortShield Interface Guide (TZ Series and SOHO W Firewalls Only)

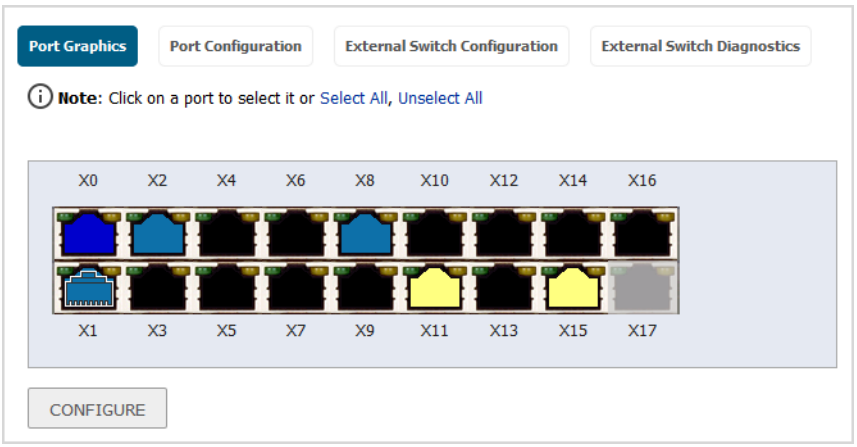
You can configure PortShield interfaces through the PortShield Interface Guide as described in the *SonicOS Quick Configuration Guide*. You can access the PortShield Interface Guide in these ways:

- Clicking **Quick Configuration Guide** on any Management Interface page. The **Configuration Guide** displays; select **PortShield Interface Guide**.
- On the **MANAGE | System Setup | Network > Interfaces** page on a TZ Series or SOHO W Security Appliance, click **PORTSHIELD WIZARD** to display the **PortShield Interface Guide**.

Configuring PortShield Interfaces on Network > PortShield Groups



Port Graphics displays a graphical representation of the current configuration of PortShield interfaces. For a description of the graphic display, see Viewing Interfaces (Ports) on Port Graphics on page 362.

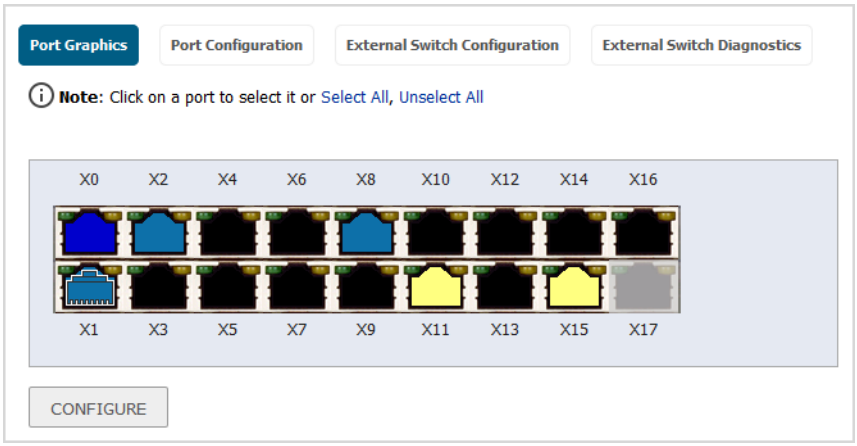


You can manually group ports using the graphical **PortShield Groups** interface by clicking on the ports you want to group. Grouping ports allows them to share a common network subnet as well as common zone settings.

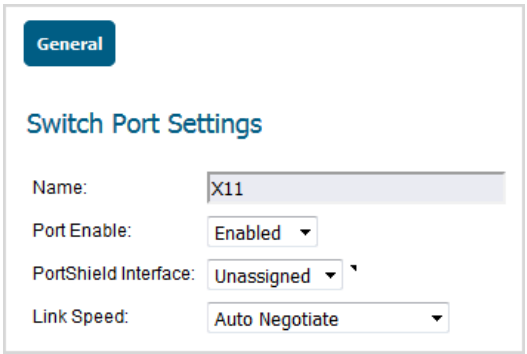
NOTE: Interfaces must be configured before being grouped with PortShield.

To configure PortShield groups:

1 In the port graphic, select the interface(s) you want to configure as part of a PortShield group. The interfaces turn yellow.



2 Click **CONFIGURE**. The **Edit Switch Port** dialog displays.



NOTE: The name of the interface for this port is dimmed and cannot be changed.

3 From **Port Enable**, select whether you want to enable or disable the interfaces. The default is **Enabled**.

4 From **PortShield Interface**, select which interface you want to assign as the master interface for this PortShield interfaces. The default is **Unassigned**.

NOTE: PortShield options may be disabled for external switch ports.

5 From **Link Speed**, select the link speed for the interfaces:

- **Auto Negotiate** (default)
- **1000 Mbps – Full Duplex**
- **100 Mbps – Full Duplex**
- **100 Mbps – Half Duplex**
- **10 Mbps – Full Duplex**
- **10 Mbps – Half Duplex**

6 Click **OK**.

Configuring External Switch PortShield Groups from Port Graphics

IMPORTANT: When an extended switch has been powered off and then the Security Appliance is restarted (rebooted), it mght take up to five minutes before the Security Appliance discovers the extended switch and reports the **Status** of the switch as **Connected**.

When configuring extended switches in a PortShield group, it might take up to five minutes for the configuration to be displayed on **Network > PortShield Groups**.

IMPORTANT: Interfaces must be configured before being grouped with PortShield.

NOTE: For more information, go to https://www.sonicwall.com/support/technical-documentation/ and search for the ***SonicWall SonicOS 6.5 X-Series/N-Series Solution Deployment Guide*** by selecting **NS*a* Series**, **NSA Series**, **SuperMassive 9000 Series**, and **TZ Series** in the **Select A Product** field.

NOTE: Extended switches are not supported on the SOHO W Security Appliance.

Network > PortShield Groups displays a graphical representation of the current configuration of PortShield interfaces on both the Security Appliance and the extended (external) switch(es). If there is one external switch, there are two graphics; for two external switches, there are three graphics, and so on. The switch graphics are labeled with the switch model and the external switch ID: 1, 2, 3, 4.

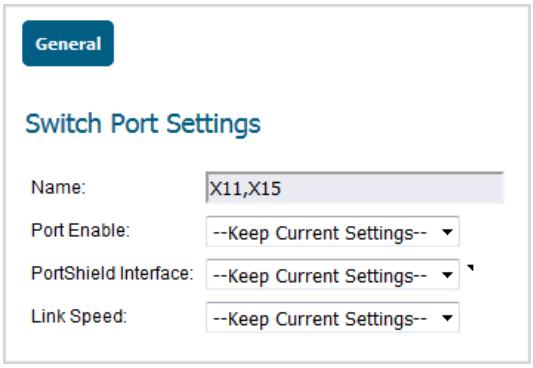
You can manually group ports on the Security Appliance and switches together using the graphical PortShield Groups interface by clicking on the ports you want to group. Grouping ports allows them to share a common network subnet as well as common zone settings.

To configure PortShield groups with external switches:

1 Configure the ports on the Security Appliance by following the procedure in Configuring PortShield Interfaces on **Network > PortShield Groups** on page 371.

2 In the port graphic for the external switch, select the interface(s) you want to configure as part of the PortShield group. The interfaces turn yellow.

3 Click **Configure**. The **Edit Multiple Switch Ports** dialog displays.



The **Name** field is dimmed and cannot be modified. It displays the names of both the Security Appliance's and external switch's ports you selected (*n* is the selected port):

- Firewall ports are named **X*n***.
- External switch 1 ports are named **ES1 : *n***.
- External switch 2 ports are named **ES2 : *n***.
- External switch 3 ports are named **ES3 : *n***.
- External switch 4 ports are named **ES4 : *n***.

4 From **Port Enable**, select:

- **Disabled**
- **Enabled**
- **—Keep Current Settings—** (default) – By default, all ports on the extended switch are enabled.

5 From **PortShield Interface**, select which interface you want to assign as the master interface for these PortShield interfaces:

- **Unassigned**
- Port name

  IMPORTANT: For a port to be an interface, it must be configured with an IP address. Otherwise, the port is not listed in **PortShield Interface**.

- **—Keep Current Settings—** (default)

  NOTE: PortShield options could be disabled for external switch ports. Ports that are portshielded here are configured automatically as access VLANs for the corresponding PortShield VLAN.

6 From **Link Speed**, select the link speed for the interfaces:

- **Auto Negotiate**
- **1000 Mbps – Full Duplex**
- **100 Mbps – Full Duplex**
- **100 Mbps – Half Duplex**
- **10 Mbps – Full Duplex**
- **10 Mbps – Half Duplex**
- **—Keep Current Settings—** (default) – By default, the link speed for all ports on the extended switch are set to **Auto Negotiate**.

7 Click **OK**.

15

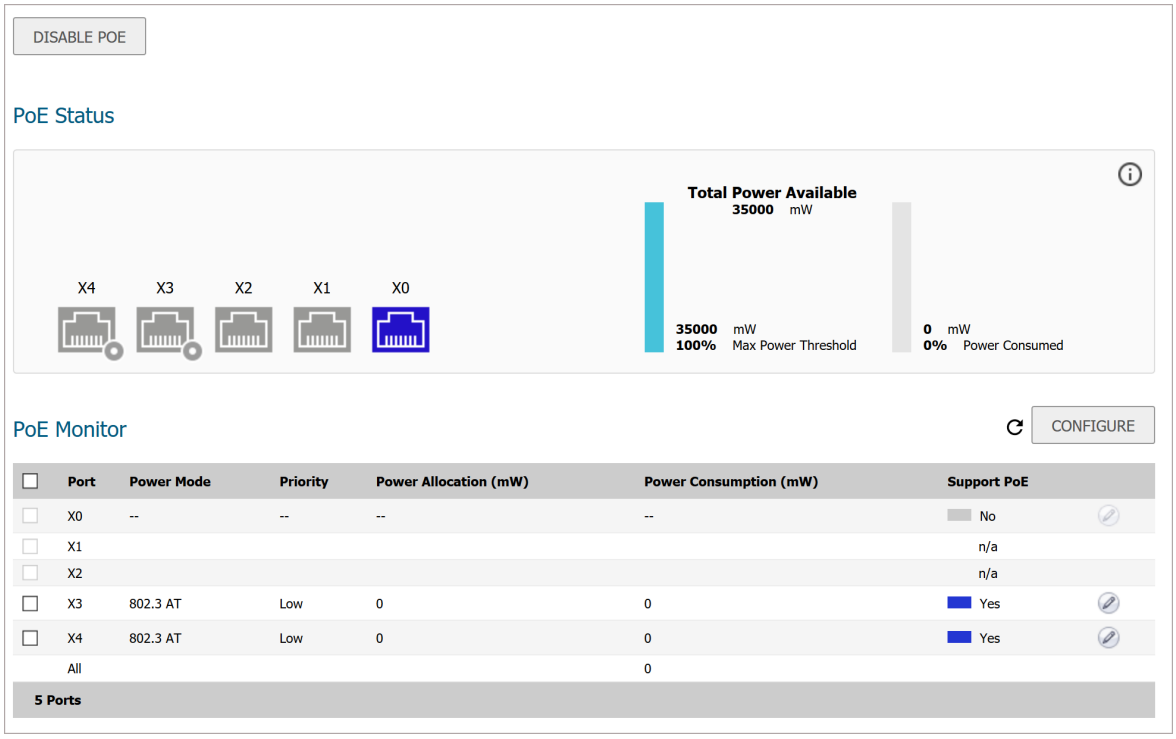
Configuring PoE

IMPORTANT: Only the TZ600P and TZ300P firewalls display the **MANAGE | System Setup** > **Network > PoE Settings** page.

Topics:

- Network > PoE Settings on page 375
  - Enabling PoE on page 376
  - Configuring Generic PoE Settings on page 377
  - Enabling PoE on page 376

Network > PoE Settings



The **MANAGE | Network** > **Network > PoE** page provides an easy-to-use, centralized dashboard to configure and monitor:

- Overall power consumption
- Enable/disable the whole POE controller or an individual port

- Set the power threshold of individual ports
- Set power mode (802.3 AT or 802.3 AF)
- Set the power priority
- Change multiple ports simultaneously)

Topics:

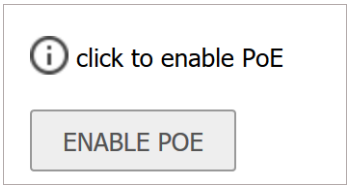
- Enabling PoE on page 376
- Configuring Generic PoE Settings on page 377

Enabling PoE

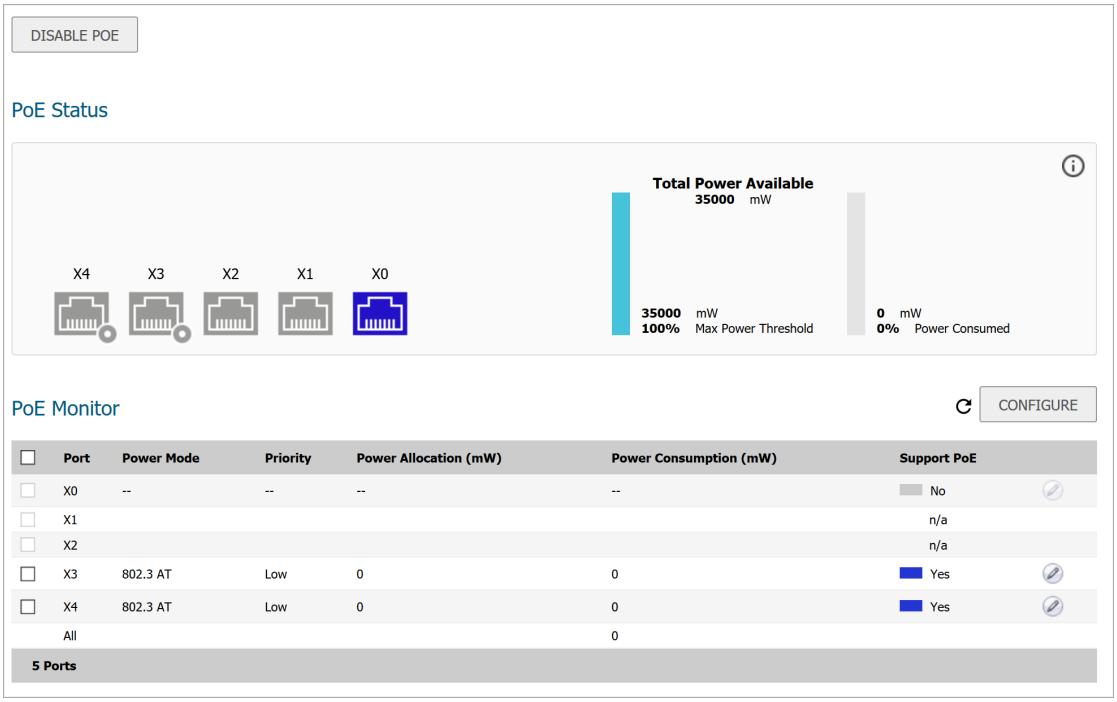
Before using PoE on your TZ600P/TZ300P Security Appliance, you must first enable it.

To enable PoE on your Security Appliance:

1 Navigate to **MANAGE | System Setup > Network > PoE Settings**.



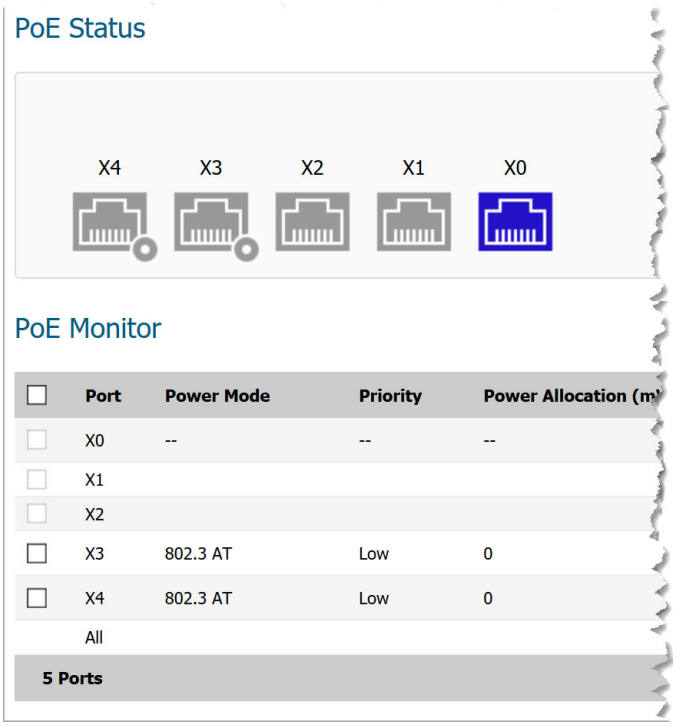
2 Click **ENABLE POE**. The page expands.



The **Network > PoE Settings** page has two sections:

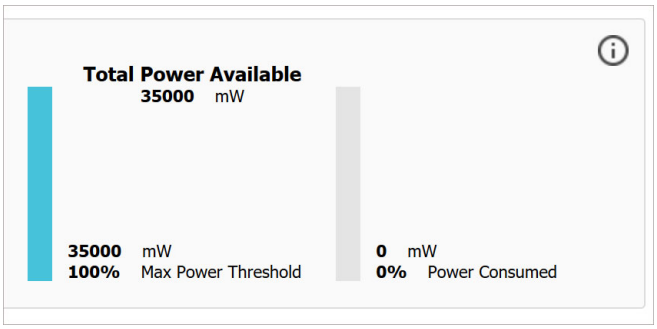
- **PoE Status** – Provides graphic visualization by displaying a:

- Status indicator, with the PoE interfaces (ports) designated by a circle icon and different colors.



These ports are also shown in the **PoE Monitor** table.

- Real-time power consumption monitor showing the total power available, the maximum power threshold percentage, and the percentage of power consumed.



Clicking the Information icon displays more information about the ports.

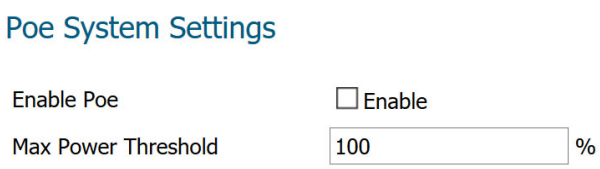
- **PoE Monitor** – lists the ports, and for each port, the power mode, priority, power allocations (mW), power consumption, whether it supports PoE, and an **Edit** icon for the PoE ports.

Configuring Generic PoE Settings

To configure generic PoE settings for a port:

1. Navigate to **MANAGE | System Setup > Network > PoE Settings**.
2. Scroll to the **PoE Monitor** section.

3 Click **CONFIGURE**. The **Poe System Settings** dialog displays.



4 To enable PoE, select **Enable Poe Enable**.

5 Specify the percentage for the maximum power threshold in the **Max Power Threshold** field.

6 Click **OK**.

Configuring Power Settings for a Port

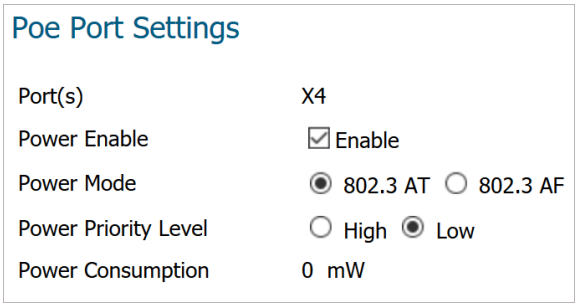
To configure power settings for a port:

1 Navigate to **MANAGE | System Setup > Network > PoE Settings**.

2 Scroll to the **PoE Monitor** section.

3 Select a port on which to configure the power settings.

4 Click the **Edit** icon for the port. The **Poe Port Settings** dialog displays.



5 To enable power on the port, select **Power Enable Enable**. This option is selected by default.

6 Chose power mode from the **Power Mode** options:

- **802.3 AT** (default)
- **802.3 AF**

7 Select the priority for the port's priority from the **Power Priority Level** options:

- **High**
- **Low** (default)

8 Click **OK**.

16

Setting Up Failover and Load Balancing

Topics:

- **Network > Failover & Load Balancing** on page 379
  - About Failover and Load Balancing on page 379
  - How Failover and Load Balancing Work on page 380
  - Multiple WAN (MWAN) on page 381
  - **Network > Failover & Load Balancing** on page 382
  - Configuring Failover and LB Groups on page 384
  - Configuring Probe Settings for Group Members on page 388

Network > Failover & Load Balancing

Topics:

- About Failover and Load Balancing on page 379
- How Failover and Load Balancing Work on page 380
- Multiple WAN (MWAN) on page 381
- **Network > Failover & Load Balancing** on page 382
- Configuring Failover and LB Groups on page 384
- Configuring Probe Settings for Group Members on page 388

About Failover and Load Balancing

Failover and Load Balancing (LB) (together, FLB) is a mechanism that actively monitors WAN connections and acts accordingly on failure/recovery of the WAN interface(s). The overall effect is a system-wide response to failure/recovery of WAN connections. Even if you only have one WAN, you still benefit because of faster recovery procedures performed on that one WAN as normal part of FLB (for more information about FLB with one WAN, see Knowledge Base article, SW13851, *Can I disable global Load Balancing if only one WAN is used on the firewall?*). In essence, FLB provides a highly-available system.

For FLB, multiple WAN members are supported (*N*–1), where *N* is the total number of interfaces on a hardware platform). For example:

- Primary WAN Ethernet Interface
- Alternate WAN #1

- Alternate WAN #2
- Alternate WAN #<*n–1*> …

IMPORTANT: It is recommended that Load Balancing be enabled at all times, even if there is only one WAN. For more information, see *Can I disable global Load Balancing if only one WAN is used on the firewall?* (SW13851).

The **Primary WAN Ethernet Interface** has the same meaning as the previous concept of "Primary WAN." It is the highest ranked WAN interface in the LB group. The **Alternate WAN #1** corresponds to "Secondary WAN," it has a lower rank than the Primary WAN, but a higher rank than the next two alternates. The others, **Alternate WAN #2** and **Alternate WAN #<*n–1*>**, are new, with **Alternate WAN #<*n–1*>** being the lowest ranked among the four WAN members of the LB group.

How Failover and Load Balancing Work

Topics:

- WAN Interface Failure on page 380
- WAN Interface Recovery on page 380

WAN Interface Failure

This is what FLB does when a WAN interface failure had been detected (linkDown or probing-failure or no-IP-settings):

1. Graceful shutdown of the interface (call the stop API, if one is provided; for example, pppoe-stop, dialup-stop).
2. Trigger the disabling of routes associated with the failed interface (except for the ones marked `do not disable on link down`).
3. Flush dynamic ARP entries using the failed interface.
4. Flush the cache entries using the failed interface as the outbound interface.
5. Update the WAN default route to point to an alternate WAN, if available. Update status data (this is part of recovery procedure).
    - Address Objects used by other apps such as CASS gets updated as well.
    - Security Services depend on this for failover capability.
6. Notify interested parties (VPN, BWM, CASS, DDNS, DNS).
7. Actively monitor status of failed interface, attempt recovery such as restarting WAN connection (call the start API, if provided; for example, pppoe-start, dial-start).

WAN Interface Recovery

This is what FLB does when a WAN interface recovery had been detected (linkUp or probing-success or IP-change):

1. On linkUp, jump-start the interface connection (call the start API, if provided; for example, pppoe-start, dial-start). In most cases, this would be in a connected state already, but if it is not, FLB attempts to push it to start. It may do a graceful shutdown and restart if a hung condition is detected (timer-based).
2. When connectivity is confirmed (simple linkUp or probing), trigger enabling of routes associated with the interface.

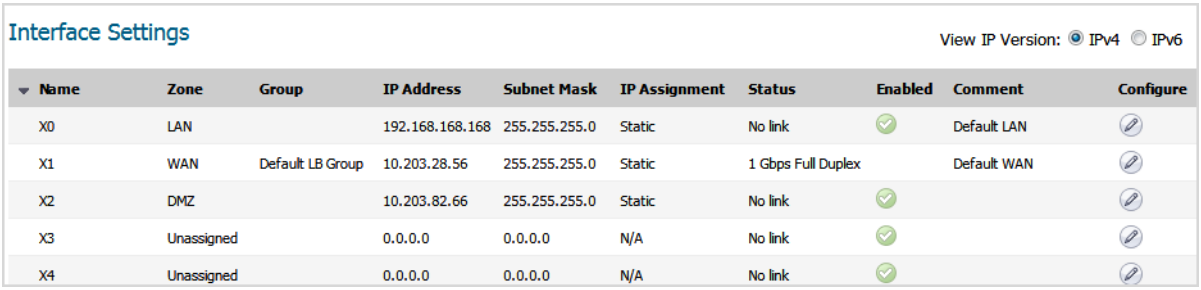
3 Add ARP entries (if any are needed).
   - Send out unsolicited ARP response (for interface) to update neighboring devices.
4 If needed, update the WAN default route (for example, preempt) to use the best available WAN. Update status data.
   - Address Objects used by other apps such as CASS gets updated as well.
   - Security Services depend on this for failover capability.
5 Notify interested parties (VPN, BWM, CASS, DDNS, DNS).
6 Continue monitoring status of interface.

Multiple WAN (MWAN)

The Multiple WAN (MWAN) feature allows you to configure all but one of the appliance's interfaces for WAN network routing (one interface must remain configured for the LAN zone for local administration). All of the WAN interfaces can be probed using the SNWL Global Responder host.

Network Interfaces

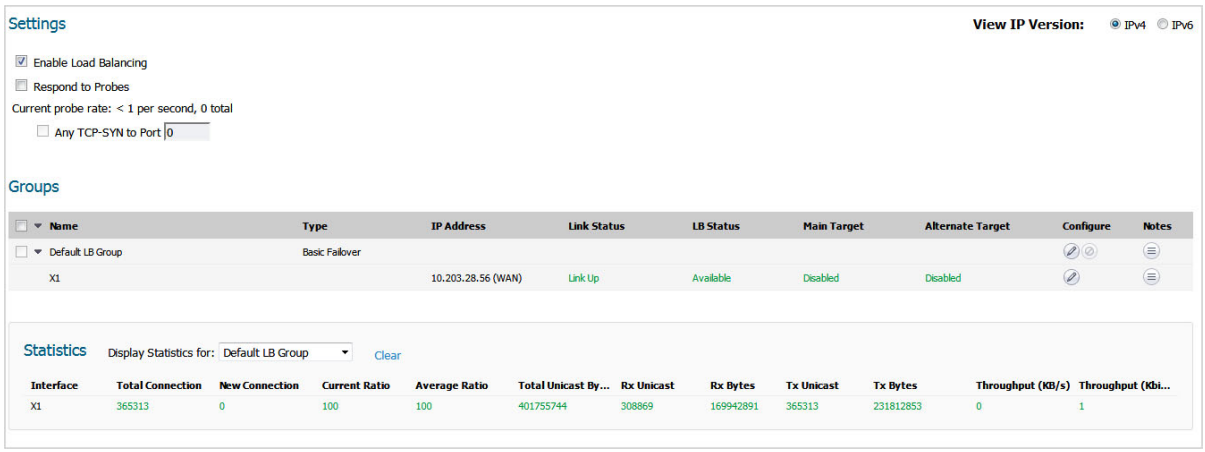
MANAGE | Network > Interfaces allows more than two WAN interfaces to be configured for routing. It is possible to configure WAN interfaces in **Network > Interfaces**, but not include them in **Network > Failover & Load Balancing**. Only the Primary WAN Ethernet Interface is required to be part of the LB group whenever LB has been enabled. Any WAN interface that does not belong to the LB group is not included in the LB function, but performs normal WAN routing functions.



NOTE: A virtual WAN interface may belong to the LB group. However, prior to using within the LB group, please ensure that the virtual WAN network is fully routable like that of a physical WAN.

A default gateway IP is required on the WAN interface if any destination is required to be reached through the WAN interface that is not part of the WAN subnet IP address space, regardless whether a default route is received dynamically from a routing protocol of a peer device on the WAN subnet.

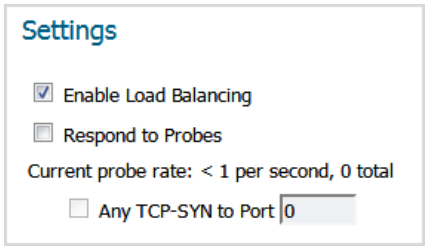
Network > Failover & Load Balancing



Topics:

- Settings on page 382
- Groups on page 383
- Statistics on page 384

Settings



- **Enable Load Balancing**—This option must be enabled for the user to access the LB Groups and LB Statistics sections of the Failover & Load Balancing configuration. If disabled, no options for Failover & Load Balancing are available to be configured. This option is enabled by default.

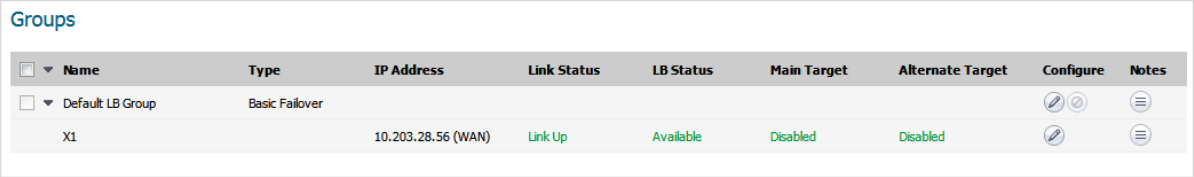
  IMPORTANT: It is recommended that Load Balancing be enabled at all times, even if there is only one WAN. For more information, see *Can I disable global Load Balancing if only one WAN is used on the firewall?* (SW13851).

- **Respond to Probes**—When enabled, the appliance can reply to probe request packets that arrive on any of the appliance's interfaces. This option is not selected by default.

  The current probe rate and total number of probes are displayed.

  - **Any TCP-SYN to Port**—This option is available when the **Respond to Probes** option is enabled. When selected, the appliance only responds to TCP probe request packets having the same packet destination address TCP port number as the configured value. This option is not selected by default.

Groups



LB Members added to a LB Group take on certain “roles.” A member can only work in one of the following roles:

- **Primary**—Only one member can be the Primary per Group. This member always appears first or at the top of the member list.

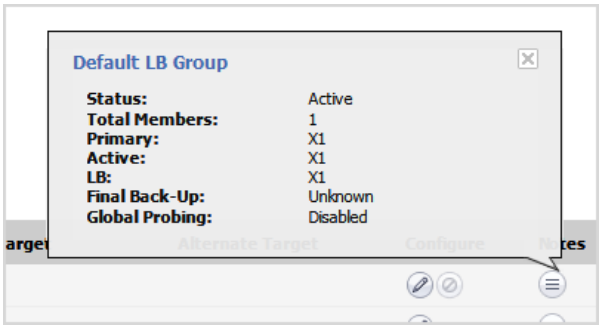
  NOTE: Although a group can be configured with an empty member list, it is impossible to have members without a Primary.

- **Alternate**—More than one member can be an Alternate; however, it is not possible to have a Group of only Alternate members.
- **Last-Resort**—Only one member can be designed as Last-Resort. Last-Resort can only be configured with other group members.

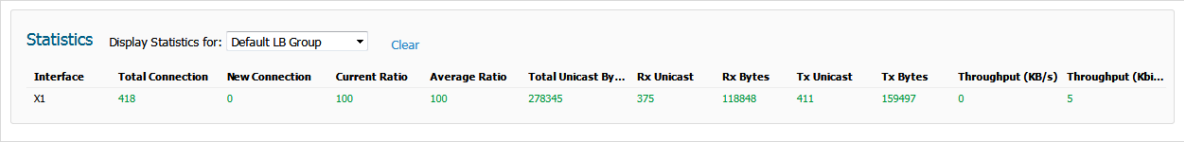
Each member in a group has a rank. Members are displayed in descending order of rank. The rank is determined by the order of interfaces as they appear in the Member List for the group. The order is important in determining the usage preferences of the Interfaces, as well as the level of precedence within the group. Thus, no two interfaces within a group have the same or equal rank; each Interface has a distinct rank.

Groups Table

- **Expand/Collapse** icon – Click to expand or collapse the group to show the members.
- **Checkbox** – Used to select a group; the default group cannot be selected.
- **Type** – The type of failover; only for groups, not members.
- **IP Address** – The IP address of the group member.
- **Link Status** – Displays whether the link is Link Up or Link Down.
- **LB Status** – Displays the status of load balancing.
- **Main Target** – Displays whether probing is performed on the main target.
- **Alternate Target** – Displays whether probing is performed on the alternate target.
- **Configure** – Displays the **Edit** icon and, for groups, the **Delete** icon (the default group cannot be deleted, so the **Delete** icon is dimmed).
- **Notes** – Displays the **Notes** icon, which, when moused over, displays a popup balloon with status about the group.



Statistics



From the **Display Statistics for** drop-down menu, select the LB group for which you want to view statistics.

The Load Balancing **Statistics** table displays the following LB group statistics for the firewall:

- Interface –
- Total Connections –
- New Connection –
- Current Ratio –
- Average Ratio –
- Total Unicast Bytes –
- Rx Unicast –
- Rx Bytes –
- Tx Unicast –
- Tx Bytes –
- Throughput (KB/s) –
- Throughput (Kbits/s) –

Click **Clear** on the top right of the **Statistics** table to clear its information.

Configuring Failover and LB Groups

Topics:

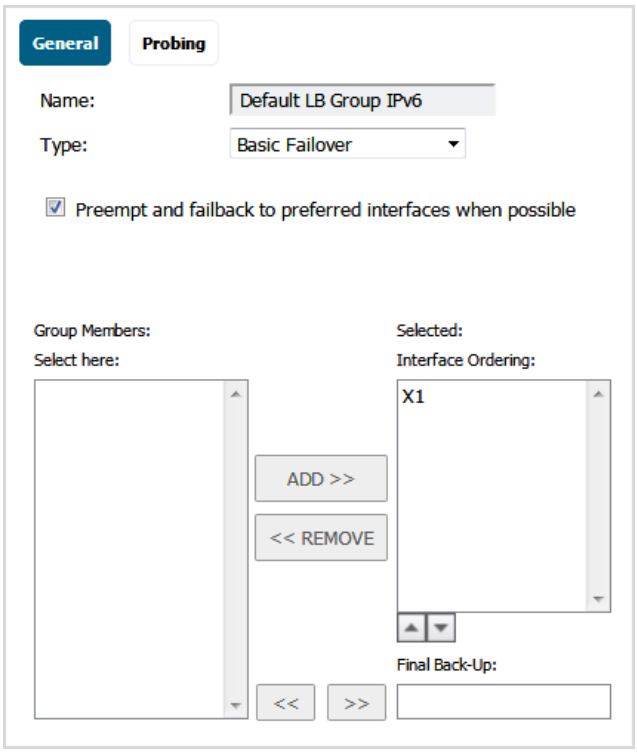
- General Settings on page 385
- Probing Settings on page 387

General Settings

To configure the Group settings:

1 Navigate to **MANAGE | Network > Failover & Load Balancing**.

2 Click the **Configure** icon of the Group you wish to configure. The **Edit LB Group** dialog displays.



3 Edit the display name of the Group in the **Name** field. The name of the default group cannot be changed and the field is dimmed.

4 From the **Type** drop-down menu, choose the type (or method) of LB; options change depending on the type selected:

- **Basic Failover**—The four WAN interfaces use rank to determine the order of preemption when the **Preempt** checkbox has been enabled. Only a higher-ranked interface can preempt an Active WAN interface. This is selected by default.
- **Round Robin**—This option now allows you to re-order the WAN interfaces for Round Robin selection. The default order is:
  - Primary WAN
  - Alternate WAN #1
  - Alternate WAN #2
  - Alternate WAN #3

  The Round Robin then returns to the Primary WAN to continue the order.
- **Spill-over**—The bandwidth threshold applies to the Primary WAN. When the threshold is exceeded, new traffic flows are allocated to the Alternates in a Round Robin manner. If the

Primary WAN bandwidth goes below the configured threshold, Round Robin stops, and outbound new flows are again sent out only through the Primary WAN.

NOTE: Existing flows remain associated with the Alternates (as they are already cached) until they time out normally.

- **Ratio**—A percentages can be set for each WAN in the LB group. To avoid problems associated with configuration errors, ensure that the percentage corresponds correctly to the WAN interface it indicates.

5 Depending on what you selected from the **Type** drop-down menu, one of these options display:

Type selection	Option
Basic Failover	**Preempt and failback to preferred interfaces when possible** Select to enable rank to determine the order of preemption. Selected by default.
Spill-over	**When bandwidth exceeds *BandwidthLimit* Kbit/s on *PrimaryInterface*, new flows will go to the alternate group members in Round Robin manner** Specify the bandwidth for the Primary in the field. If this value is exceeded, new flows are then sent to alternate group members according to the order listed in the **Selected** column. The default value is **0**.
Round Robin, **Spill-over**, and **Ratio**	**Use Source and Destination IP Address binding** The option is especially useful when using HTTP/HTTPS redirection or in a similar situation. For example, connection A and connection B need to be on the same WAN interface, the source and destination IP addresses in Connection A are the same as those for connection B, but a different service is being used. In this case, source and destination IP address binding is required to keep both the connections on the same WAN interface so that the transactions do not fail. This option is not selected by default.

6 Add, delete, and order member interfaces in the **Group Members: Select here:/Selected** lists. The use of the selected members in the **Selected** list depends on the **Type** selected:

- **Basic Failover: Interface Ordering:**
- **Round Robin: Interface Pool:**
- **Spill-over: Primary/Alt. Pool:**
- **Ratio: Interface Distribution:**

7 Add members by selecting a displayed interface from the **Group Members:** column, and then clicking **Add>>**.

8 You can order the entries in the **Selected** column by:

a Selecting an entry.

b Clicking **Up/Down**.

9 If you selected **Ratio**, instead of ordering the entries, you can specify the ratio of bandwidth for each interface. See Configuring Bandwidth as a Ratio on page 387.

IMPORTANT: To avoid problems associated with configuration errors, ensure that the percentage corresponds correctly to the WAN interface it indicates.

10 Enter a percentage of bandwidth to be assigned to an interface in the percent **(%)** field. The total bandwidth for all interfaces should add up to 100%. The total percentage of bandwidth allocated is displayed.

11 You can modify the ratio by clicking **Modify Ratio** or have the ratios adjusted automatically by clicking **Auto Adjust**.

12 You can delete members from the **Selected** column by:

a Selecting the displayed interface.

b Clicking **<<Remove**.

NOTE: The interface at the top of the list is the Primary.

The Interface Rank does not specify the operation performed on the individual member. The operation that is performed is specified by the Group Type.

13 Optionally, enter this setting:

- **Final Back-Up**—An entry in this setting is an interface of "last resort," that is, an interface that is used only when all other interfaces in the **Selected**: group are either unavailable or unusable. To specify a Final Back-Up interface, select an entry in the Group Members list, and then click the double right arrow. To remove a **Final Back-Up** interface, click the double left arrow.

14 Click **OK**.

Configuring Bandwidth as a Ratio

If **Ratio** is selected, **Add >>** is replaced by a percent **(%)** field and **Double Right Arrow** and **Up/Down Arrow** are replaced with **Auto Adjust**.

Enter a percentage of bandwidth to be assigned to the interface. The total percentage of bandwidth allocated is displayed.

IMPORTANT: To avoid problems associated with configuration errors, ensure that the percentage corresponds correctly to the WAN interface it indicates.

If multiple interfaces are selected, you can either:

- Click **Auto Adjust** to distribute the bandwidth equally among the interfaces.
- Enter a percentage of bandwidth to be assigned to each interface.

To modify the bandwidth percentage for an interface:

1 Select the interface in the **Selected** column.

2 Click **Modify Ratio**.

3 Enter a new percentage in the percent **(%)** field.

4 Click **Modify Ratio** again. The percentage for the bandwidth and the total bandwidth allocated are updated.

Probing Settings

When Logical probing is enabled, test packets can be sent to remote probe targets to verify WAN path availability. A new option has been provided to allow probing through the additional WAN interfaces: Alternate WAN #3 and Alternate WAN #4.

NOTE: VLANs for alternate WANs do not support QoS or VPN termination.

To configure the probing options for a specific Group:

1 Navigate to **MANAGE | Network > Failover & Load Balancing**

2 Click the **Configure** icon of the Group you wish to configure. The **Edit LB Group** dialog displays.

3 Click **Probing**.



4 Modify the following settings:

- **Check Interface every:** *n* **sec** —The interval of health checks in units of seconds. The default value is **5** seconds.
- **Deactivate Interface after:** *n* **missed intervals**—The number of failed health checks after which the interface sets to Failover. The default value is **6** seconds.
- **Reactivate Interface after:** *n* **successful intervals**—The number of successful health checks after which the interface sets to Available. The default value is **3** seconds.
- **Probe responder.global.SonicWall.com on all interfaces in this group**—Enable this checkbox to automatically set Logical/Probe Monitoring on all interfaces in the Group. When enabled, TCP probe packets are sent to the global SNWL host that responds to SNWL TCP packets, `responder.global.SonicWall.com`, using a target probe destination address of `204.212.170.23:50000`. When this checkbox is selected, the rest of the probe configuration enables built-in settings automatically. The same probe is applied to all four WAN Ethernet interfaces.

  NOTE: The Dialup WAN probe setting also defaults to the built-in settings.

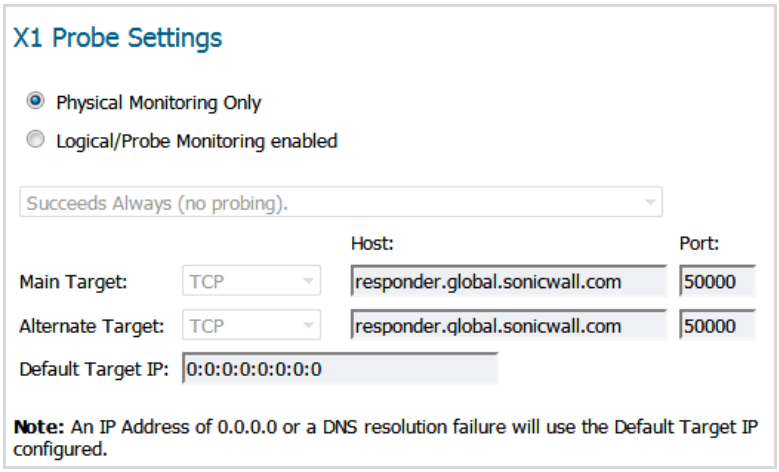
5 Click **OK**.

Configuring Probe Settings for Group Members

To configure the Group Member probe settings:

1 Navigate to **MANAGE | Network > Failover & Load Balancing**

2 Click the **Configure** icon of the Group member you wish to configure. The **Probe Settings** dialog displays.



3 Choose the type of probing to be done:

- **Physical Monitoring Only** (default; all other options are dimmed). Go to Step 9.
- **Logical/Probe Monitoring enabled** – all other options become available.

4 From **Logical/Probe Monitoring**, select when the probe succeeds:

- **Probe succeeds when either Main Target or Alternate Target responds.**
- **Probe succeeds when both Main Target and Alternate Target respond.**
- **Probe succeeds when Main Target responds.**
- **Succeeds Always (no probing)**. – Default; all other options are dimmed. Go to Step 9.

5 From **Main Target**, select:

- **Ping (ICMP)**
- **TCP** (default)

a In the **Main Target Host** field, enter the host name. The default is **responder.global.SonicWall.com**.

b In the **Main Target Port** field, enter the applicable port. The default is **50000**.

6 If Probe succeeds when Main Target Responds was selected, go to Step 8.

7 From the **Alternate Target** drop-down menu, select:

NOTE: The **Alternate Target** options are available only when **Probe succeeds when either Main Target or Alternate Target responds** or **Probe succeeds when both Main Target and Alternate Target respond** is selected for **Logical/Probe Monitoring** enabled.

- **Ping (ICMP)**
- **TCP** (default)

a In the **Alternate Target Host** field, enter the host name. The default is **responder.global.SonicWall.com**.

b In the **Alternate Target Port** field, enter the applicable port. The default is **50000**.

8 In the **Default Target IP** field, enter the IP address of the default target.

NOTE: This option is dimmed if S**ucceeds Always (no probing)** is selected for **Logical/Probe Monitoring enabled**.

An IP Address of `0.0.0.0` or a DNS resolution failure uses the configured Default Target IP.

9 Click **OK**.

17

Configuring Network Zones

- About Zones on page 390
  - How Zones Work on page 391
  - Predefined Zones on page 392
  - Security Types on page 392
  - Allow Interface Trust on page 393
  - Enabling SonicWall Security Services on Zones on page 393
  - Effect of Wireless and Non-Wireless Controller Modes on page 394
- **Network > Zones** on page 396
  - The Zone Settings Table on page 396
  - Adding a New Zone on page 397
  - Configuring a Zone for Guest Access on page 400
  - Configuring a Zone for Open Authentication and Social Login on page 402
  - Configuring a Zone for Captive Portal Authentication with Radius on page 403
  - Configuring a Zone for a Customized Policy Message on page 405
  - Configuring a Zone for a Customized Login Page on page 406
  - Configuring the WLAN Zone on page 407
  - Configuring the RADIUS Server on page 409
  - Configuring DPI-SSL Granular Control per Zone on page 410
  - Deleting a Zone on page 412

About Zones

A zone is a logical grouping of one or more interfaces designed to make management, such as the definition and application of Access Rules, a simpler and more intuitive process than following strict physical interface scheme. Zone-based security is a powerful and flexible method of managing both internal and external network segments, allowing the administrator to separate and protect critical internal network resources from unapproved access or attack.

A network security zone is simply a logical method of grouping one or more interfaces with friendly, user-configurable names, and applying security rules as traffic passes from one zone to another zone. Security zones provide an additional, more flexible, layer of security for the firewall. With the zone-based security, the administrator can group similar interfaces and apply the same policies to them, instead of having to write the same policy for each interface. For more information on configuring interfaces, see **Network > Interfaces** on page 268.

SonicOS zones allows you to apply security policies to the inside of the network. This allows you to do this by organizing network resources to different zones, and allowing or restricting traffic between those zones. This way, access to critical internal resources, such as payroll servers or engineering code servers, can be strictly controlled.

Zones also allow full exposure of the NAT table to allow you control over the traffic across the interfaces by controlling the source and destination addresses as traffic crosses from one zone to another. This means that NAT can be applied internally, or across VPN tunnels, which is a feature that users have long requested. Firewalls can also drive VPN traffic through the NAT policy and zone policy, because VPNs are now logically grouped into their own VPN zone.

Topics:

- How Zones Work on page 391
- Predefined Zones on page 392
- Security Types on page 392
- Allow Interface Trust on page 393
- Enabling SonicWall Security Services on Zones on page 393

How Zones Work

An easy way to visualize how security zones work is to imagine a large new building, with several rooms inside the building, and a group of new employees that do not know their way around the building. This building has one or more exits, which can be thought of as the WAN interfaces. The rooms within the building have one or more doors, which can be thought of as interfaces. These rooms can be thought of as zones inside each room are a number of people. The people are categorized and assigned to separate rooms within the building. People in each room going to another room or leaving the building, must talk to a door person on the way out of each room. This door person is the inter-zone/intra-zone security policy, and the door person's job to consult a list and make sure that the person is allowed to go to the other room, or to leave the building. If the person is allowed (for example, the security policy allows them in), they can leave the room through the door (the interface).

Upon entering the hallway, the person needs to consult with the hallway monitor to find out where the room is, or where the door out of the building is located. This hallway monitor provides the routing process because the monitor knows where all the rooms are located, and how to get in and out of the building. The monitor also knows the addresses of any of the remote offices, which can be considered the VPNs. If the building has more than one entrance/exit (WAN interfaces), the hallway monitor can direct people to use the secondary entrance/exit, depending upon how they have been told to do so (for example, only in an emergency, or to distribute the traffic in and out of the entrance/exits). This function can be thought of as WAN Load Balancing.

There are times that the rooms inside the building have more than one door, and times when there are groups of people in the room who are not familiar with one another. In this example, one group of people uses only one door, and another group uses the other door, even though groups are all in the same room. Because they also do not recognize each other, in order to speak with someone in another group, the users must ask the door person (the security policy) to point out which person in the other group is the one with whom they wish to speak. The door person has the option to not let one group of people talk to the other groups in the room. This is an example of when zones have more than one interface bound to them, and when intra-zone traffic is not allowed.

Sometimes, people want to visit remote offices, and people might arrive from remote offices to visit people in specific rooms in the building. These are the VPN tunnels. The hallway and doorway monitors check to see if this is allowed or not, and allow traffic through. The door person can also elect to force people to put on a costume before traveling to another room, or to exit, or to another remote office. This hides the true identity of the person, masquerading the person as someone else. This process can be thought of as the NAT policy.

Predefined Zones

The predefined zones on your firewall depend on the device. The predefined security zones on the SonicWall Security Appliance are not modifiable:

This zone	Has this function
DMZ	Normally used for publicly accessible servers and can consist of one to four interfaces, depending on your network design.
LAN	Cn consist of multiple interfaces, depending on your network design. Even though each interface has a different network subnet attached to it, when grouped together, they can be managed as a single entity.
MGMT	Used for appliance management and includes only the MGMT interface. Interfaces in other zones can also be enabled for SonicOS management, but the MGMT zone/interface provides the added security of a separate zone just for management.
MULTICAST	Provides support for IP multicasting, which is a method for sending IN packets from a single source simultaneously to multiple hosts.
SSLVPN	Used for secure remote access using the SonicWall NetExtender client.
VPN	A virtual zone used for simplifying secure, remote connectivity.
WLAN	Provides support to SonicWall SonicPoints and SonicWaves. When assigned to the Opt port, it enforces SonicPoint Enforcement, automatically dropping all packets received from non-SonicPoint devices. The WLAN zone supports: • Discovery Protocol (SDP) to automatically poll for and identify attached SonicPoints and SonicWaves • SonicWall Simple Provisioning Protocol to configure SonicPoints and SonicWaves using profiles • Wireless and guest service configurations
WAN:	Can consist of multiple interfaces. If you are using the Security Appliance's WAN failover capability, you need to add the second Internet interface to the WAN zone.

NOTE: Even though you can group interfaces together into one security zone, this does not preclude you from addressing a single interface within the zone.

Security Types

Each zone has a security type, which defines the level of trust given to that zone:

Trusted	Provides the highest level of trust—meaning that the least amount of scrutiny is applied to traffic coming from trusted zones. Trusted security can be thought of as being on the LAN (protected) side of the Security Appliance. The LAN zone is always Trusted.
Management	Unique to the MGMT zone and MGMT interface and also provides the highest level of trust.
Encrypted	Used exclusively by the VPN and SSLVPN zones. All traffic to and from an Encrypted zone is encrypted.
Wireless	Applied to the WLAN zone or any zone where the only interface to the network consists of SonicWall SonicPoint and SonicWave devices. Wireless security type is designed specifically for use with SonicPoints and SonicWaves. Placing an interface in a Wireless zone activates SDP (SonicWall Discovery Protocol) and SSPP (SonicWall Simple Provisioning Protocol) on that interface for automatic discovery and provisioning of SonicPoints and SonicWaves. Only traffic that passes through a SonicPoint or SonicWaveis allowed through a Wireless zone; all other traffic is dropped.

Public	Offers a higher level of trust than an Untrusted zone, but a lower level of trust than a Trusted zone. Public zones can be thought of as being a secure area between the LAN (protected) side of the Security Appliance and the WAN (unprotected) side. The DMZ, for example, is a Public zone because traffic flows from it to both the LAN and the WAN. By default, traffic from DMZ to LAN is denied, but traffic from LAN to ANY is allowed. This means only LAN-initiated connections have traffic between DMZ and LAN. The DMZ only has default access to the WAN, not the LAN.
Untrusted	Represents the lowest level of trust. It is used by both the WAN and the virtual Multicast zone. An Untrusted zone can be thought of as being on the WAN (unprotected) side of the Security Appliance. By default, traffic from Untrusted zones is not permitted to enter any other zone type without explicit rules, but traffic from every other zone type is permitted to Untrusted zones.

Allow Interface Trust

The **Allow Interface Trust** setting in the **Add Zone** dialog automates the creation of Access Rules to allow traffic to flow between the interface of a zone instance. For example, if the LAN zone has both the **LAN** and **X3** interfaces assigned to it, checking **Allow Interface Trust** on the LAN zone creates the necessary Access Rules to allow hosts on these interfaces to communicate with each other.

Enabling SonicWall Security Services on Zones

You can enable SonicWall Security Services for traffic across zones. For example, you can enable SonicWall Intrusion Prevention Service for incoming and outgoing traffic on the WLAN zone to add more security for internal network traffic. You can enable these SonicWall Security Services on zones:

Enforce Content Filtering Service	Enforces content filtering on multiple interfaces in the same Trusted and Public security types for WLAN zones.
Enforce Client Anti-Virus Service	Enforces anti-virus protection on multiple interfaces in the same Trusted and Public security types for WLAN zones.
Enable Gateway Anti-Virus	Enforces gateway anti-virus protection on multiple interfaces in the same Trusted and Public security types for WLAN zones.
Enable IPS	Enforces intrusion detection and prevention on multiple interfaces in the same Trusted and Public security types for WLAN zones.
Enable App Control Service	Enforces application control policy services on multiple interfaces in the same Trusted and Public security types for WLAN zones.
Enable Anti-Spyware Service	Enforces anti-spyware detection and prevention on multiple interfaces in the same Trusted and Public security types for WLAN zones.
Enforce Global Security Clients	Enforces Global Security Client (GSC) protection on multiple interfaces in the same Trusted and Public security types for WLAN zones.
Create Group VPN	Creates a GroupVPN policy for the zone, which is displayed in the VPN Policies table on **MANAGE | Connectivity > VPN > Base Settings**. You can customize the GroupVPN policy on **VPN > Base Settings**. If you clear **Create Group VPN**, the GroupVPN policy is removed from **VPN > Base Settings**. For more information about creating VPN policies, see *SonicOS 6.5 Connectivity*.

Enable SSL Control	Enables SSL Control on the zone. All new SSL connections initiated from that zone are now subject to inspection. SSL Control must first be enabled globally on **MANAGE | Firewall Settings | SSL Control**. For more information about SSL Control, see *SonicOS 6.5 Security Configuration*.
Enable SSLVPN Access	Enables SSLVPN secure remote access on the zone.

Effect of Wireless and Non-Wireless Controller Modes

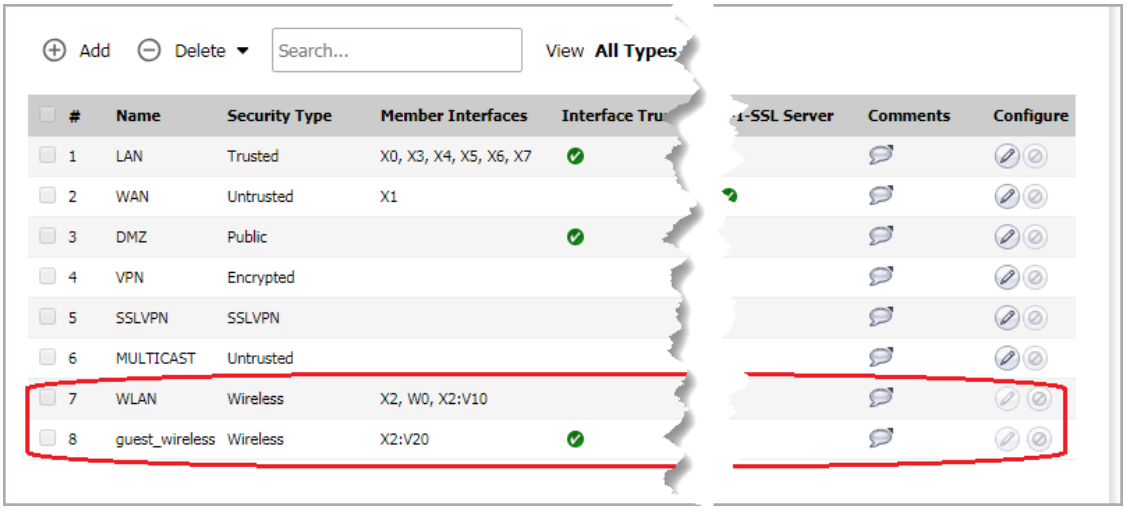
Topics:

- Effects of Enabling Non-Wireless Controller Mode on page 394
- Effects of Enabling Wireless Controller Mode on page 395

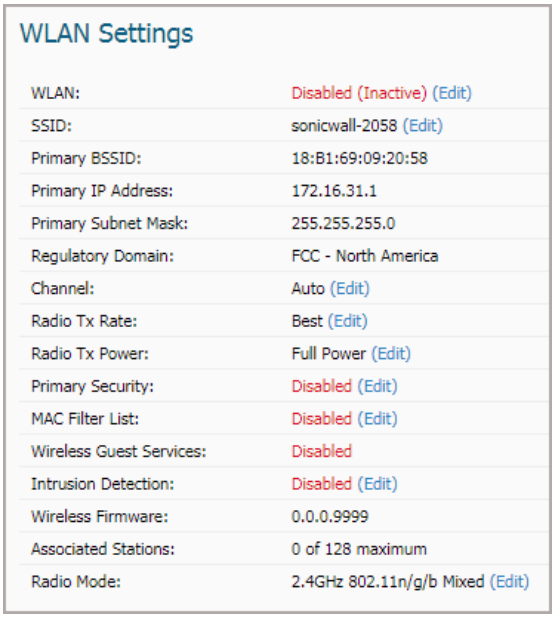
Effects of Enabling Non-Wireless Controller Mode

Enabling Nnn-Wireless Controller Mode affects the **Network > Zones** page. Attempts to enable or delete the affected features are denied.

- The **Edit** and **Delete** icons for wireless zones become dimmed on the **MANAGE | System Setup > Network > Zones** page.



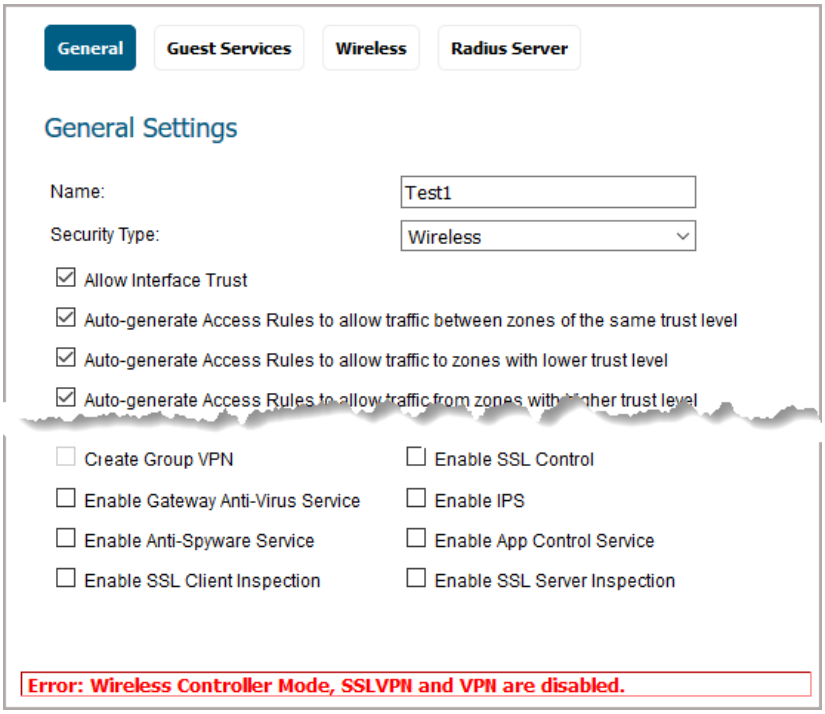
- Internal wireless zones are disabled.



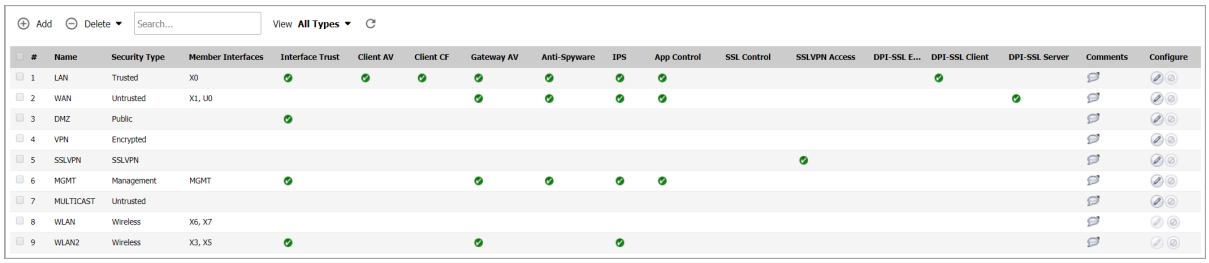
Effects of Enabling Wireless Controller Mode

Enabling Wireless Controller Mode affects the **Network > Zones** page. Attempts to enable or delete the affected features are denied.

- The **Edit** and **Delete** icons for VPN and SSL VPN zones become dimmed on the **MANAGE | System Setup > Network > Zones** page.
- Any attempt to enable a zone with VPN and/or SSL VPN results in an error.



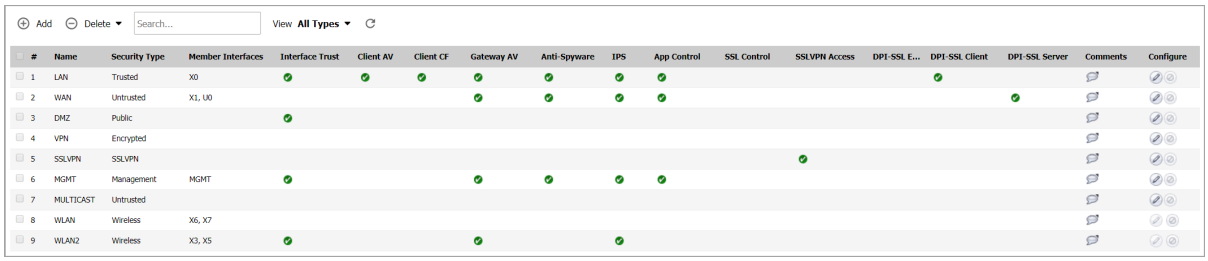
Network > Zones



- The Zone Settings Table on page 396
- Adding a New Zone on page 397
- Deleting a Zone on page 412
- Configuring a Zone for Guest Access on page 400
- Configuring a Zone for Open Authentication and Social Login on page 402
- Configuring a Zone for Captive Portal Authentication with Radius on page 403
- Configuring a Zone for a Customized Policy Message on page 405
- Configuring a Zone for a Customized Login Page on page 406
- Configuring the WLAN Zone on page 407

The Zone Settings Table

The **Zone Settings** table displays a listing of all the SonicWall Security Appliance's default predefined zones as well as any zones you create. The table displays the following status information about each zone configuration:



Name	Name of the zone. The predefined **LAN**, **WAN**, **WLAN**, **VPN**, **SSLVPN**, **MGMT**, **MULTICAST**, and **Encrypted** zone names cannot be changed.
Security Type	Security type: **Trusted**, **Untrusted**, **Public**, **Wireless**, or **Encrypted**.
Member Interfaces	Interfaces that are members of the zone.
Interface Trust	Check mark indicates the **Allow Interface Trust** setting is enabled for the zone.
Client AV	Check mark indicates SonicWall Client Anti-Virus is enabled for traffic coming in and going out of the zone. SonicWall Client Anti-Virus manages an anti-virus client application on all clients on the zone.
Client CF	Check mark indicates Client Content Filtering services are enabled.
Gateway AV	Check mark indicates SonicWall Gateway Anti-Virus is enabled for traffic coming in and going out of the zone. SonicWall Gateway Anti-Virus manages the anti-virus service on the firewall.

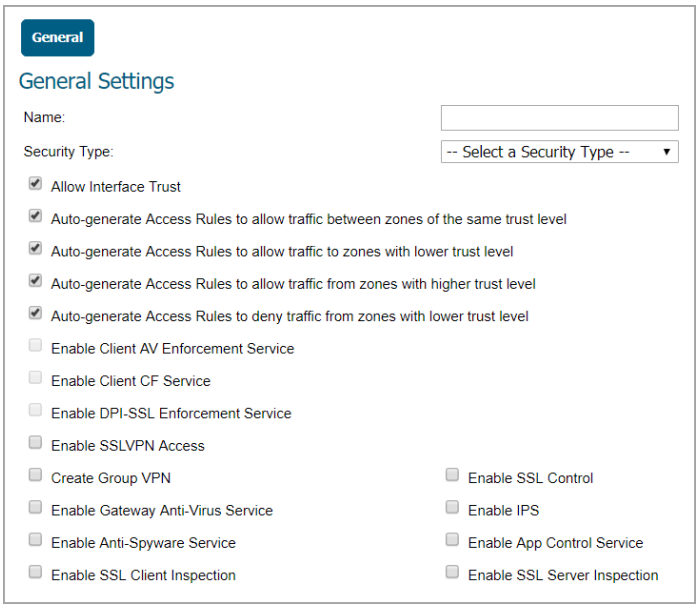
Anti-Spyware	Check mark indicates SonicWall Anti-Spyware detection and prevention is enabled for traffic through interfaces in the zone.
IPS	Check mark indicates SonicWall Intrusion Prevention Service is enabled for traffic coming in and going out of the zone.
App Control	Check mark indicates App Control Service is enabled for traffic coming in and going out the zone.
SSL Control	Check mark indicates SSL Control is enabled for traffic coming in and going out the zone. All new SSL connections initiated from that zone is now subject to inspection.
SSL VPN Access	Check mark indicates SSL VPN secure remote access is enabled for traffic coming in and going out the zone.
DPI-SSL Client	Check mark indicates granular DPI-SSL on a per-zone basis rather than a global basis for DPI-SSL clients.
DPI-SSL Server	Check mark indicates granular DPI-SSL on a per-zone basis rather than global basis for DPI-SSL servers.
Comments	Mousing over the **Comment** icon displays the comment entered when the Zone was configured.
Configure	Clicking the **Edit** icon displays the **Edit Zone** dialog. Clicking the **Delete** icon deletes the zone. The delete icon is dimmed for the predefined zones; you cannot delete these zones.

Adding a New Zone

To add a new zone:

1 Navigate to **MANAGE | System Setup > Network > Zones**.

2 Click the **Add** icon. The **Add Zone** dialog displays.



3 Type a name for the new zone in the **Name** field.

4 From **Security Type**, select:

Trusted	Zones with the highest level of trust, such as internal LAN segments.
Public	Zones with a lower level of trust requirements, such as a DMZ interface.
Wireless	WLAN interface.
SSLVPN	Interfaces on which Content Filtering, Client AV enforcement, and Client CF services are enabled. **NOTE:** Selecting this security type disables the **Enable SSLVPN Access** and **Create Group VPN** options on this dialog.

5 To allow intra-zone communications, select **Allow Interface Trust**. An Access Rule allowing traffic to flow between the interfaces of a Zone instance is created automatically. This option is selected by default.

6 To have SonicOS automatically generate access rules to allow traffic between this zone and other zones of equal trust, select **Auto-generate Access Rules to allow traffic between zones of the same trust level**. For example, `CUSTOM_LAN -> CUSTOM _LAN` or `CUSTOM_LAN -> LAN`. This option is selected by default.

NOTE: For this option and the following Access Rules options, see *SonicOS 6.5 Policies* for information about Access Rules.

7 To have SonicOS automatically generate access rules to allow traffic between this zone and other zones of lower trust, select **Auto-generate Access Rules to allow traffic to zones with lower trust level**. For example, `CUSTOM_LAN -> WAN` or `CUSTOM_LAN -> DMZ`. This option is selected by default.

8 To have SonicOS automatically generate access rules to allow traffic between this zone and other zones of higher trust, select **Auto-generate Access Rules to allow traffic from zones with higher trust level**. For example, `LAN -> CUSTOM_DMZ` or `CUSTOM_LAN -> CUSTOM_DMZ`. This option is selected by default.

9 To have SonicOS automatically generate access rules to deny traffic between this zone and zones of lower trust, select **Auto-generate Access Rules to deny traffic from zones with lower trust level**. For example, `WAN -> CUSTOM_LAN` or `DMZ -> CUSTOM_LAN`. This option is selected by default.

10 To enforce managed Client Anti-Virus protection on clients connected to multiple interfaces in the same Trusted, Public, or WLAN zones using the Client Anti-Virus client on your network hosts, select **Enable Client AV Enforcement Service**. This option is not selected by default.

NOTE: This option is dimmed and unavailable until you select a security type from **Security Type**.

NOTE: For this option and the following Security Services options, see *SonicOS 6.5 Security Configuration* for more information about these services.

11 To enforce managed Client Content Filtering on clients connected to multiple interfaces in the same Trusted, Public, or WLAN zones using the Client CF client on your network hosts, select **Enable Client CF Service**. This option is not selected by default.

NOTE: This option is dimmed and unavailable until you select a security type from **Security Type**.

12 To enforce enhanced NGAV (Next Generation AV) such as DPI-SSL Enforcement or SentinelOne AV enforcement, select **Enable DPI-SSL Enforcement Service**. This option is not selected by default. For more information about NGAV, see *SonicOS 6.5 Security Configuration*.

13 To enable SSLVPN secure remote access on the zone, select **Enable SSLVPN Access**. This option is not selected by default.

NOTE: This option is dimmed if **SSLVPN** is selected for **Security Type**.

14 To create a SonicWall Group VPN Policy for this zone automatically, select **Create Group VPN**. You can customize the Group VPN Policy in **MANAGE | Connectivity > VPN > Settings**. This option is not selected

by default. This option is available UNTIL **SSLVPN** is selected for **Security Type**, but after the Security Type is changed to one of the other types, it remains dimmed and unavailable.

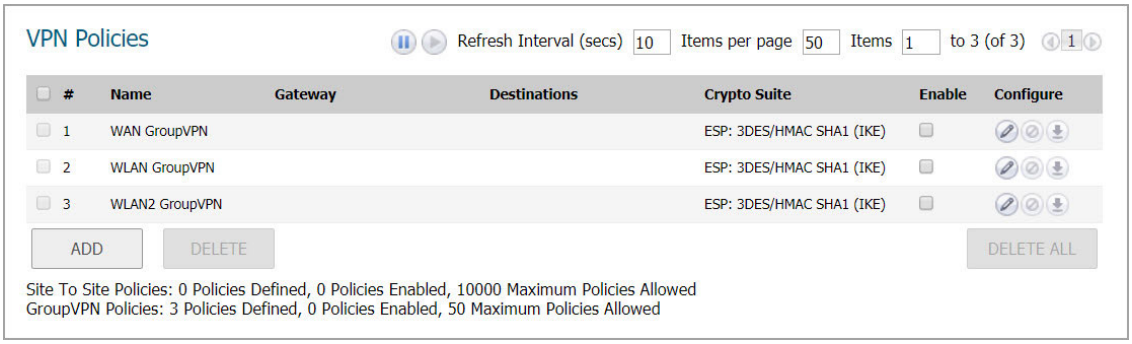
CAUTION: Disabling Create Group VPN removes any corresponding Group VPN policy.

NOTE: This option is dimmed if **SSLVPN** is selected for **Security Type**.

For this and other connectivity options, see *SonicOS 6.5 Connectivity* for more information.

Disabling Group VPN for WAN/WLAN VPN policies, deletes all VPN policies. Re-enabling the **Create Group VPN** option automatically creates a new, enabled VPN policy. Disabling VPN policies globally does not also delete auto-rules. If you do not want to VPN polices at all, globally disable VPN, and then delete all policies that correlate with VPN.

GroupVPN policies appear in the **VPN Policies** table located in **MANAGE | Connectivity | VPN > Base Settings**. WAN/WLAN GroupVPN policies are disabled by default when the firewall is booted with the factory default.



15 To enable SSL Control on the zone, select **Enable SSL Control**. All new SSL connections initiated from that zone are now subject to inspection. This option is not selected by default.

NOTE: SSL Control must first be enabled globally on **MANAGE | Security Configuration > Firewall > SSL Control**.

16 To enforce gateway anti-virus protection on your Security Appliance for all clients connecting to this zone, select **Enable Gateway Anti-Virus Service**. SonicWall Gateway Anti-Virus manages the anti-virus service on the Security Appliance. This option is not selected by default.

17 To enforce intrusion detection and prevention on multiple interfaces in the same Trusted, Public, or WLAN zones. select **Enable IPS**. This option is not selected by default.

18 To enforce anti-spyware detection and prevention on multiple interfaces in the same Trusted or Public security type for WLAN zones, select **Enable Anti-Spyware Service**. This option is not selected by default.

19 To enforce application control policy services on multiple interfaces in the same Trusted or Public security type for WLAN zones, select **Enable App Control Service**. This option is not selected by default. For more information about App Control, see *SonicOS 6.5 Policies*.

20 To enable granular DPI-SSL on a per-zone basis rather than globally for DPI-SSL clients, select **Enable SSL Client Inspection**. This option is not selected by default.

21 To enable granular DPI-SSL on a per-zone basis rather than globally for DPI-SSL servers, select **Enable SSL Server Inspection**. This option is not selected by default.

22 Click **OK**. The new zone is now added to the Security Appliance.

Configuring a Zone for Guest Access

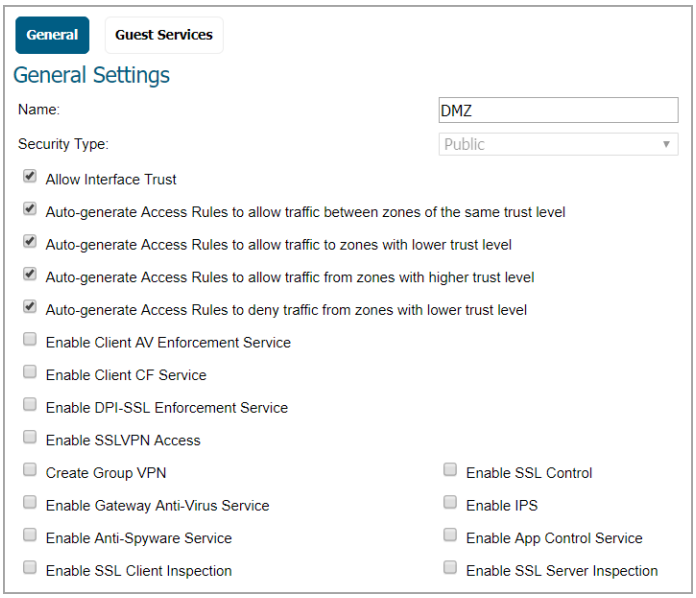
IMPORTANT: You cannot configure an Untrusted, Encrypted, SSL VPN, or Management zone for guest access.

SonicWall User Guest Services provides an easy solution for creating wired and wireless guest passes and/or locked-down Internet-only network access for visitors or untrusted network nodes. This functionality can be extended to wireless or wired users on the WLAN, LAN, DMZ, or public/semi-public zone of your choice.

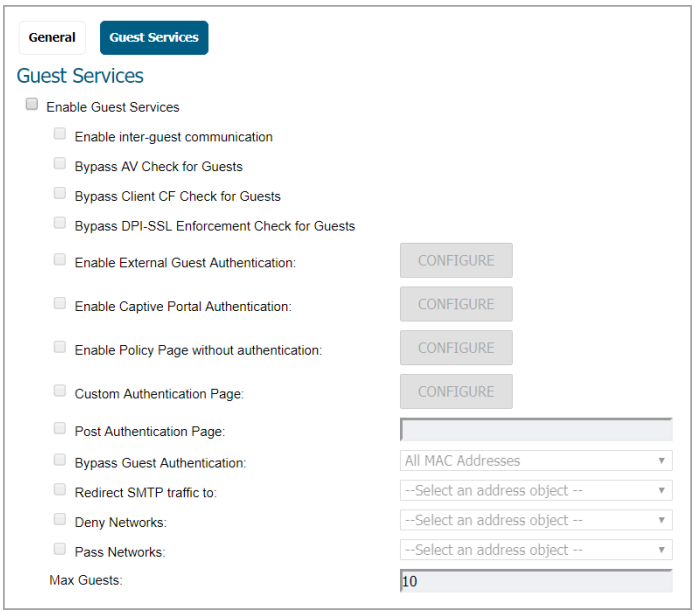
To configure Guest Services feature:

1 Navigate to **MANAGE | System Setup > Network > Zones**.

2 Click **Edit** for the zone you wish to add Guest Services to. The **Edit Zone** dialog displays.



3 Click **Guest Services**. Only the **Enable Guest Services** option is available.



4 Click **Enable Guest Services**. All other options become available, but are not selected by default.

5 Select from the following configuration options for Guest Services:

Enable inter-guest communication	Allows guests to communicate directly with other users who are connected to this zone.
Bypass AV Check for Guests	Allows guest traffic to bypass Anti-Virus protection.
Bypass Client CF Check for Guests	Allows guest traffic to bypass Client CF enforcement.
Bypass DPI-SSL Enforcement Check for Guests	Allows guest traffic to bypass DPI-SSL enforcement.
Enable External Guest Authentication	Requires guests connecting from the device or network you select to authenticate before gaining access. Selecting this option makes **CONFIGURE** available. Clicking **CONFIGURE** displays the **External Guest Authentication** dialog. For information about configuring this option, see Configuring Social Login in SonicOS on page 801. **NOTE:** When this option is selected, the following four options become dimmed and unavailable.
Enable Captive Portal Authentication	Allows you to create a customized login page with RADIUS authentication. Selecting this option makes **CONFIGURE** available. Clicking **CONFIGURE** displays the **Customize Login Page** dialog. For information about configuring this option, see Configuring a Zone for Captive Portal Authentication with Radius on page 403.
Enable Policy Page without authentication	Directs users to a guest services usage policy page when they first connect to a SonicPoint or SonicWave in the WLAN zone. Guest users are authenticated by accepting the policy instead of providing a user name and password. Selecting this option makes **CONFIGURE** available. To set up an HTML customizable policy usage page, click **CONFIGURE**. The **Customize Policy Message** dialog displays. For information about configuring this option, see Configuring a Zone for a Customized Policy Message on page 405.
Custom Authentication Page	Redirects users to a custom authentication page when they first connect to the network. Selecting this option makes **CONFIGURE** available. To set up the custom authentication page, click **CONFIGURE** to display the **Customize Login Page** dialog. For information about configuring this option, see Configuring a Zone for a Customized Login Page on page 406.
Post Authentication Page	Directs users to the specified page immediately after successful authentication. Selecting this option makes its field available. Enter a URL for the post-authentication page in the field.
Bypass Guest Authentication	Allows the Guest Services feature to integrate into environments already using some form of user-level authentication. This feature automates the authentication process, allowing wireless users unrestricted wireless Guest Services without requiring authentication. When selected, this option's drop-down menu becomes available; select: • **All MAC Addresses** (default) • An Address Object • An Address Group • **Create new MAC object** – Displays the **Add Address Object** dialog.[a] **NOTE:** This feature should only be used when unrestricted Guest Service access is desired, or when another device upstream is enforcing authentication.

Redirect SMTP traffic to	Redirects SMTP traffic incoming on this zone to an SMTP server you specify. When selected, this option's drop-down menu becomes available; select: • An Address Object • **Create new address object** – Displays the **Add Address Object** dialog.[a]
Deny Networks	Blocks traffic to the networks you name. When selected, this option's drop-down menu becomes available; select: • An Address Object • An Address Object group • **Create new address object** [a] • **Create new address object group** [a]
Pass Networks	Allows traffic through the Guest Service-enabled zone to the selected networks automatically. When selected, this option's drop-down menu becomes available; select: • An Address Object • An Address Object group • **Create new address object** [a] • **Create new address object group** [a] **NOTE:** Displays the **Add Address Object** dialog.
Max Guests	Specifies the maximum number of guest users allowed to connect to this zone. The minimum number is 1, the maximum number is 4500, and the default setting is **10**.
Wireless Zone Guest Services Options	Displays only for the WLAN zone or for a custom zone with a **Security Type** of **Wireless**.
Enable Dynamic Address Translation	Grants access to non-DHCP guests. This option is not selected by default.

a. For information about creating Address Objects and Address Object Groups, see *SonicOS Policies*.

6 Click **OK** to apply these settings to this zone.

Configuring a Zone for Open Authentication and Social Login

SonicOS supports Open Authentication (OAuth) and Social Login:

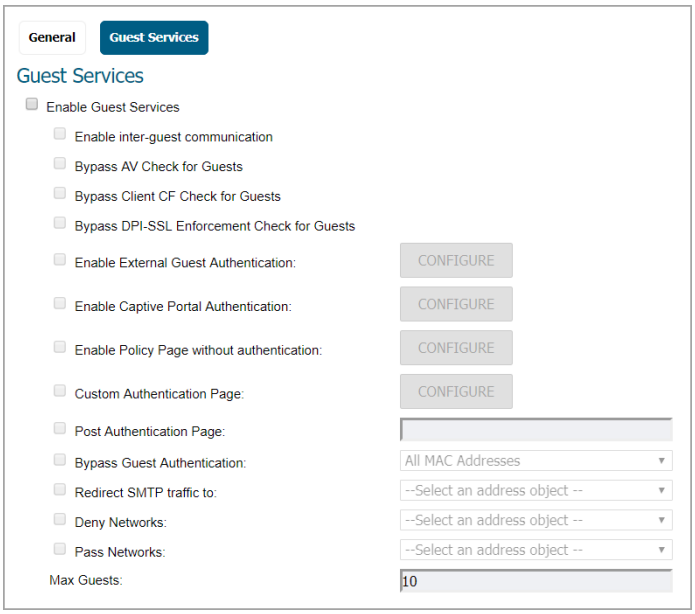
- Oauth assists users in sharing data between applications.
- Social Login simplifies the login process for various social media

To use these features, you create a zone, as described in the Configuring Open Authentication, Social Login, and LHM on page 792.

Configuring a Zone for Captive Portal Authentication with Radius

To configure captive portal authentication with RADIUS:

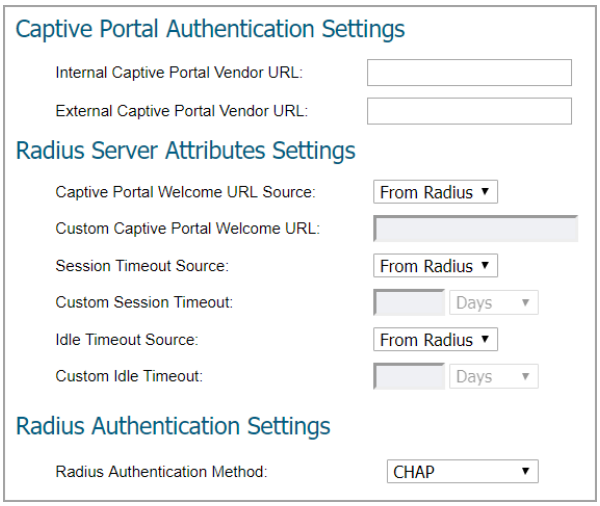
1 On the **Add/Edit Zone** dialog, click **Guest Services**.



2 Select **Enable Guest Services**. The options become available.

3 Select **Enable Captive Portal Authentication**. **CONFIGURE** becomes available.

4 Click **CONFIGURE**. The **Customize Login Page** dialog displays.



5 In the **Captive Portal Authentication Settings** section:

   a Enter the internal captive portal vendor's URL in the **Internal Captive Portal Vendor URl** field.

   b Enter the external captive portal vendor's URL in the **External Captive Portal Vendor URl** field.

6 In the **Radius Server Attributes Settings** section:

   a Select the source for the captive portal welcome URL from **Captive Portal Welcome URL Source**:

- **From Radius** (default); go to Step c.
- **Custom**; the next option becomes available

b Enter the welcome URL in the **Custom Captive Portal Welcome URL** field.

c Select the source for the session timeout limit from Session Timeout Source:

- **From Radius** (default); go to Step f.
- **Custom**; the next option becomes available

d Select the type of session timeout duration from **Custom Session Timeout**:

- **Minutes**
- **Hours**
- **Days** (default)

e Enter the limit in the field.

f Select the source for the idle timeout from **Idle Timeout Source**:

- **From Radius** (default); go to Step 7.
- **Custom**; the next option becomes available

g Select the type of idle timeout duration from **Custom Session Timeout**:

- **Minutes**
- **Hours**
- **Days** (default)

h Enter the limit of the duration in the field.

7 In the **Radius Authentication Settings** section, select the authentication method from **Radius Authentication Method**:

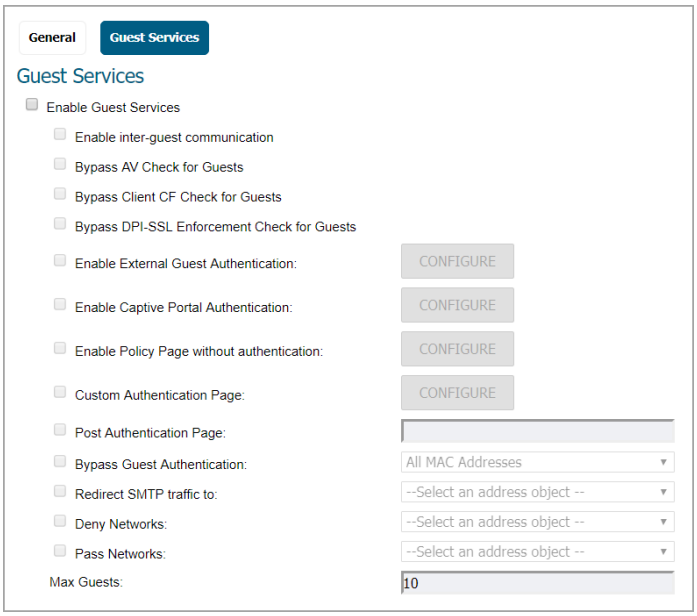
- **CHAP** (default)
- **PAP – Encrypted**
- **PAP – ClearText**

8 Click **OK**.

Configuring a Zone for a Customized Policy Message

To configure a customized policy message:

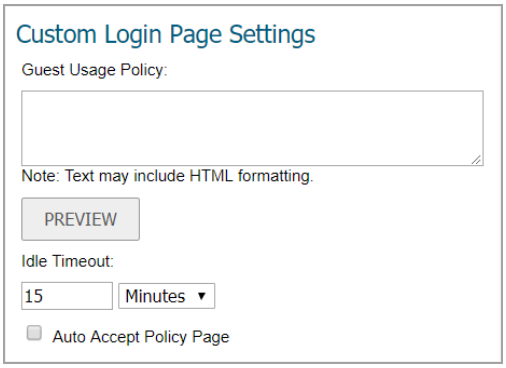
1 On the **Add/Edit Zone** dialog, click **Guest Services**.



2 Select **Enable Guest Services**. The options become available.

3 Select **Enable Policy Page without authentication**. **CONFIGURE** becomes available.

4 Click **CONFIGURE**. The **Customize Login Page** dialog displays.



5 Enter your policy for guest usage in the **Guest Usage Policy** field. The text may include HTML formatting.

6 To preview your policy message, click **PREVIEW**.

7 To specify an idle timeout, enter the timeout value in the **Idle Timeout** field.

8 Select the type of timeout:

- **Seconds**
- **Minutes** (default)
- **Hours**

- **Days**

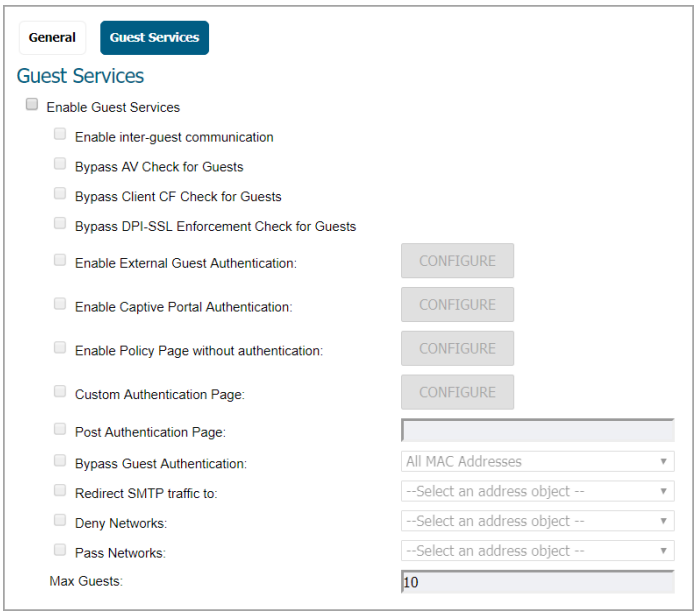
9 Select **Auto Accept Policy Page**. This option is not selected by default.

10 Click **OK**.

Configuring a Zone for a Customized Login Page

To configure a customized login page:

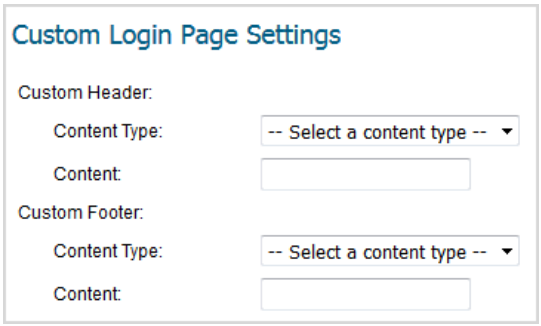
1 On the **Add/Edit Zone** dialog, click **Guest Services**.



2 Select **Enable Guest Services**. The options become available.

3 Select **Custom Authentication Page**. **CONFIGURE** becomes available.

4 Click **CONFIGURE**. The **Customize Login Page** dialog displays.



5 For **Custom Header**, select from **Content Type**:

- **URL**
- **Text**

6 Enter the URL or text in the **Content** field.

7 For **Custom Footer**, select from **Content Type**:

- **URL**

- **Text**

8 Enter the URL or text in the **Content** field.

9 Click **OK**.

Configuring the WLAN Zone

1 Navigate to **MANAGE | System Setup > Network > Zones**.

2 If you are configuring:

- A new zone, click **Add...**
- An existing zone, click the **Edit** icon for the WLAN zone.

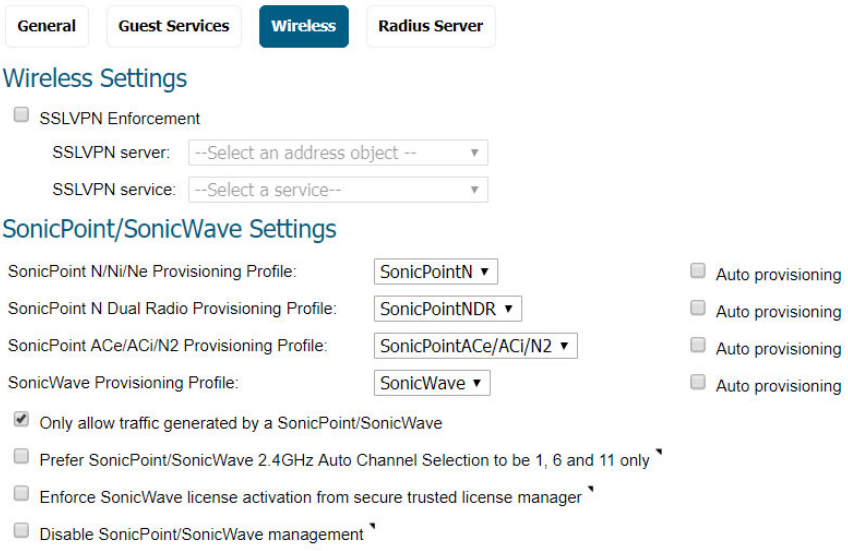
The **Add/Edit Zone** dialog displays.

> **NOTE:** Depending on the zone, there also may be views available for **Guest Services**, **Wireless**, and **Radius Server**.
>
> How to configure the **General** view is described in Adding a New Zone on page 397.

3 If creating a new zone, select **Wireless** from **Security Type**. **Guest Services**, **Wireless,** and **Radius Server** appear.

4 To automate the creation of Access Rules to allow traffic to flow between the interfaces of a zone instance, select **Allow Interface Trust**. For example, if the LAN zone has both the **LAN** and **X3** interfaces assigned to it, checking **Allow Interface Trust** on the LAN zone creates the necessary Access Rules to allow hosts on these interfaces to communicate with each other. This option is not selected by default.

5 Click **Wireless**.



6 In the **Wireless Settings** section, to require that all traffic that enters into the WLAN zone be authenticated through a SonicWall SSL VPN appliance, select **SSLVPN Enforcement**. Selecting this option makes the following two options available. This option is not selected by default.

7 From **SSL VPN server**, select an address object to direct traffic to the SonicWall SSL VPN appliance or create a new one. For information about creating Address Objects and Address Object Groups, see *SonicOS 6.5 Policies*.

8 From **SSL VPN service**, select the service or group of services to allow clients authenticated through the SSL VPN.

9 In the **SonicPoint/SonicWave Settings** section, select the **SonicPoint/SonicWave Provisioning Profile** to apply to all SonicPoints/SonicWaves connected to this zone. Whenever a SonicPoint/SonicWave connects to this zone, it is provisioned automatically by the settings in the SonicPoint/SonicWave Provisioning Profile, unless you have individually configured it with different settings. For information SonicPoint/SonicWave provisioning profiles, see *SonicOS 6.5 Connectivity*.

> **NOTE:** For the following four settings, optionally select **Auto provisioning** to allow SonicPoints/SonicWaves attached to the profile to be provisioned automatically when the profile is modified. This option is not selected by default.

10 Select the **SonicPointN/Ni/Ne Provisioning Profile** when you want to apply to all SonicPointN/Ni/Nes connected to this zone. Whenever a SonicPointN/Ni/Ne connects to this zone, it is automatically provisioned by the settings in the SonicPoint Provisioning Profile, unless you have individually configured it with different settings. The default provisioning profile is **SonicPointN**.

11 Select **SonicPoint N Dual Radio Provisioning Profile** when you want to apply to all SonicPointNDRs connected to this zone. Whenever a SonicPointNDR connects to this zone, it is automatically provisioned by the settings in the SonicPointNDR Provisioning Profile, unless you have individually configured it with different settings. The default provisioning profile is **SonicPointNDR**.

12 Select **SonicPointACe/ACi/N2 Provisioning Profile** when you want to apply to all SonicPointACe/ACi/N2s connected to this zone. Whenever a SonicPointACe/ACi/N2 connects to this zone, it is automatically provisioned by the settings in the SonicPointACe/ACi/N2 Provisioning Profile, unless you have individually configured it with different settings. The default provisioning profile is **SonicPointACe/ACi/N2**.

13 Select **SonicWave Provisioning Profile** when you want to apply to all SonicPointNDRs connected to this zone. Whenever a SonicPointNDR connects to this zone, it is automatically provisioned by the settings in the SonicPointNDR Provisioning Profile, unless you have individually configured it with different settings. The default provisioning profile is **SonicWave**.

14 Check **Only allow traffic generated by a SonicPoint/SonicWave** to allow only traffic from SonicWall SonicPoints to enter the WLAN zone interface. This allows maximum security of your WLAN. This option is selected by default. Clear this option if you want to allow any traffic on your WLAN zone regardless of whether the traffic is from a wireless connection.

> **TIP:** To allow any traffic on your WLAN zone regardless of whether it is from a wireless connection, clear **Only allow traffic generated by a SonicPoint / SonicPointN**.

> **NOTE:** For Guest Services configuration information, see Configuring a Zone for Guest Access on page 400.
>
> For RADIUS server configuration information, see Configuring the RADIUS Server on page 409.

15 Optionally, select **Prefer SonicPoint/SonicWave 2.4Hz Auto Channel Selection to be 1.6 and 11 only**. This option is not selected by default.

> **IMPORTANT:** Enable this option only when SonicPointN/AC 2.4Hz Auto Channel selection is preferred to be 1, 6, and 11.

16 Select **Enforce SonicWave license activation from secure trusted license manager**.

> **CAUTION: This option enforces license activation from a secure trusted license manager; manual license keyset input is not allowed. Change this setting only under the direction of Technical Support.**

17 Select **Disable SonicPoint/SonicWave management** to disable all management capabilities on this WLAN.

18 To:

- Configure the RADIUS server, go to Configuring the RADIUS Server on page 409.
- Apply these settings to the WLAN zone., click **OK**.

Configuring the RADIUS Server

1 Navigate to **MANAGE | System Setup > Network > Zones**.

2 If you are configuring:

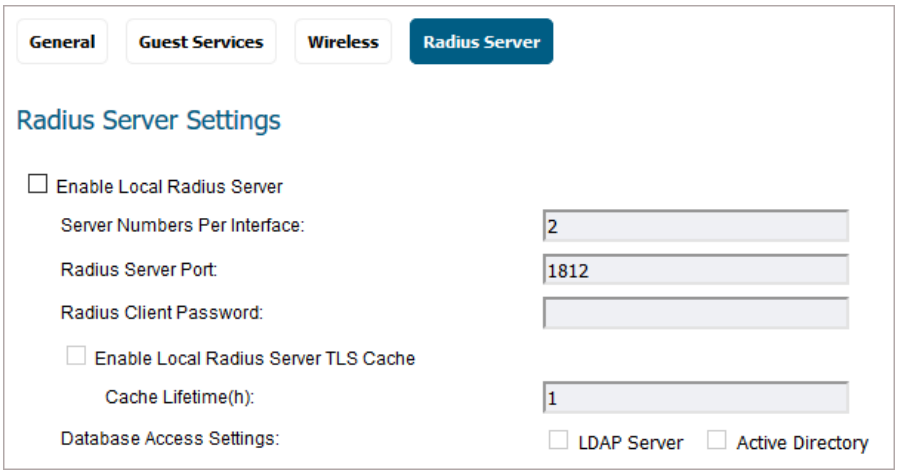
- A new zone, click **Add...**
- An existing zone, click the **Edit** icon for the WLAN zone.

The **Add/Edit Zone** dialog displays.

> **NOTE:** Depending on the zone, there also may be views available for **Guest Services**, **Wireless**, and **Radius Server**.
>
> How to configure the **General** view is described in Adding a New Zone on page 397.

3 If creating a new zone, select **Wireless** from **Security Type**. **Guest Services**, **Wireless,** and **Radius Server** appear.

4 Click **Radius Server**.



5 Click **Enable Local Radius Server**. The other options become available.

6 Enter the number of RADIUS servers numbers per interface in **Server Numbers Per Interface**. The minimum number is 1, the maximum is 512, and the default is **2**.

7 Enter the port for the RADIUS server in the **Radius Server Port** field. The default is **1812**.

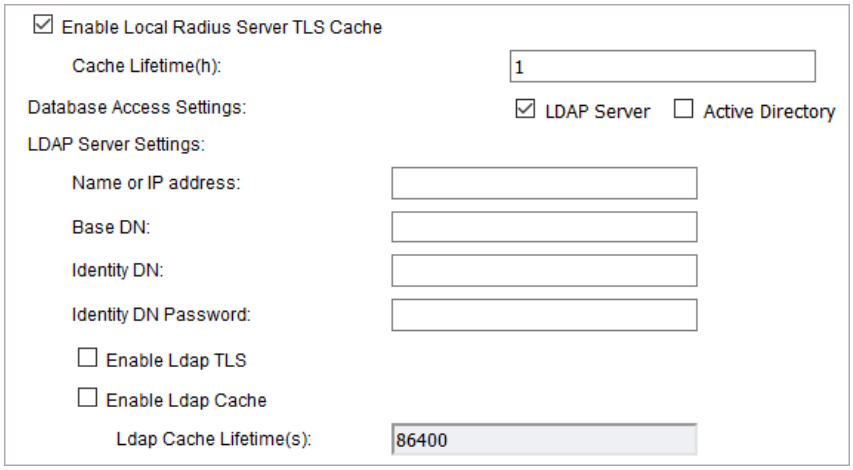
8 Enter the password for the RADIUS client in the **Radius Client Password** field.

9 Optionally, select **Enable Local Radius Server TLS Cache** lifetime. This option is not selected by default. The **Cache Lifetime(h)** field becomes available.

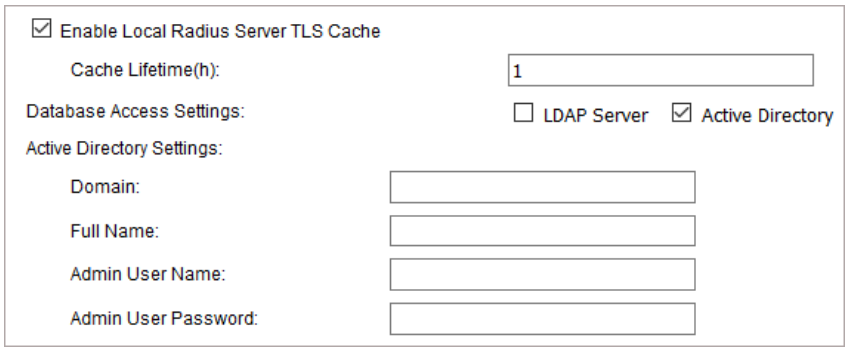
a Enter the lifetime, in hours, in the **Cache Lifetime(h)** field. The minimum and default is **1** hour; the maximum is 99999 hours.

10 Choose the database access method from **Database Access Settings**:

- **LDAP Server** – More options appear; go to Step 11.



- **Active Directory** – More options appear; go to Step 18.



11 Enter the name or IP address of the LDAP server in the **Name or IP address** field.

12 Enter the base distinguished name in the **Base DN** field.

13 Enter the Identity distinguished name in the **Identity DN** field.

14 Enter the distinguished name password in the Identity **DN Password** field.

15 To enable LDAP Transport Layer Security (TLS), select **Enable Ldap TLS**. This option is not selected by default.

16 To enable LDAP cache, select **Enable Ldap Cache**. The Ldap Cache Lifetime(s) field becomes active.

   a Enter the lifetime, in seconds in the **Ldap Cache Lifetime(s)** field; the minimum is 1, the maximum is 99999, and the default is **86400**.

17 Go to Step 22.

18 Enter the domain name in the **Domain** field.

19 Enter the full name of the Active Directory in the **Full Name** field.

20 Enter the user name of the administrator user in the **Admin User Name** field.

21 Enter the password of the administrator user in the **Admin User Password** field.

22 Click **OK**.

Configuring DPI-SSL Granular Control per Zone

DPI-SSL granular control allows you to enable DPI-SSL on a per-zone basis rather than globally. You can enable both DPI-SSL Client and DPI-SSL Server per zone. For further information, see *SonicOS 6.5 Security Configuration*

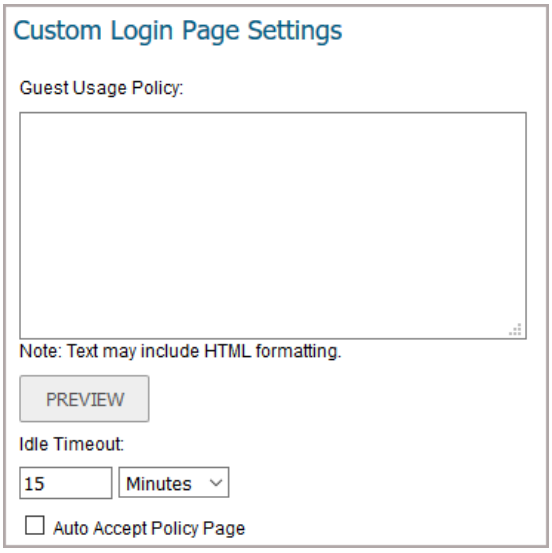
Enabling Automatic Redirection to the User-Policy Page

SonicOS 6.5 allows you to redirect a guest automatically to your guest-user policy page. If you enable this feature, also known as the zero-touch policy page redirection, the guest user is redirected automatically to your guest-user policy page. If you disable the feature, the guest must click **ACCEPT**.

To enable automatic redirection to the user-policy page:

1. Navigate to **MANAGE | System Setup > Network > Zones**.
2. Click either the:
   - **Add** icon to add a new zone.
   - **Edit** icon of an existing zone.

   The **Add Zone/Edit Zone** dialog displays.
3. Click **Guest Services**.
4. Click **Enable Guest Services**.
5. Click **Enable Policy Page without authentication**.
6. Click **CONFIGURE**. The **Custom Policy Message** dialog displays.



7. Select **Auto Accept Policy Page**. This option is not selected by default.
8. Click **OK**.
9. Finish configuring the zone.
10. Click **OK**.

Deleting a Zone

To delete a user-created zone:

1 Navigate to **MANAGE | System Setup | Network > Zones**.

> **NOTE:** The **Delete** icon is unavailable for predefined zones. You cannot delete these zones. Any zones that you create can be deleted.

2 Clicking the **Delete** icon in the zone's **Configure** column.

To delete one or more user-created zones:

1 Navigate to **MANAGE | System Setup | Network > Zones**.

> **NOTE:** The checkboxes are unavailable for predefined zones. You cannot delete these zones. Any zones that you create can be deleted.

2 Select the zones to delete.

3 From **Delete**, select which zones to delete:

- **Delete Selected**
- **Delete All**

18

Configuring Wire Mode VLAN Translation

Topics:

- **Network > VLAN Translation** on page 413
  - About VLAN Translation on page 413
  - Creating and Managing VLAN Maps on page 415

Network > VLAN Translation

NOTE: VLAN Translation is available on all platforms that support Wire Mode.

NOTE: VLAN Translation and Wire Mode over VLAN interfaces cannot be enabled at the same time.

- About VLAN Translation on page 413
- Creating and Managing VLAN Maps on page 415

About VLAN Translation

The VLAN Translation (mapping) feature allows traffic arriving on a VLAN to a Wire Mode interface operating in Secure mode to be mapped to a different VLAN on the outgoing paired interface. Re-routing some of the traffic coming into the SonicWall Security Appliance onto different VLANS allows you to perform further analysis, processing, or merely remapping traffic. This feature is supported on all Wire Mode-capable devices.

An advantage of Wire Mode, that is, you can pre-provision the VLAN mapping. This allows you to have the mapping in place before the interface receives traffic. You also can add and delete mapping on an active Wire Mode interface.

Topics:

- Mapping Modes on page 413
- Mapping Persistence on page 414
- Map Multiple Interface Pairs on page 414

Mapping Modes

You can create a VLAN mapping in these modes:

- Unidirectional mapping – For example, use to:
  - Secure printing from a less-secure network to a high-secure network

- Transfer application and operating system updates from a less-secure network to a high-secure network
- Monitor multiple networks in a SOC (security operations center)
- Provide time synchronisation in high-secure networks
- Transfer files
- Provide a “you have mail” alert to a high-secure network from a less-secure network

- Bidirectional mapping – For example, use to setup a two-way connection to and from devices through the Security Appliance, for example, TCP.

Mapping Persistence

The VLAN map created for a pair of interfaces is persistent over reload and is stored as part of the configuration. If the wire-mode pair (secure mode) have mapping associated with them, the wire mode cannot be changed unless the mapping policy is deleted.

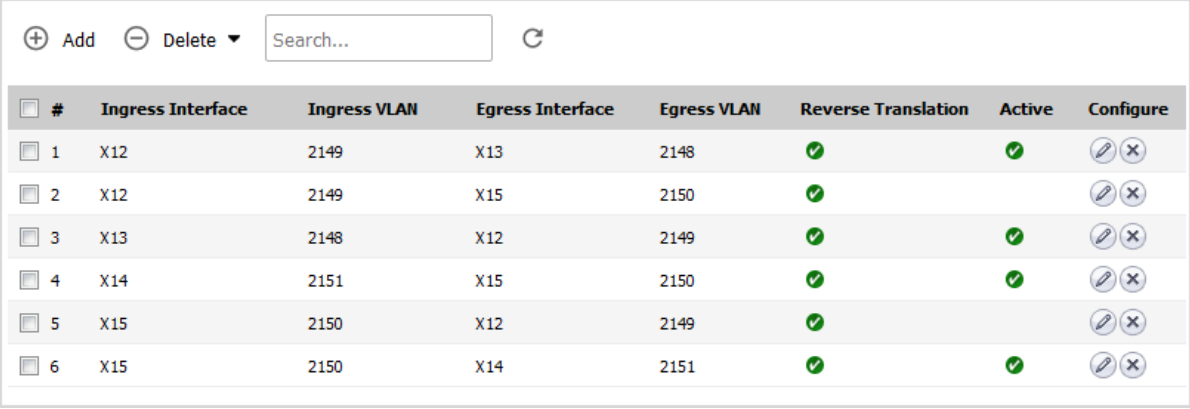
Map Multiple Interface Pairs

You can create VLAN mapping for multiple pairs of interfaces at the same time. These interfaces must form part of an existing Secure Wire Mode pair at the time of the VLAN mapping creation. You can also create mappings for an interface with multiple interfaces, but only the mappings for the current active Wire Mode pair are in use at any given time.

If the paired interface is changed, the message, `Cannot change wire-mode pair interface when WireMode VLAN entries exist for the interface`, displays.

Example

Multiple interface pairs mapping



#	Ingress Interface	Ingress VLAN	Egress Interface	Egress VLAN	Reverse Translation	Active	Configure
1	X12	2149	X13	2148	✓	✓	
2	X12	2149	X15	2150	✓		
3	X13	2148	X12	2149	✓	✓	
4	X14	2151	X15	2150	✓	✓	
5	X15	2150	X12	2149	✓		
6	X15	2150	X14	2151	✓	✓	

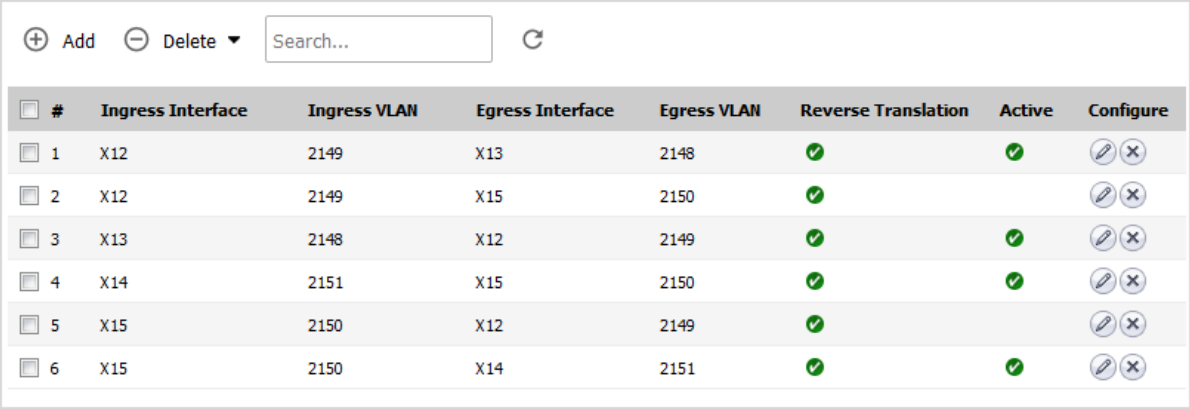
In Multiple interface pairs mapping, a mapping exists for X12 to X13 (policy 1) as well as X12 to X15 (policy 2).

As only X12 and X13 (policies 1 and 3) and X14 and X15 (policies 4 and 6) are currently forming a Wire Mode pair, only policies 1, 3, 4, and 6 are active as indicated by the green checkmark in the active column.

NOTE: The wire-mode pair interfaces cannot change if Wire Mode VLAN entries exist for the interface.

Creating and Managing VLAN Maps

Network > VLAN Translation allows you to create and manage the VLAN mapping of interfaces.



#	Ingress Interface	Ingress VLAN	Egress Interface	Egress VLAN	Reverse Translation	Active	Configure
1	X12	2149	X13	2148	✓	✓	
2	X12	2149	X15	2150	✓		
3	X13	2148	X12	2149	✓	✓	
4	X14	2151	X15	2150	✓	✓	
5	X15	2150	X12	2149	✓		
6	X15	2150	X14	2151	✓	✓	

Add icon	Displays the **Add VLAN Translation** dialog.
Delete icon	Displays the **Delete** drop-down menu: • **Delete Selected** • **Delete All**
Search field	Allows you to display only those VLAN translations of interest.
Refresh icon	Refreshes the VLAN Translation table.
Policy number and checkbox	Number of the policy and its associated checkbox.
Ingress Interface	Name of the incoming interface.
Ingress VLAN	VLAN tag of the incoming interface.
Egress Interface	Name of the interface to which traffic is mapped.
Egress VLAN	VLAN tag of the interface to which traffic is mapped.
Reverse Translation	Indicates whether the mapping is unidirectional or bidirectional: • **Disabled** – Unidirectional; column blank. • **Enabled** – Bidirectional; green checkmark.
Active	Status of the mapped pair: • **Active** – The Wire Mode pair is mapped and active; green checkmark. • **Inactive** – The Wire Mode pair is mapped but not active (pre-provisioned); column blank.
Configure	Displays **Edit** and **Delete** icons for a mapped pair.

Topics:

- Creating a VLAN Map on page 415
- Managing VLAN Mappings on page 419

Creating a VLAN Map

You can create a unidirectional VLAN map before or after a Wire Mode pair. Creating a VLAN map is a two-step process:

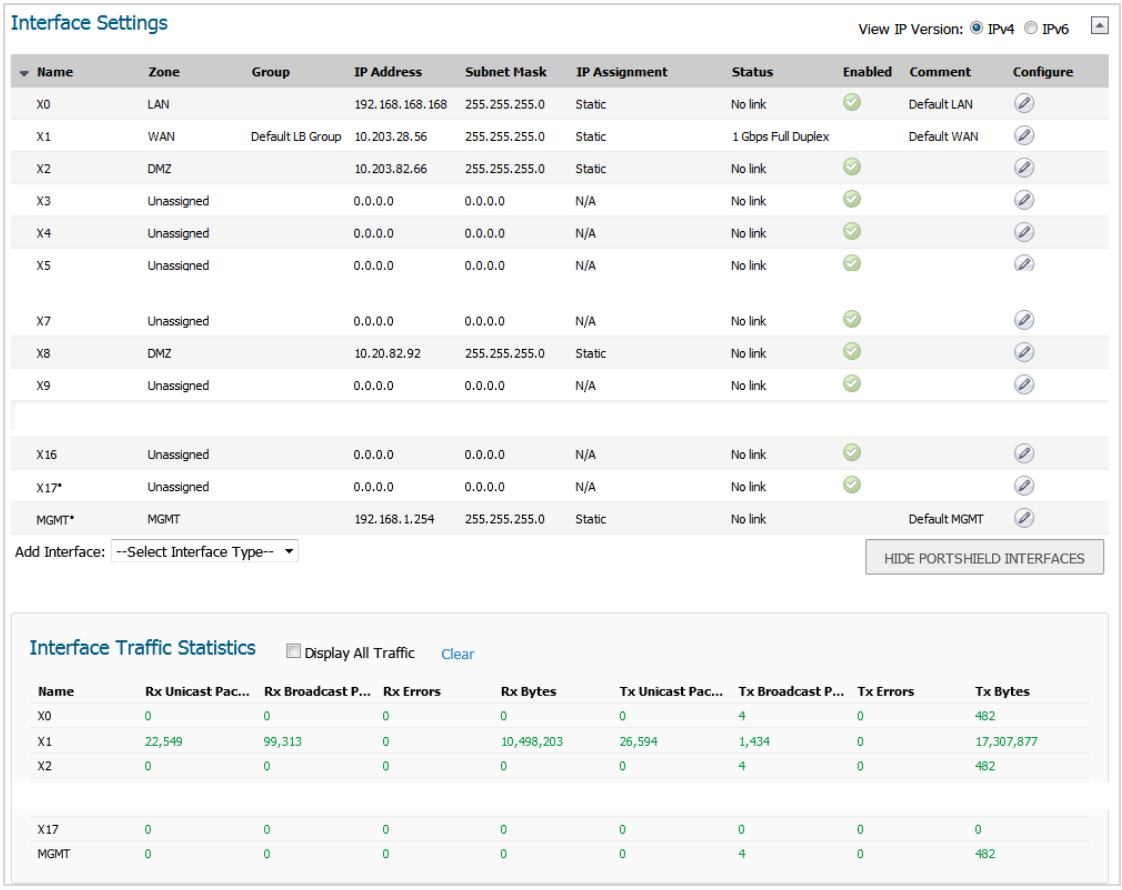
1 Creating a Wire Mode Pair in Secure Mode on page 416

2 Creating the VLAN Mapping on page 418

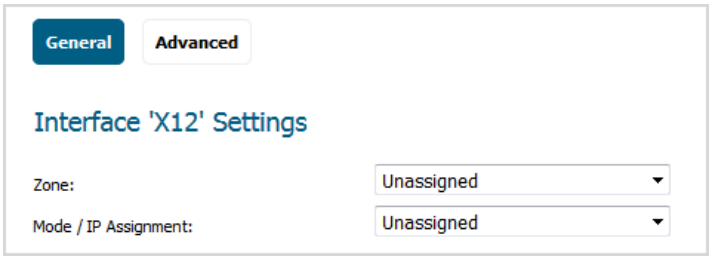
Creating a Wire Mode Pair in Secure Mode

To create a Wire Mode pair in secure mode:

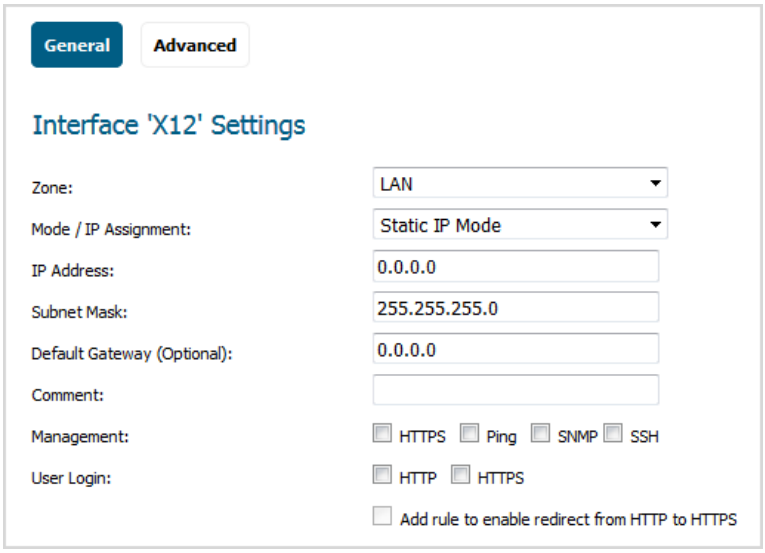
1 Navigate to **MANAGE | System Setup | Network > Interfaces**.



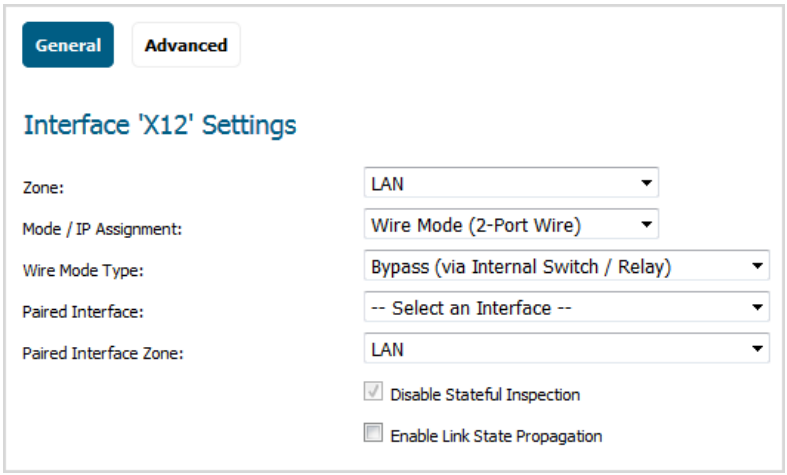
2 Click the **Edit** icon for the interface to be part of the Wire Mode pair. The **Edit Interface** dialog displays.



3 Select the zone for the Wire Mode pair from **Zone**. The options change.



4 Select **Wire Mode (2-Port Wire)** from **Mode / IP Assignment**. The options change again.



5 Select **Secure (Active DPI of Inline Traffic)** from **Wire Mode Type**.

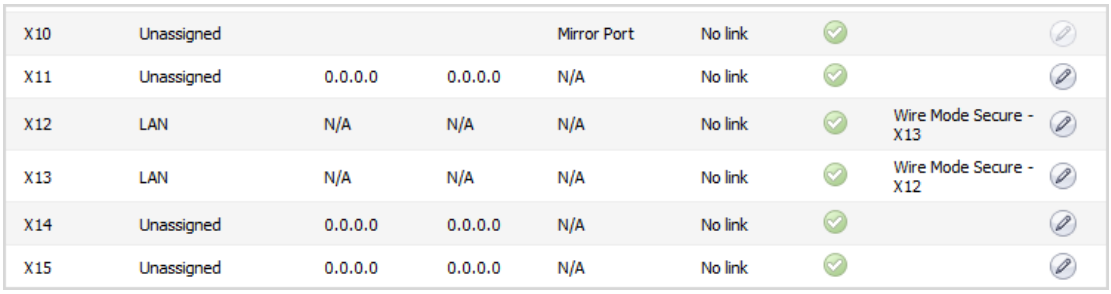
6 Select the interface to pair with the current interface from the **Paired Interface** drop-down menu.

TIP: Ensure the interface you pair with is unassigned.

7 Select the zone for the paired interface from **Paired Interface Zone**. The default is **LAN**.

8 Configure the other options as if configuring a regular Wire Mode pair as described in Configuring Wire and Tap Mode on page 310 and Configuring Wire and Tap Mode on page 310.

9 Click **OK**. The **Network > Interfaces** page is updated.



X10	Unassigned			Mirror Port	No link		
X11	Unassigned	0.0.0.0	0.0.0.0	N/A	No link		
X12	LAN	N/A	N/A	N/A	No link	Wire Mode Secure - X13	
X13	LAN	N/A	N/A	N/A	No link	Wire Mode Secure - X12	
X14	Unassigned	0.0.0.0	0.0.0.0	N/A	No link		
X15	Unassigned	0.0.0.0	0.0.0.0	N/A	No link		

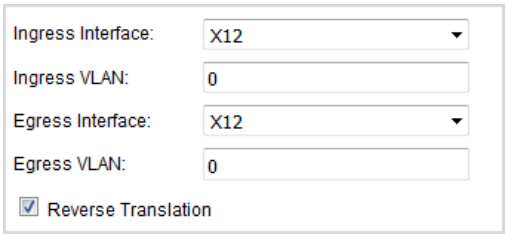
Creating the VLAN Mapping

To create a VLAN mapping:

1 Navigate to **Network > VLAN Translation**.



2 Click the **Add** icon. The **Add VLAN Translation** dialog displays.



3 Select the Wire Mode interface in the pair on which you expect to receive traffic from **Ingress Interface**.

4 Set **Ingress VLAN** to the VLAN on which you expect to receive traffic for mapping.

5 Select the Wire Mode interface in the pair on which you want to map traffic to the **Egress Interface** drop-down menu.

6 Set **Egress VLAN** to the VLAN to which you expect to map traffic.

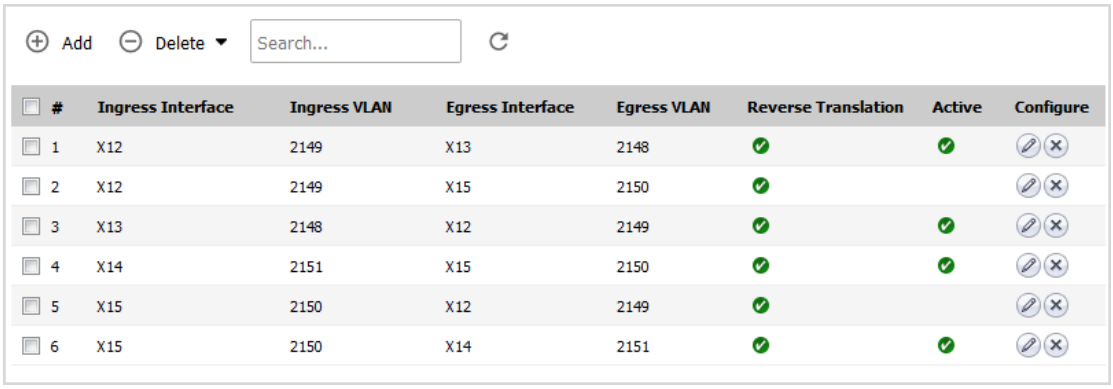
7 To create a:

- Unidirectional mapping, ensure the **Reverse Translation** checkbox is not selected. For example, to map VLAN X on interface A to VLAN Y on interface B.

  NOTE: This option is selected by default.

- Bidirectional mapping, select the **Reverse Translation** checkbox. For example, to map VLAN Y on interface B to VLAN X on interface A as well as map VLAN X on interface A to VLAN Y on interface B.

8 Click **Add**. The **Wiremode VLAN Translation** table is updated.



#	Ingress Interface	Ingress VLAN	Egress Interface	Egress VLAN	Reverse Translation	Active	Configure
1	X12	2149	X13	2148	✓	✓	
2	X12	2149	X15	2150	✓		
3	X13	2148	X12	2149	✓	✓	
4	X14	2151	X15	2150	✓	✓	
5	X15	2150	X12	2149	✓		
6	X15	2150	X14	2151	✓	✓	

Managing VLAN Mappings

Topics:

- Editing Mappings on page 419
- Filtering Mappings on page 419
- Deleting Mappings on page 419

Editing Mappings

To edit a mapping, click its **Edit** icon in the **Configuration** column. The **Edit VLAN Translation** dialog displays. You can change any of the mappings except the **Reverse Translation** setting.

Filtering Mappings

If you have a lot of VLAN mappings, you can display only those of interest by:

1 Entering an interface name or VLAN tag in the **Search** field.

2 Pressing **Enter**.

Only those mappings meeting the search criterion are displayed.

To redisplay all the mappings:

1 Delete the criterion from the **Search** field.

2 Press **Enter**.

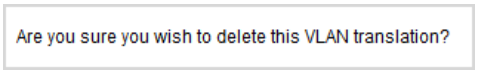
Deleting Mappings

To delete mappings:

1 To delete:

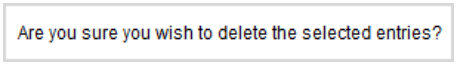
- A single mapping by:
  - Clicking i**ts Delete icon in the Configuration column.**

    A confirmation message displays:



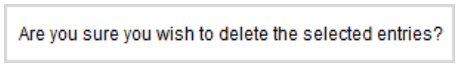
- **Clicking its Selection** checkbox and then selecting **Delete** Selected from the **Delete** drop-down menu.

  A confirmation message displays:



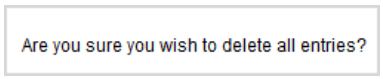
- Multiple mappings by clicking their **Selection** checkboxes and then selecting **Delete** Selected from the **Delete** drop-down menu.

  A confirmation message displays:



- All mappings by selecting **Delete** Selected from the **Delete All** drop-down menu.

  A confirmation message displays:



2 Click **OK**.

If a policy is bidirectional, then both directions are deleted if one is deleted.

19

Configuring DNS Settings

Topics:

- **Network > DNS** on page 421
  - About Split DNS on page 421
  - Managing DNS Servers on page 423
  - DNS and IPv4 on page 432

Network > DNS

The Domain Name System (DNS) is a distributed, hierarchical system that provides a method for identifying hosts on the Internet using alphanumeric names called fully qualified domain names (FQDNs) instead of using difficult to remember numeric IP addresses. **MANAGE | System Setup | Network > DNS** allows you to manually configure your DNS settings, if necessary.

NOTE: For complete information on the SonicOS implementation of IPv6, see IPv6 on page 932.

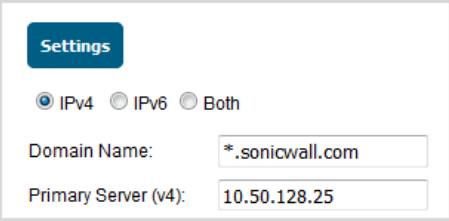
Topics:

- About Split DNS on page 421
- Managing DNS Servers on page 423
- DNS and IPv4 on page 432

About Split DNS

Split DNS is an enhancement that allows you to configure a set of servers and associate them to a given domain name (which can be a wildcard). When SonicOS DNS Proxy receives a query that matches the domain name, the name is transmitted to the designated DNS server. Split DNS example shows how this works:

Split DNS example



- This topology has two firewalls with network connectivity:
  - One firewall is connected to the Internet.
  - Another is a VPN tunnel connected to the corporation network.

- Default DNS queries go to the public ISP DNS Server.
- All queries to `*.sonicwall.com` go to the DNS server located behind the VPN tunnel.

For viewing and configuring split DNS entries, see Configuring Domain-Specific DNS Servers for Split DNS on page 426.

By adding a split DNS entry, all queries to `sonicwall.com` are sent to the specific server (see Configuring Domain-Specific DNS Servers for Split DNS on page 426).

Multiple DNS servers could be configured to handle queries to `sonicwall.com` as well.

About Per-Partition DNS Servers and Split DNS

With or without authentication partitions, it is usually necessary to use a domain's own DNS servers to resolve the names of devices in the domain, and occasionally there can also be a need to use different external DNS servers to resolve external host names. Now, with multiple authentication partitions, this situation is exacerbated as those partitions usually require using different DNS servers to resolve the host names in the different partitions.

NOTE: Use of a domain's own DNS servers can be required unexpectedly because LDAP referrals usually give the referred server by DNS name, even when the LDAP servers are configured by IP address.

An example where different external DNS servers to resolve external host names was required involved external-using cloud services that could not be resolved by the internal domain's DNS servers.

The Split DNS feature is used directly by the SonicWall Security Appliance to resolve the names of devices in domains without the need to enable DNS Proxy, including for multiple unrelated domains with authentication partitioning.

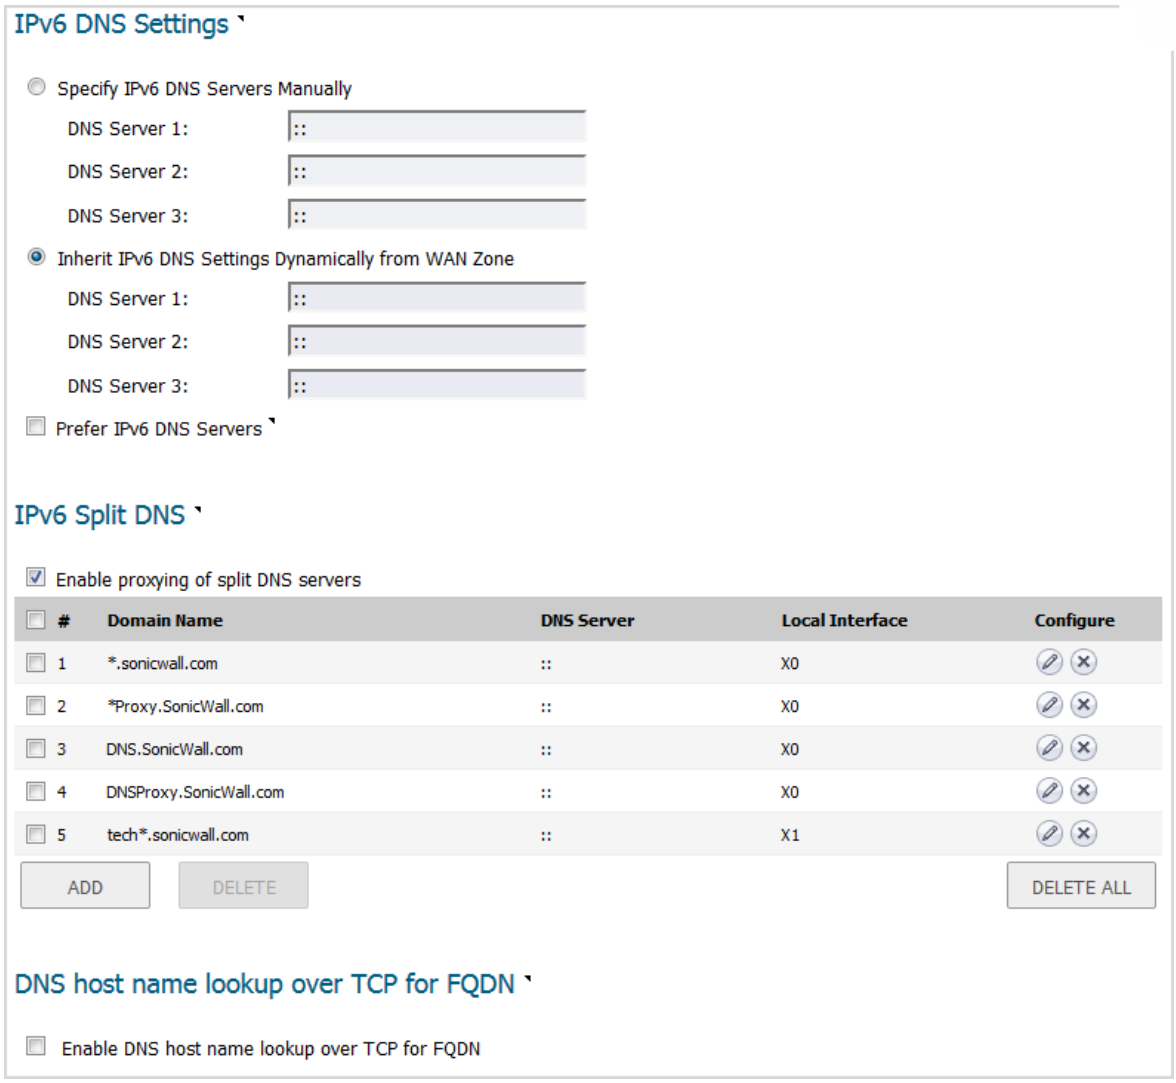
DNS servers configured in Split DNS are used directly for DNS lookups of host names in internal domains as follows:

- This applies for anything that has entries in the Security Appliance's main DNS Cache:
  - SMTP servers
  - SYSLOG servers
  - Web Proxy servers and User (internal) Proxy servers
  - GMS and GMS standby
  - POP servers
  - RADIUS authentication and accounting servers
  - LDAP servers
  - SSO / Terminal Services agents and RADIUS accounting clients
- If partitioning is enabled and a partition has one domain or one tree of parent/sub-domains (AKA one AD Forest), then if Split DNS servers are configured for the partition's top-level domain, then those are copied into the internal partition structure. Those DNS servers are then used to resolve the names of agents, servers, and clients in the partition.
- If partitioning is enabled and a partition is configured with multiple separate domains (which is allowed but is not common), then no DNS servers are copied into the partition structure, relying instead on the mechanism described below.
- If partitioning is disabled or a partition has no DNS servers set, or for resolving items not associated with a partition, the DNS servers to use are selected per-request through the API provided by Split DNS.

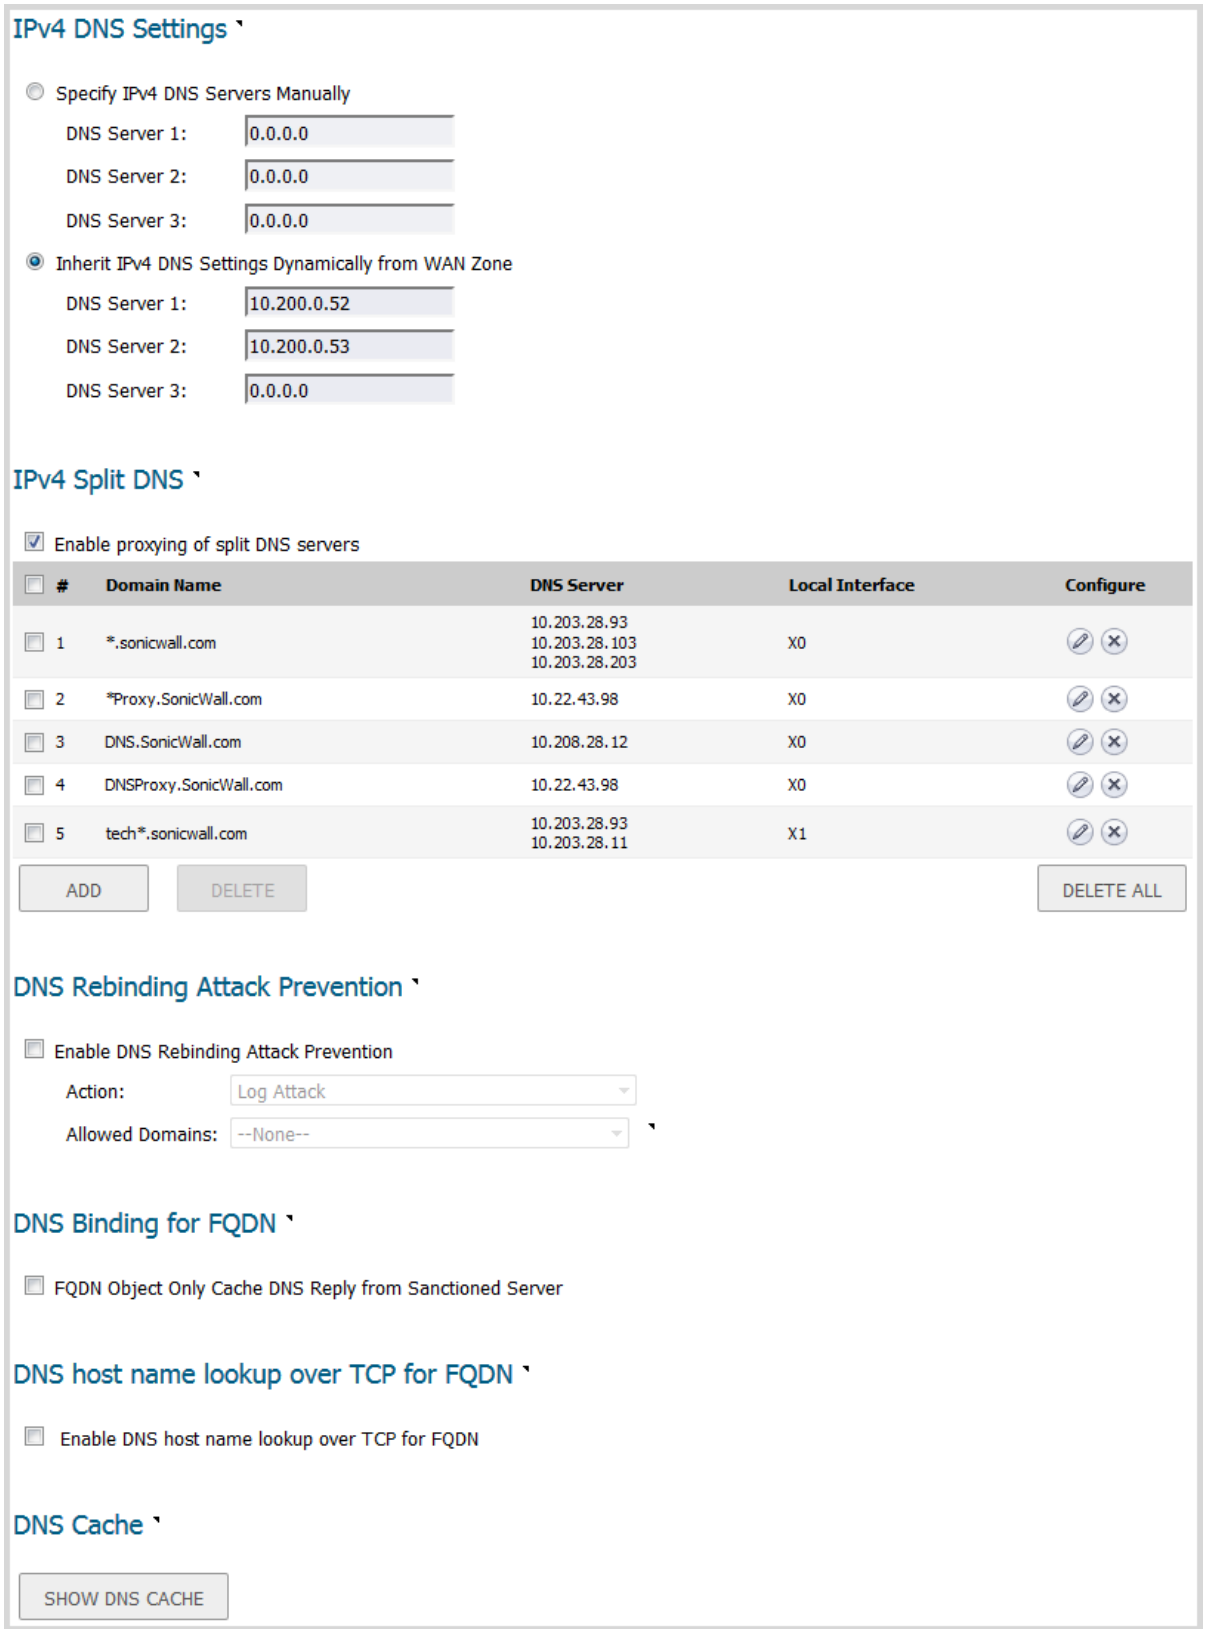
Managing DNS Servers

The options on **Network > DNS** change depending on whether you specify IPv6 or IPv4; see: IPv6 **Network > DNS** and IPv4 **Network > DNS**.

IPv6 Network > DNS



IPv4 Network > DNS



The two versions of the management interface page have the **DNS Settings**, **Split DNS**, and **DNS host name lookup over TCP for FQDN** sections in common and are described together.

Topics:

- Selecting IP Version on page 425
- Specifying which DNS Servers are Used on page 425
- Configuring Domain-Specific DNS Servers for Split DNS on page 426
- Enabling DNS Host Name Lookup over TCP for FQDN on page 430

Selecting IP Version

To select the IP version:

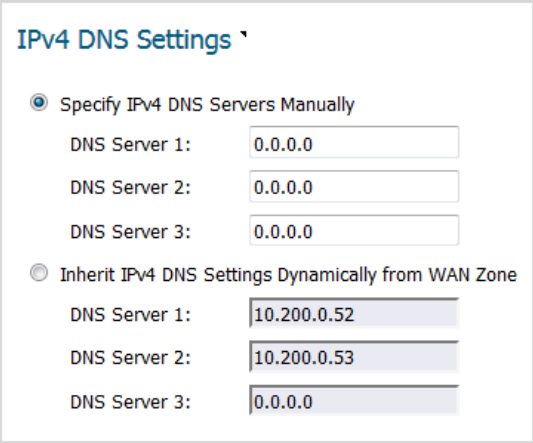
1. Navigate to **MANAGE | System Setup > Network > DNS**.
2. From **View IP Version** at the top right of the page, choose:
   - **IPv4**
   - **IPv6**

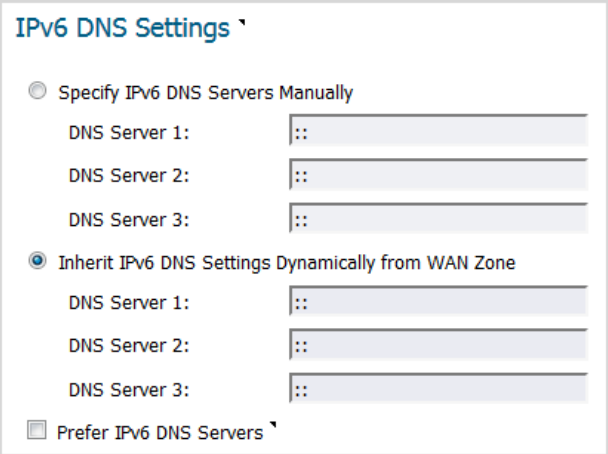
The options on **Network > DNS** change depending on whether you specify IPv6 or IPv4; see IPv6 **Network > DNS** and IPv4 **Network > DNS**.

Specifying which DNS Servers are Used

Regardless of the IP version, you can specify how SonicOS selects the DNS servers. The method is the same for both IP versions.

IPv4 DNS Settings/IPv6 DNS Settings section





To specify which DNS servers are used:

1. Navigate to **Network > DNS**.
2. In the **IPv4/IPv6 DNS Settings** section, select one of the following:
   - To manually specify the DNS servers.
     a) Select **Specify IPv4/IPv6 DNS Servers Manually**.
     b) Enter up to three IP addresses into the **DNS Server** fields.
     c) If you are using:

- IPv4 go to Step 4.
- IPv6, go to Step 3.

- To use the DNS Settings configured for the WAN zone:

  a) Select **Inherit IPv4/IPv6 DNS Settings Dynamically from WAN Zone**. This is the default. The IP address(es) are populated into the **DNS Server** fields automatically.

  b) For IPv4, go to Step 4.

3 To use only IPv6 servers, select **Prefer IPv6 DNS Servers**. This option is not selected by default.

SonicOS DNS supports these server types:

- DNS_SYSTEM_BEHAVIOR – the system default behavior, which depends on the setting of this option.
- DNS_PREFER_V4_DNSSERVER – IPv4 DNS servers preferred unless there is a failure, then IPv6 DNS servers are requested.
- DNS_PREFER_V6_DNSSERVER: – IPv6 DNS servers preferred unless there is a failure, then IPv4 DNS servers are requested.

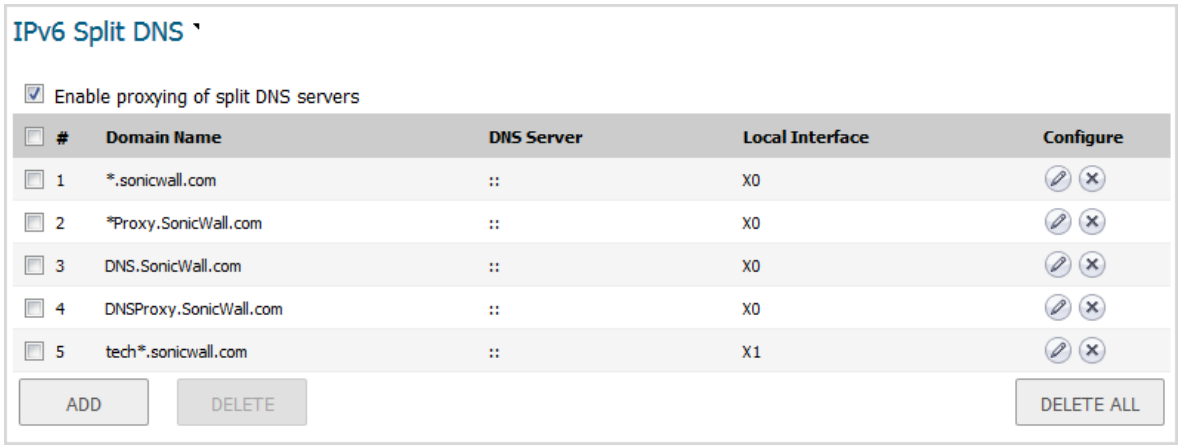
CAUTION: Select this option only if you have configured the IPv6 DNS server correctly.

4 Click **Accept** to save your changes.

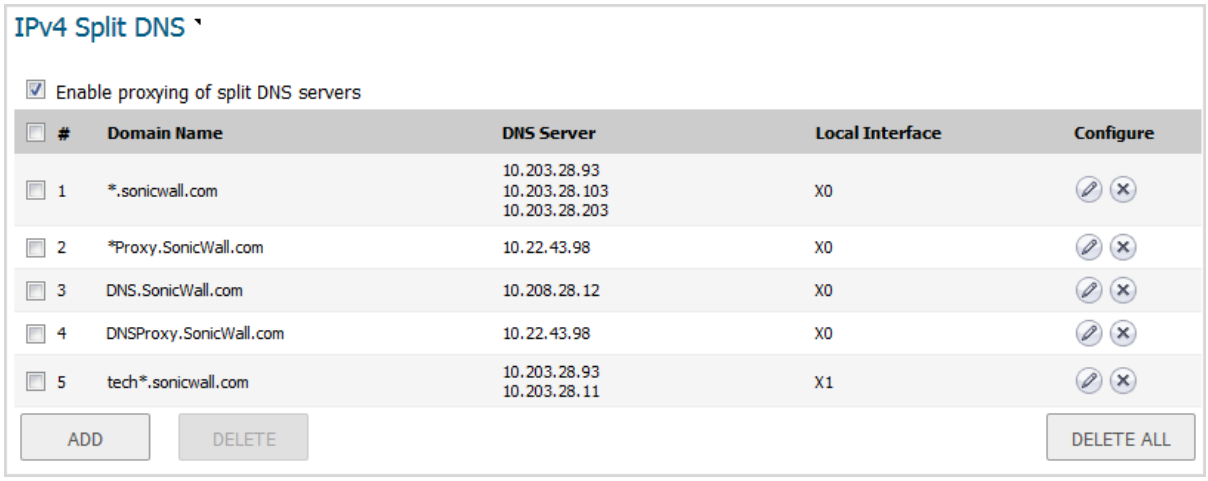
Configuring Domain-Specific DNS Servers for Split DNS

You can optionally configure separate domain-specific DNS servers to use with either IPv6 or IPv4. The method is the same for both IP versions. Any differences are noted.

IPv6 Split DNS section



IPv4 Spit DNS section



Domain name	Name of the DNS Server.
DNS Server	IPv4/IPv6 IP address of the DNS Server. **NOTE:** The status of the DNS servers are displayed on the **Network > DNS Proxy** page.
Local Interface	Interface assigned to the DNS Server.
Configure	Contains **Edit** and **Delete** icons for each server.

Topics:

- Adding a DNS Server on page 427
- Editing Split DNS Entries on page 430
- Deleting Split DNS Entries on page 430

Adding a DNS Server

To add domain-specific DNS servers and associate them to a given domain name:

IMPORTANT: The maximum number of entries for Split DNS is 32. If the list is full, new entries cannot be added.

1 Navigate to **MANAGE | System Setup > Network > DNS**.

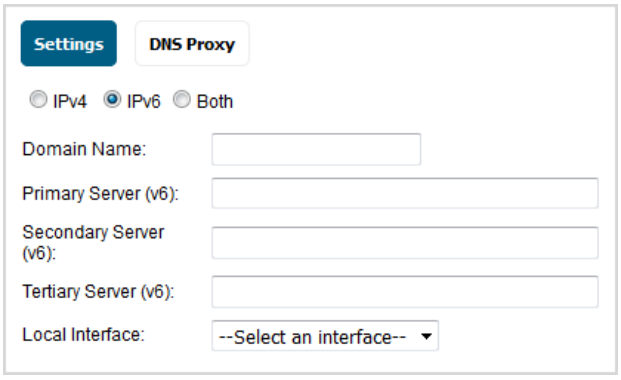
2 Choose the IP version from **View IP Version**.

3 To enable proxying of split DNS servers, select **Enable proxying of split DNS servers**. This option is selected by default.

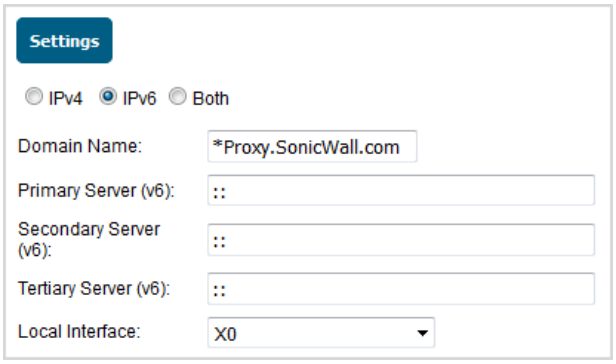
4 Under the **Split DNS** table, click **Add**. The **Add Split DNS Entry** dialog displays.

TIP: If you selected DNS Proxy, a page for it, **DNS Proxy**, also displays on the **Add Split DNS Entry** dialog.

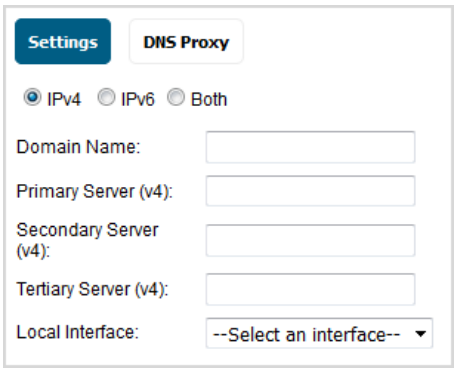
IPv6 Add Split DNS Entry—DNS Proxy enabled

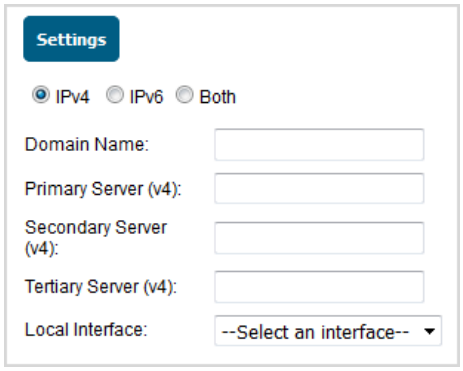


IPv6 Add Split DNS Entry—DNS Proxy disabled

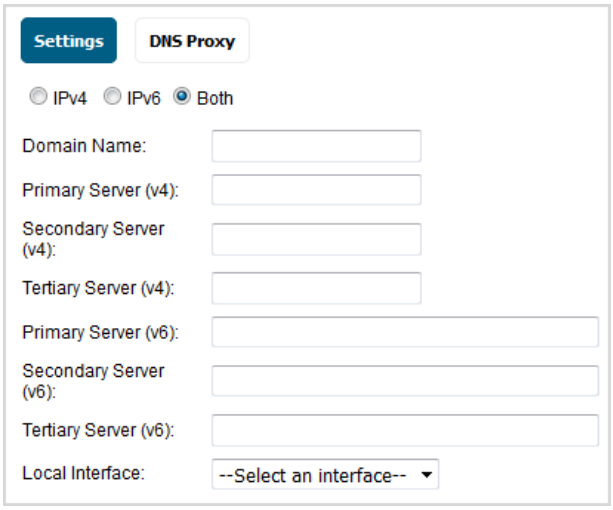


IPv4 Add Split DNS Entry





IPv6 and IPv4 Add Split DNS Entry—DNS Proxy enabled



5 Choose the IP version:

- **IPv4**
- **IPv6**
- **Both**

6 Enter the domain name in the **Domain Name** field. The name can contain a wildcard (`*`; for example, `*.SonicWall.com`).

7 To configure one or more IPv4/IPv6 Split DNS Servers for this domain, enter the IP addresses in the appropriate fields:

- **Primary Server (v4/v6)**
- **Secondary Server (v4/v6)** (optional)
- **Tertiary Server (v4/v6)** (optional)

8 Select an interface from **Local interface**.

9 if you have not enabled DNS Proxy, go to Step 13.

10 Click **DNS Proxy**.



11 To specify a Time to Live, select **Manually set TTL value in DNS reply**. This option is not selected by default. If this option is not selected, the TTL value is the same as that from the DNS response; if it is set, the TTL value is the same as the setting.

NOTE: This option applies only when Split DNS is used by DNS Proxy.

12 Enter the maximum time for the cache entry to exist. The minimum is 1 second, the maximum is 9999999999999999 seconds.

13 Click **OK**.

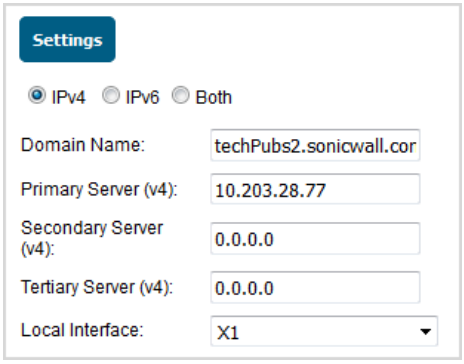
TIP: The DNS servers display in the **Split DNS** table of both IP versions regardless of which IP version was chosen when configuring them.

Editing Split DNS Entries

To edit a Split DNS entry.

1 Navigate to **MANAGE | System Setup > Network > DNS**.

2 In the **Split DNS** table, click the entry's **Edit** icon. The **Edit Split DNS Entry** dialog displays.



3 Make the changes.

4 Click **OK**.

Deleting Split DNS Entries

To delete a Split DNS entry:

1 Click the entry's **Delete** icon.

To delete two or more Split DNS entries:

1 Select the checkboxes of the entries to be deleted. **Delete** becomes available.

2 Click **Delete**.

To delete all Split DNS entries:

1 Click **Delete All**.

Enabling DNS Host Name Lookup over TCP for FQDN

By default, DNS queries are sent over UDP. The DNS response can include a Truncated flag if the response length exceeds the maximum allowed by UDP.

When the **Enable DNS host name lookup over TCP for FQDN** option is:

- Enabled and the Truncated flag is set in the DNS response, SonicOS sends an additional DNS query over TCP to determine the full DNS response for multiple IP addresses.
- Disabled, DNS queries are sent over UDP, and SonicOS only processes the IP addresses in the DNS response packet, although the Truncated flag is set in the response.

The DNS query times out after one second if no DNS response over TCP is received from the DNS server.

This option is used to gain more IP addresses when sending DNS queries from FQDN over TCP while the Security Appliance receives DNS responses over UDP.

To enable DNS host name lookup over TCP for FQDN:

1 Navigate to **MANAGE | System Setup** > **Network > DNS**.

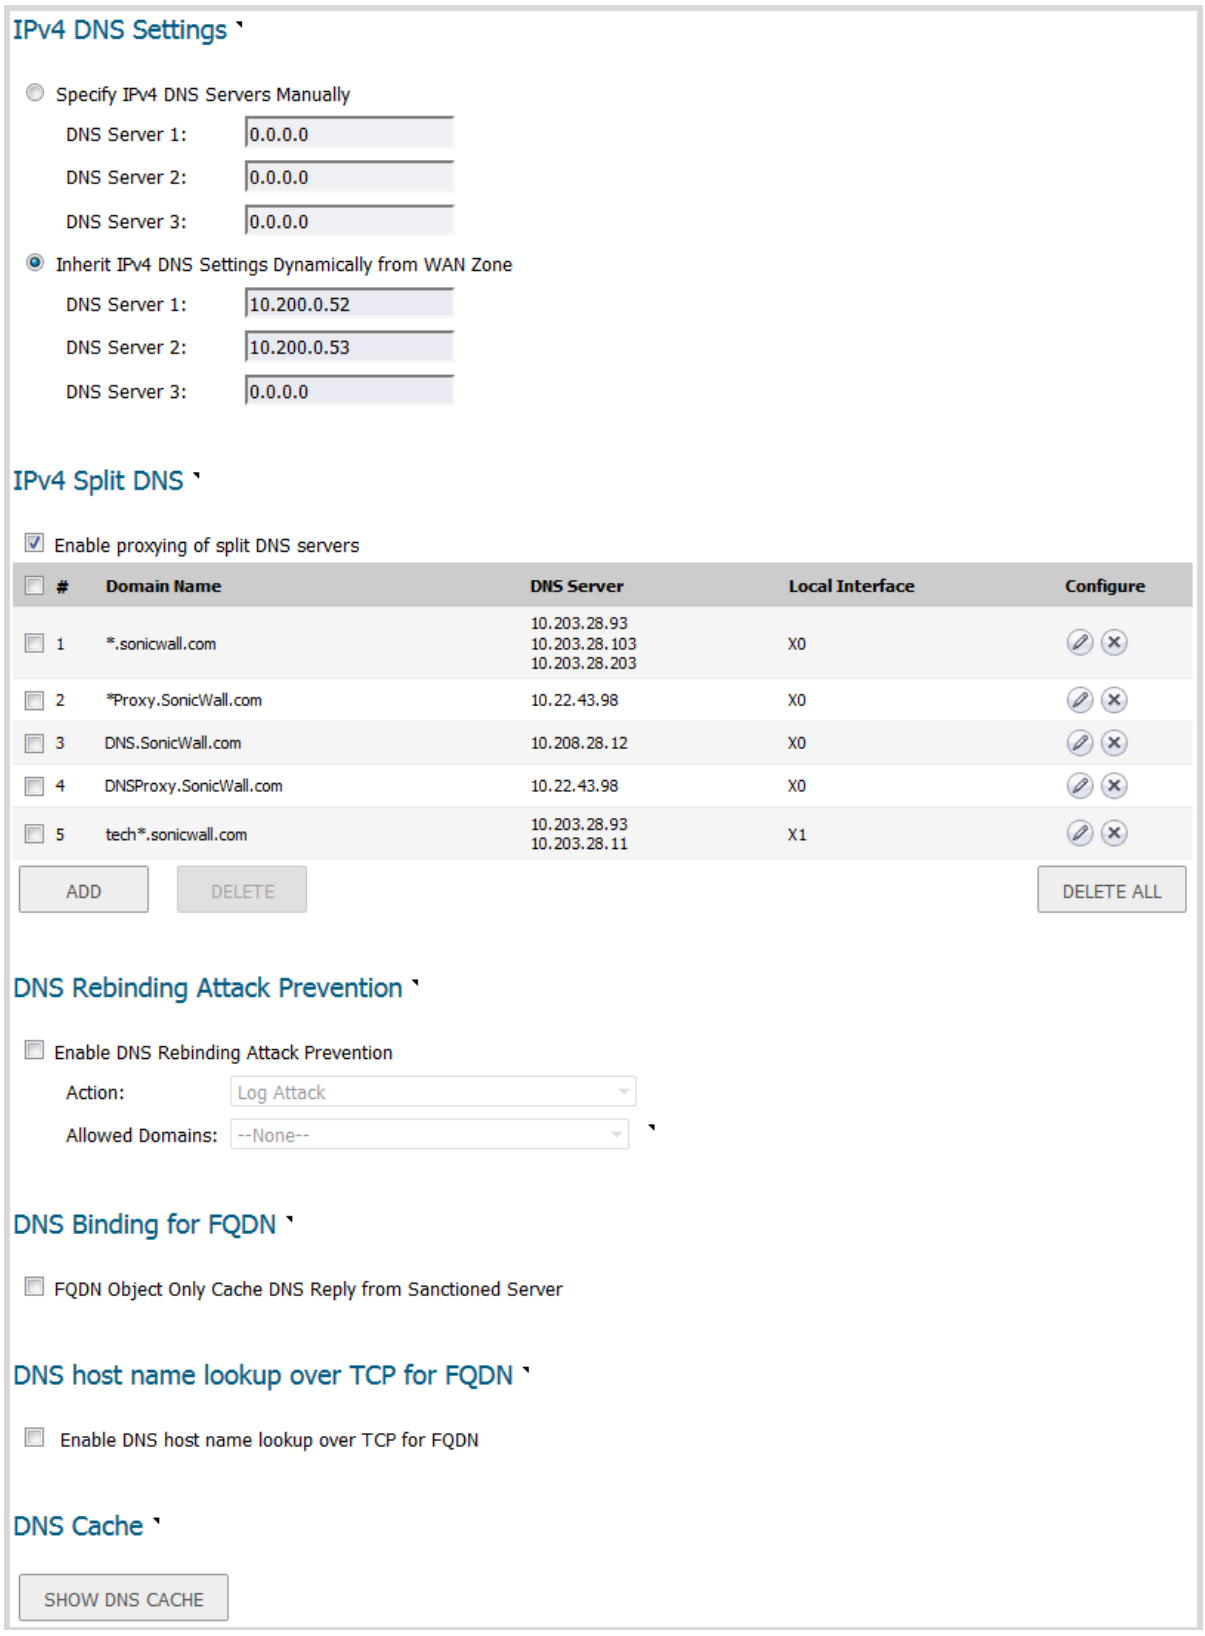
2 Scroll to the **DNS host name lookup over TCP for FQDN** section.



3 Select **Enable DNS host name lookup over TCP for FQDN**. This option is not selected by default.

4 Click **Accept**.

DNS and IPv4



The IPv4 **Network > DNS** page has these sections in addition to those in common with the IPv6 **Network > DNS** page:

- DNS Rebinding Attack Prevention on page 433
- DNS Binding for FQDN on page 433
- DNS Cache on page 434

DNS Rebinding Attack Prevention

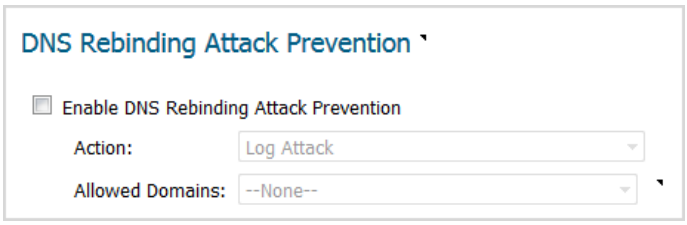
DNS rebinding is a DNS-based attack on code embedded in web pages. Normally requests from code embedded in web pages (JavaScript, Java, and Flash) are bound to the web-site they are originating from (see Same Origin Policy). A DNS rebinding attack can be used to improve the ability of JavaScript-based malware to penetrate private networks, and subvert the browser's same-origin policy.

DNS rebinding attackers register a domain that is delegated to a DNS server they control. The server is configured to respond with a very short Time to Live (TTL) parameter, which prevents the result from being cached. The first response contains the IP address of the server hosting the malicious code. Any subsequent requests contain IP addresses from private (RFC 1918) network, presumably behind a firewall, being target of the attacker. Because both are fully valid DNS responses, they authorize the sandbox script to access hosts in a private network. By iterating addresses in these short-term but still valid DNS replies, the script is able to scan the network and perform other malicious activities.

To configure DNS rebinding attack prevention:

1 Navigate to **Network > DNS**.

2 Scroll to the **DNS Rebinding Attack Prevention** section.



3 Select **Enable DNS Rebinding Attack Prevention**. This option is not selected by default. The two options become available.

4 From **Action**, select an action to perform when a DNS rebinding attack is detected:

- **Log Attack** (default)
- **Log Attack & Return a Query Refused Reply**
- **Log Attack & Drop DNS Reply**

5 From **Allowed Domains**, select an allowed domain FQDN Address Object or FQDN Address Object Group containing allowed domain-names (such as, `*.SonicWall.com`) for which locally connected/routed subnets should be considered legal responses.

You can also create new FQDN address objects or FQDN address object groups by selecting **Create new FQDN Address Object...** or **FQDN Address Object Group...**.

6 Click **Accept**.

DNS Binding for FQDN

To enable DNS binding for FQDN:

1 Navigate to **Network > DNS**.

2 Scroll to the **DNS Binding for FQDN** section.

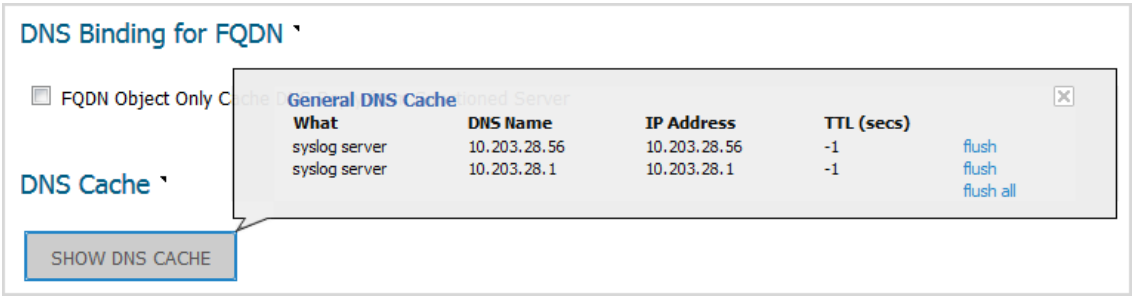


3 Select **FQDN Object Only Cache DNS Reply from Sanctioned Server**. This option is not selected by default.

4 Click **Accept**.

DNS Cache

To show the contents of the general DNS cache, click **Show DNS Cache**. A popup displays the cache contents:



What	DNS Server name: • Forward DNS cache, the host name. • Reverse DNS cache, a string representation of the IP address.
DNS Name	Domain name, such as `www.SonicWall.com`, or IP address.
IP Address	Resolved IP address.
TTL (secs)	Time to Live; the TTL value from the DNS response.
flush	Clicking this flushes the server's DNS cache entry
flush all	Clicking this flushes all DNS cache entry of all listed servers

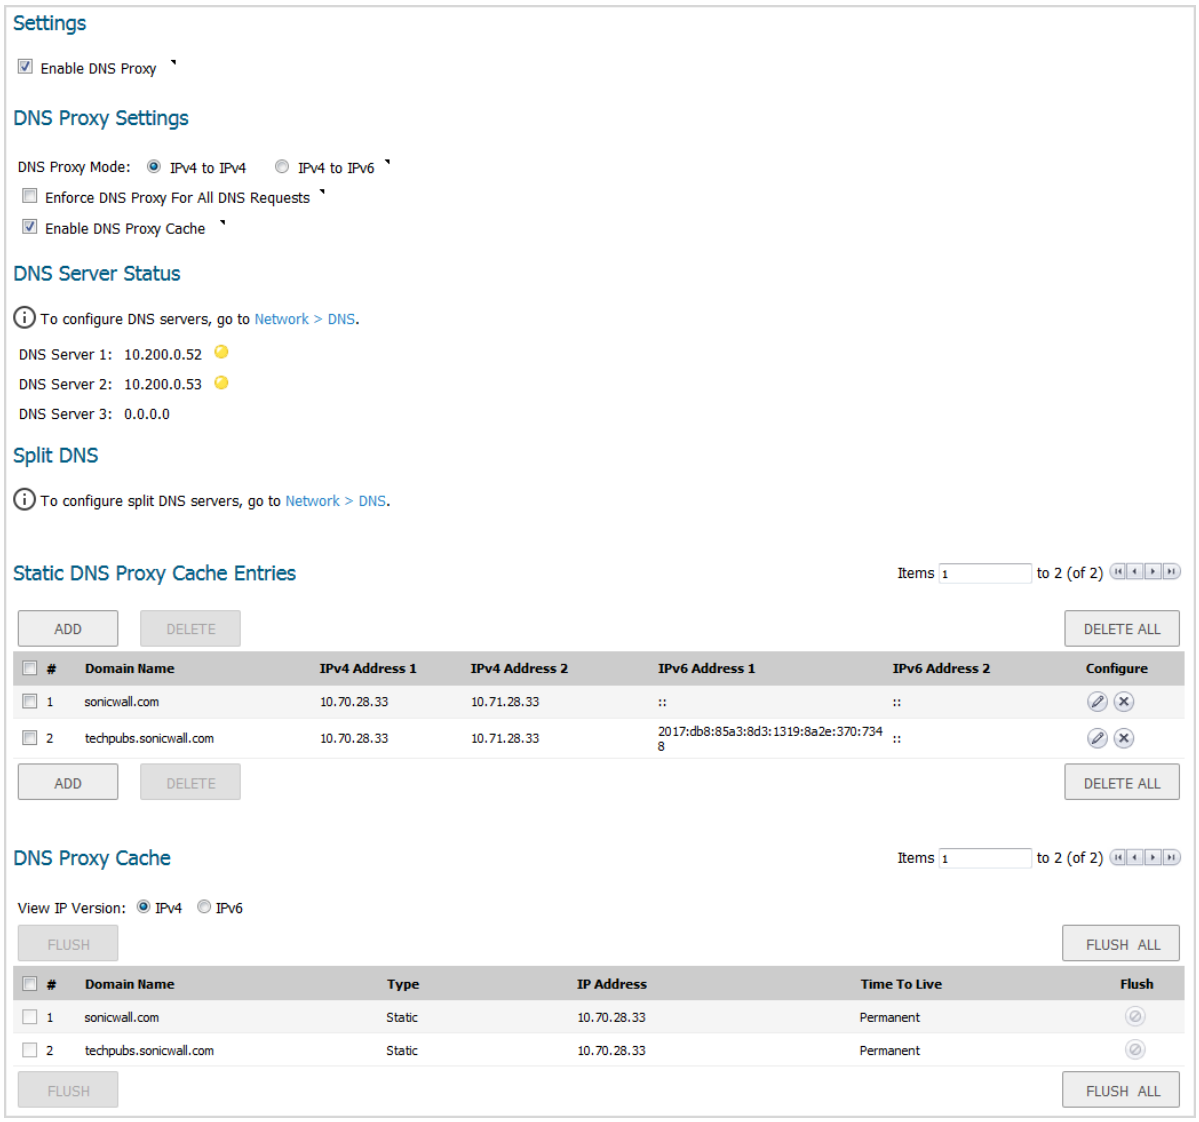
20

Configuring DNS Proxy Settings

Topics:

- **Network > DNS Proxy** on page 436
  - About DNS Proxy on page 437
  - Enabling DNS Proxy on page 439
  - Configuring DNS Proxy Settings on page 440
  - Monitoring DNS Server Status on page 441
  - Monitoring Split DNS Server Status on page 442
  - Viewing and Managing Static DNS Cache Entries on page 442
  - Viewing DNS Proxy Cache Entries on page 444

Network > DNS Proxy



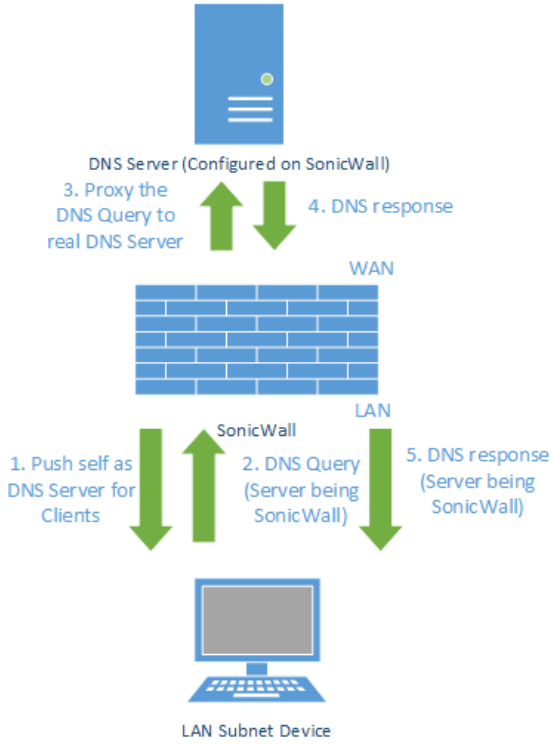
Topics:

- About DNS Proxy on page 437
- Enabling DNS Proxy on page 439
- Configuring DNS Proxy Settings on page 440
- Monitoring DNS Server Status on page 441
- Monitoring Split DNS Server Status on page 442
- Viewing DNS Proxy Cache Entries on page 444
- Viewing and Managing Static DNS Cache Entries on page 442

About DNS Proxy

An IPv4 interface can do name resolution on an IPv4 Internet, and an IPv6 interface can only do name resolution on an IPv6 Internet through DNS proxy. To allow IPv4 clients to access DNS services in a network with mixed IPv4 and IPv6 interfaces, SonicOS supports DNS proxy; see DNS Proxy.

DNS Proxy



The DNS proxy feature provides a transparent mechanism that allows devices to proxy hostname resolution requests on behalf of clients. The proxy can use existing DNS cache, which is either statically configured by you or learned dynamically, to respond to the queries directly.

The proxy can redirect the DNS queries selectively to specific DNS servers, according to partial or complete domain specifications. This is useful when VPN tunnels or PPPoE virtual links provide multiple network connectivity, and it is necessary to direct some DNS queries to one network, and other queries to another network

With DNS Proxy, LAN Subnet devices use the SonicWall Security Appliance as the DNS Server and send DNS queries to the Security Appliance. The Security Appliance proxies the DNS queries to the real DNS Server. In this way, the Security Appliance is the central management point for the network DNS traffic, providing the ability to manage the DNS queries of the network at a single point.

NOTE: To maintain security, an incoming DNS Query is proxied only after Access Rule and DPI checking.

When DNS proxy is enabled on an interface, one Allow Rule is auto-added by SonicOS. For the Access Rules associated with the interface, see *SonicOS Policies*.

When **DNS Proxy over TCP** is enabled, another Allow Rule is auto-added.

Topics:

- Supported Interfaces on page 438
- DNS Server Liveness Detection and Failover on page 438

- DNS Cache on page 438
- DHCP Server on page 439
- Enabling Log Settings on page 439
- Monitoring Packets on page 439

Supported Interfaces

The DNS proxy feature is supported on physical interfaces, VLAN interfaces, or VLAN trunk interfaces. The zone for each interface should only be LAN, DMZ, or WLAN.

DNS Server Liveness Detection and Failover

When multiple DNS servers are configured, to determine the “best” server, SonicOS considers these factors:

- DNS server priority.
- DNS server status (up, down, unknown).
- Time duration after failover.

DNS Cache

In DNS Proxy, a DNS cache memory saves the most commonly used domains and host addresses, and when it receives the DNS query that match the domain in DNS cache, the Security Appliance directly responds to clients by using the cache records, without processing DNS query and reply proxy.

There are two kinds of DNS Cache:

Static	Manually configured by you.
Dynamic	Auto-learned by SonicOS. For each DNS Query, SonicOS DNS Proxy does the deep inspection on the URI and records the valid response to the caches.

When a DNS query matches an existing cache entry, SonicOS DNS Proxy responds directly with the cached URI. This usually decreases the network traffic and, thus, improves overall network performance.

Maximum DNS Proxy Cache Size

Static DNS Proxy Cache Size

Static DNS proxy cache entry size is always 256 regardless of platform. The static DNS cache is never deleted unless it is done manually.

Dynamic DNS Proxy Cache Size

Dynamic DNS proxy cache size depends on the platform, as shown in Dynamic Cache Size.

Dynamic Cache Size

Platform	Maximum Cache Size
SM 9600/SM 9400	4096
SM 9200	2048
NSA 6600/NSA 5600/NSA 4600	2048

Dynamic Cache Size (Continued)

Platform	Maximum Cache Size
NSA 3600/NSA 2600/NSA 2650	1024
TZ600	512
TZ500/TZ500 W/TZ400/TZ400 W/TZ350/ TZ350 W/TZ300/TZ300 W	512
SOHO 250/SOHO 250 W/SOHO W	512

If the maximum DNS proxy cache size has been reached when the Security Appliance attempts to add an entry to the proxy cache, the Security Appliance:

1 Deletes the DNS proxy ache entry with the earliest expire time.

2 Adds the new DNS proxy cache entry.

High Availability Stateful Synchronization of DNS Cache

DNS proxy supports stateful synchronization of DNS proxy cache. When the DNS proxy cache is added, deleted, or updated dynamically, it synchronizes to the idle Security Appliance.

DHCP Server

When DNS Proxy is enabled on an interface, the device needs to push the interface IP as a DNS server address to clients, so the DHCP server must be configured manually, using the interface address as the **DNS Server 1** address in the **DHCP Server** settings on the **DNS/WINS** tab. The **Interface Pre-Populate** option in the **Dynamic Range Configuration** dialog makes this easy to configure; if the selected interface has enabled DNS Proxy, the DNS server IP is added automatically into the DNS/WINS page. For how to configure a DHCP server statically, see Configuring Static DHCP Entries on page 534.

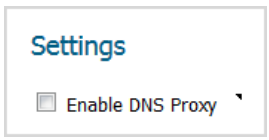
Enabling Log Settings

Several events logs are related to DNS Proxy and need to be configured as described in *SonicOS Investigate*.

Monitoring Packets

The process of DNS Proxy is monitored with **Dashboard > Packet Monitor**. For information about the Packet Monitor, see *SonicOS Investigate*.

Enabling DNS Proxy



Enabling DNS Proxy must be done first globally on the **Network > DNS Proxy** page and then on each interface. This provides a gradual control to enable the feature for different network segments independently

To enable DNS Proxy:

1 Navigate to **Network > DNS Proxy**.

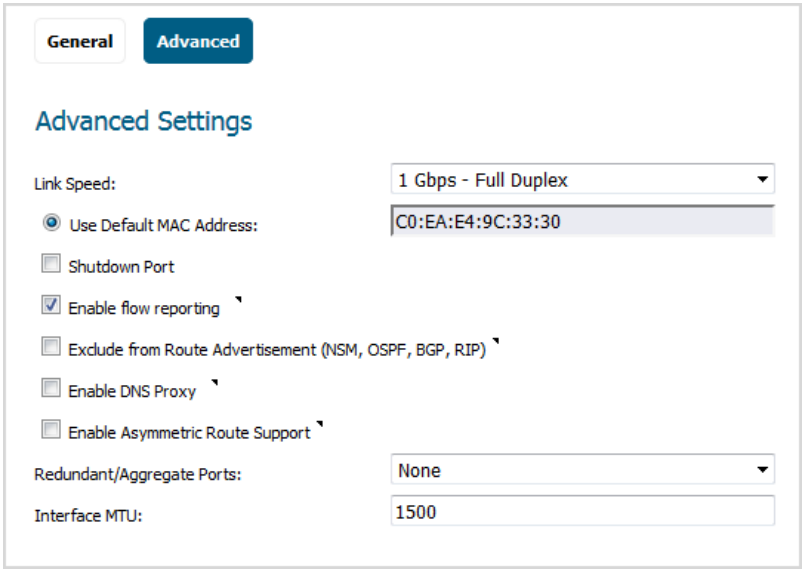
2 Select **Enable DNS Proxy**. This option is not selected by default.

3 Click **Accept**.

4 Navigate to **Network > Interfaces**.

5 Click the **Edit** icon for the interface on which to enable DNS Proxy. The **Edit Interface** dialog displays.

6 Click **Advanced**.



7 Select **Enable DNS Proxy**. This option displays only when DNS Proxy is enabled globally.

8 Click **OK**.

9 Repeat Step 5 through Step 8 for each interface on which to enable DNS Proxy.

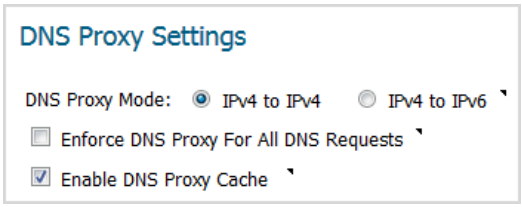
10 Click **Accept**.

For the Access Rules associated with the interface, see the *SonicOS 6.5 Policies* administration documentation.

Configuring DNS Proxy Settings

To configure DNS Proxy:

1 Navigate to **Network > DNS Proxy | DNS Proxy Settings**.



2 From **DNS Proxy Mode**, choose the IP version for sending/receiving DNS proxy packets between the Security Appliance and the DNS servers:

- **IPv4 to IPv4** (default)
- **IPv4 to IPv6**

3 To allow all types of DNS requests, including stack DNS packets sent by SonicOS, to be processed by DNS Proxy, including the forwarding of DNS queries with a destination address of outside DNS servers, select

Enforce DNS Proxy for All DNS Requests. If this option is disabled, only those requests destined for SonicWall Security Appliances are processed. This option is not selected by default.

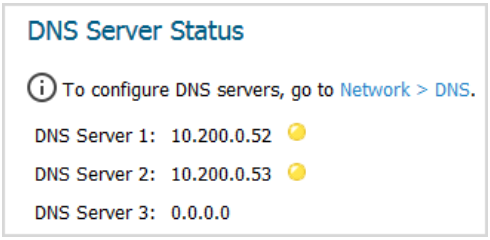
NOTE: This option affects only DNS over UDP. If this option is not selected, only DNS proxy requests destined for a SonicWall Security Appliance are enabled.

4 For DNS over UDP requests only, select **Enable DNS Cache**. This option is selected by default.

5 Click **Accept**.

NOTE: There are several advanced settings, such as DNS Proxy protocol, that can be configured. For more information about these settings, contact Technical Support.

Monitoring DNS Server Status

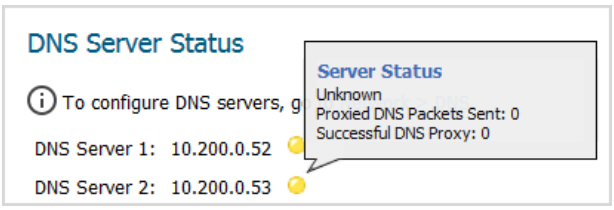


NOTE: A configured DNS Server has its IP address displayed. If a server is not configured, the IP address is 0.0.0.0. To configure a server, click the link to **Network > DNS**; see Configuring DNS Settings on page 421.

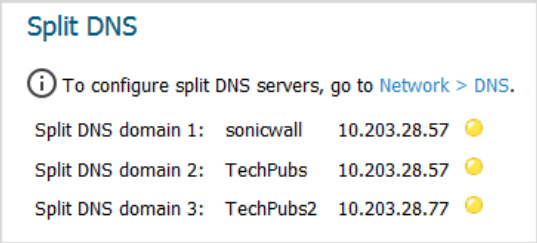
You monitor the status of each configured upstream DNS Servers in the **DNS Server Status** section. The server status is decided by DNS reply from the server:

Up (green LED)	The reply was successful.
Unknown (yellow LED)	A DNS reply has not been received by the server.
Down (red LED)	The failure count exceeded the limit of 20. The status remains down until the next successful DNS query.

Moving the mouse over the LED displays a popup with further information about the number of proxied DNS packets sent and the number of successful DNS Proxy queries:



Monitoring Split DNS Server Status

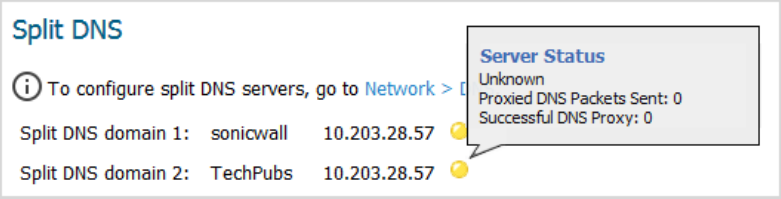


NOTE: A configured split DNS server has its IP address displayed. To configure a split server, click the link to **Network > DNS**; see Configuring DNS Settings on page 421.

You monitor the status of each configured upstream DNS Servers in the **Split DNS** section. The server status is decided by DNS reply from the server:

Up (green LED)	The reply was successful.
Unknown (yellow LED)	A DNS reply has not been received by the server.
Down (red LED)	The failure count exceeded the limit of 20. The status remains down until the next successful DNS query.

Moving the mouse over the LED displays a popup with further information about the number of proxied DNS packets sent and the number of successful DNS Proxy queries:



Viewing and Managing Static DNS Cache Entries



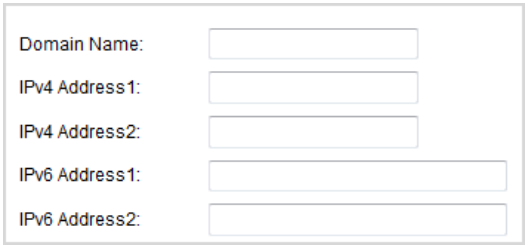
Domain Name	Name of the domain.
IPv4 Address 1	Primary IPv4 address of Static DNA cache. 0.0.0.0 if not specified.
IP4 Address 2	Secondary IPv4 address of Static DNA cache. 0.0.0.0 if not specified.
IPv6 Address 1	Primary IPv6 address of Static DNA cache. :: if not specified.
IPv6 Address 2	Secondary IPv6 address of Static DNA cache. :: if not specified.
Configure	Contains the **Edit** and **Delete** icons for each entry.

To add static DNS cache entries:

1 Navigate to **Network > DNS Proxy**.

2 Scroll to **Static DNS Proxy Cache Entries**.

3 Click **Add** either above or below the table. The **Add Static DNS Cache** dialog displays.



4 Enter a name in the **Domain Name** field.

5 For IPv4 static DNS cache, enter the primary IPv4 address in the **IPv4 Address 1** field.

6 Optionally, for IPv4 static DNS cache, enter the secondary IPv4 address in the **IPv4 Address 2** field.

7 For IPv6 static DNS cache, enter the primary IPv6 address in the **IPv6 Address 1** field.

8 Optionally, for IPv6 static DNS cache, enter the secondary IPv6 address in the **IPv4 Address 2** field.

9 Click **OK**.

10 To add another static DNS cache entry, repeat Step 4 through Step 9.

11 Click **Cancel**.

Deleting Static DNS Cache Entries

To delete a static DNS cache entry:

1 Click the entry's **Delete** icon.

To delete two or more static DNS cache entries:

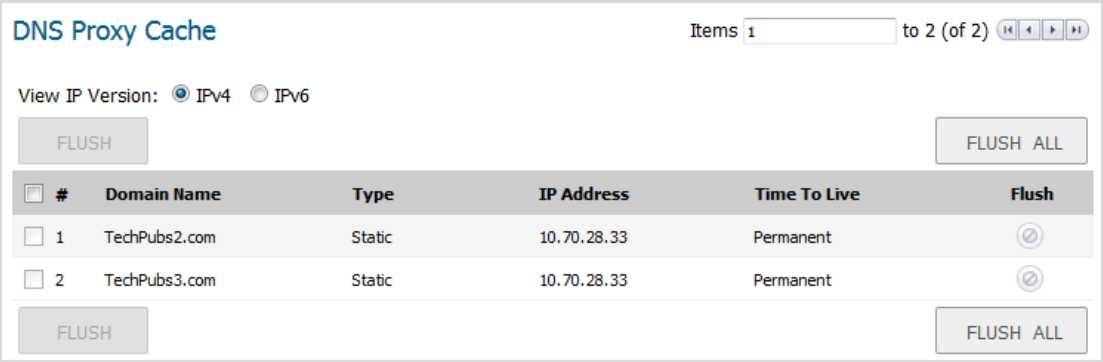
1 Select the checkboxes of the entries to be deleted. **Delete** become available.

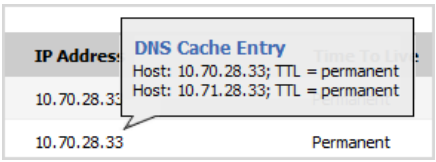
2 Click **Delete** .

To delete all static DNS cache entries:

1 Click **Delete All** .

Viewing DNS Proxy Cache Entries



View IP Version	Select either **IPv4** or **IPv6**.
Domain Name	Name of the DNS Server.
Type	**Dynamic** **Static**
IP Address	IPv4 or IPv6 address of the DNS Server. Mousing over an entry displays **Host** and Time to Live (**TTL**) information for the entry: 
Time to Live	Either: • **Expires in *n minutes x seconds*** (Dynamic DNS) • **Expired** (Dynamic DNS) • **Permanent** (Static DNS)
Flush	**Flush** icon for each entry.

Dynamic DNS cache is added automatically during the DNS Proxy process; static DNS cache is added when you configure it. Dynamic DNS cache has a TTL value and can be flushed. Static DNS cache must be deleted; see Deleting Static DNS Cache Entries on page 443

Flushing Dynamic DNS Cache Entries

To flush a dynamic DNS cache entry:

1 Click the entry's **Flush** icon.

To flush two or more dynamic DNS cache entries:

1 Select the checkboxes of the entries to be deleted. **Flush** become available.

2 Click **Flush**.

To flush all dynamic DNS cache entries:

1 Click **Flush All**.

21

Configuring DNS Security

Topics:

- About Sinkholes on page 445
- Network > DNS Security on page 445
  - Configuring DNS Security Settings on page 446
  - Deleting Entries in the Lists on page 447

About Sinkholes

A DNS sinkhole, also known as a sinkhole server, Internet sinkhole, or BlackholeDNS, is a DNS server that gives out false information to prevent the use of the domain names it represents. DNS sinkholes are effective at detecting and blocking malicious traffic, and used to combat bots and other unwanted traffic.

SonicOS now provides the ability to configure a sinkhole with black and white lists.

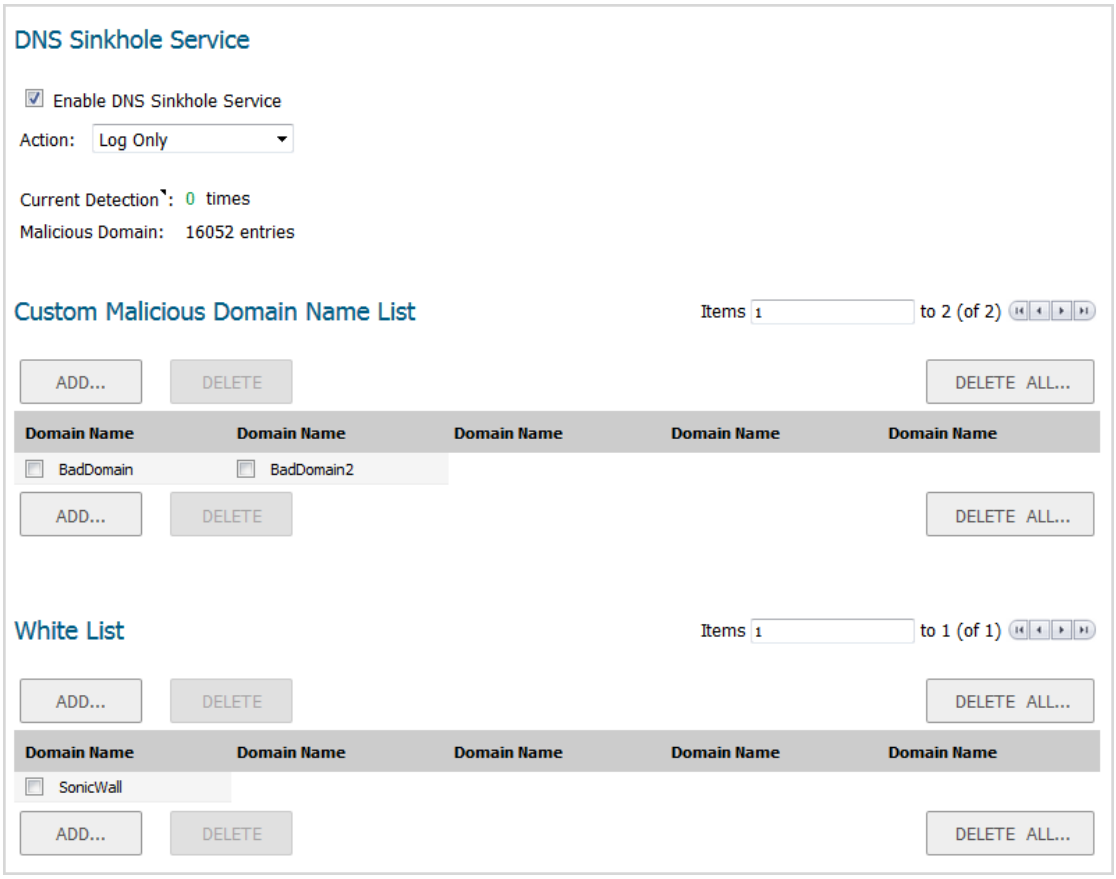
Network > DNS Security

MANAGE | System Setup > Network > DNS Security allows you to manually configure your DNS security settings.

Configuring DNS Security Settings

To configure DNS Security settings:

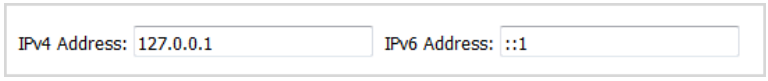
1 Navigate to **MANAGE | System Setup > Network > DNS Security**.



2 Select **Enable DNS Sinkhole Service**. This option is selected by default.

3 Select what the service is to do from **Action**:

- **Log Only**
- **Negative Reply**
- **Reply with Forged IP** – Two fields display:



a) Enter the IPv4 and IPv6 addresses in the fields.

4 Scroll to **Custom Malicious Domain Name List**.

5 CLICK **ADD**. THE **ADD ONE DOMAIN** NAME DIALOG DISPLAYS.



6 Click **OK**.

7 Enter the malicious domain name in the **Domain Name** field.

8 Repeat Step 5 through Step 7 for each domain name.

9 Scroll to **White LIst**.

10 CLICK **ADD**. THE **ADD ONE WHITE ENTRY** DIALOG DISPLAYS.



11 Enter the white-list domain name in the **Domain Name** field.



12 Click **OK**.

13 Repeat Step 10 through Step 7 for each domain name.

14 Click **ACCEPT** to save your changes.

Deleting Entries in the Lists

To delete a list's entries.

1 Navigate to **MANAGE | System Setup > Network > DNS Security**.

2 Select the entry to delete. **DELETE** and **DELETE ALL** become available.

3 Click the appropriate button.

About DNS Tunneling Detection

DNS tunneling is a method of bypassing security controls and exfiltrating data from a targeted organization. A DNS tunnel can be used as a full remote-control channel for a compromised internal host. Capabilities include Operating System (OS) commands, file transfers or even a full IP tunnel.

SonicOS provides the ability to detect DNS tunneling attacks, displays suspicious clients, and allows you to create white lists for DNS tunnel detection.

When DNS tunneling detection is enabled, SonicOS logs whenever suspicious DNS packets are dropped.

Topics:

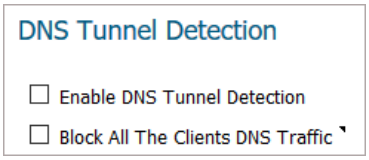
- Enabling DNS Tunneling Detection on page 447
- Viewing Detected Suspicious Clients on page 448
- Creating DNS Tunnel Detection White Lists on page 448
- Deleting DNS Tunnel Detection White List Entries on page 449

Enabling DNS Tunneling Detection

To configure DNS tunneling detection:

1 Navigate to **MANAGE | System Setup > Network > DNS Security**.

2 Scroll to the **DNS Tunnel Detection** section.



3 To enable DNS tunnel detection, select **Enable DNS Tunnel Detection**.

4 To block all the detected clients' DNS traffic, select **Block All The Clients DNS Traffic**.

5 Click **ACCEPT**.

Viewing Detected Suspicious Clients

SonicOS displays information about all hosts that have established a DNS tunnel in the **Detected Suspicious Clients Information** table.

TIP: This table is populated only if DNS tunnel detection is enabled. Hosts are dropped only if blocking clients DNS traffic is enabled. See Enabling DNS Tunneling Detection on page 447.



IP Address	IP address of the suspicious client
MAC Address	MAC address of the suspicious client
Detection Method	DNS type used to detect suspicious clients: • **Normal DNS Type**: A, AAAA, CNAME • **Corner DNS Type**: such as TXT, NULL, SRV, PRIVATE, MX
Interface	Interface on which the host establishing the DNS tunnel was detected
Block	Indicates whether the host was blocked

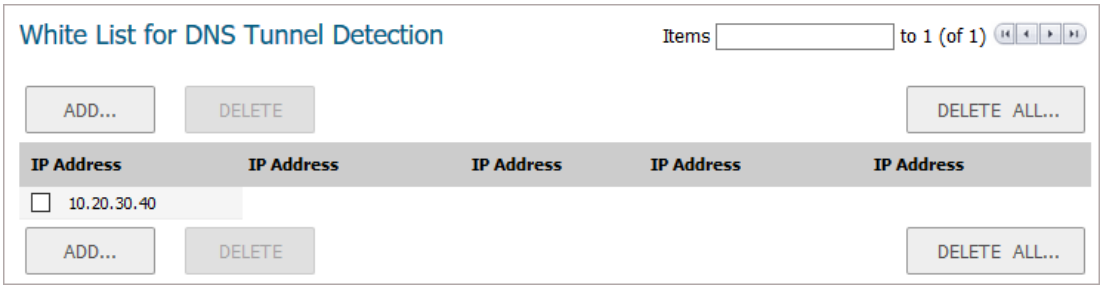
Creating DNS Tunnel Detection White Lists

You can create white lists for IP address you consider safe. If a detected DNS tunnel IP address matches an address in the white list, DNS tunnel detection is bypassed.

To create a DNS

1 Navigate to **MANAGE | System Setup > Network > DNS Security**.

2 Scroll to **White List for DNS Tunnel Detection**.



3 Click **ADD**. The **Add One White List Entry** dialog displays.



4 Click **OK**.

5 Repeat Step 3 and Step 4 for each white list entry.

Deleting DNS Tunnel Detection White List Entries

To delete all DNS tunnel detection white list entries:

1 Navigate to **MANAGE | System Setup > Network > DNS Security**.

2 Scroll to **White List for DNS Tunnel Detection**.

3 Click **DELETE ALL**.

To delete one or more DNS tunnel detection white list entries:

1 Navigate to **MANAGE | System Setup > Network > DNS Security**.

2 Scroll to **White List for DNS Tunnel Detection**.

3 Select one or more DNS tunnel detection white list entries. **DELETE** becomes available.

4 Click **DELETE**.

22

Configuring Route Advertisements and Route Policies

IMPORTANT: For SD-WAN routing and route policies, see Configuring SD-WAN Route Policies on page 591.

Topics:

- About Routing on page 451
  - About Metrics and Administrative Distance on page 451
  - Route Advertisement on page 452
  - ECMP Routing on page 453
  - Policy-based Routing on page 453
  - Policy-based TOS Routing on page 453
  - PBR Metric-based Priority on page 454
  - Policy-based Routing and IPv6 on page 455
  - OSPF and RIP Advanced Routing Services on page 455
  - Drop Tunnel Interface on page 463
- **Network > Routing** on page 464
  - **Network > Routing > Settings** on page 464
  - **Network > Routing > Route Advertisement** on page 466
  - **Network > Routing > OSPFv2** on page 467
  - **Network > Routing > RIP** on page 469
  - **Network > Routing > OSPFv3** on page 470
  - **Network > Routing > RIPng** on page 472
- Configuring Routing on page 472
  - Prioritizing Routes by Metric on page 473
  - Configuring Metrics for Default Routes Learned through Router Advertisement on page 473
  - Configuring Route Advertisement on page 474
  - Configuring Static and Policy-based Routes on page 475
  - Configuring a Static Route for a Drop Tunnel Interface on page 479
  - Configuring OSPF and RIP Advanced Routing Services on page 481
  - Configuring BGP Advanced Routing on page 490

About Routing

SonicWall Security Appliances support these routing protocols:

- RIPv1 (Routing Information Protocol)
- RIPv2
- OSPFv2 (Open Shortest Path First)
- OSPFv3
- PBR (Policy-Based Routing)

Topics:

- About Metrics and Administrative Distance on page 451
- Route Advertisement on page 452
- ECMP Routing on page 453
- Policy-based TOS Routing on page 453
- PBR Metric-based Priority on page 454
- Policy-based Routing and IPv6 on page 455
- OSPF and RIP Advanced Routing Services on page 455
- Policy-based Routing and IPv6 on page 455

About Metrics and Administrative Distance

Metrics and administrative distance affect network performance, reliability, and circuit selection.

About Metrics

A *metric* is a weighted cost assigned to static and dynamic routes. Metrics determine the best route among several, usually the gateway with the lowest metric. This gateway is usually the default gateway.

Metrics have a value between 1 and 254; see Metric value descriptions. Lower metrics are considered better and take precedence over higher costs. SonicOS adheres to Cisco-defined metric values for directly connected interfaces, statically encoded routes, and all dynamic IP routing protocols.

Metric value descriptions

Metric value	Description
1	Static Route
5	EIGRP Summary
20	External BGP
90	EIGRP
100	IGRP
110	OSPF
115	IS-IS
120	RIP
140	EGP

Metric value descriptions

Metric value	Description
170	External EIGRP
200	Internal BGP

About Administrative Distance

Administrative distance (admin distance) is a value that influences which source of routes should be used for two identical routes from different sources. The lower the administrative distance value, the more trusted the route.

The admin distance, when set, is only used by the ZebOS components when choosing which routes to:

- Populate into PBR
- Redistribute to other routing protocols when a static route competes with a route received from a particular routing protocol.

The admin distance is not used for prioritizing routes within PBR itself, so unless dynamic routing is in use, the admin distance set for a static route has no effect. When dynamic routing is being used, the admin distance provides a mechanism by which static routes defined in PBR can be compared to otherwise equivalent dynamic routes possibly received from protocols such as OSPF, RIP, or BGP. By default, the admin distance of a PBR static route inserted into the network services module (NSM) is equal to the metric defined for the PBR route. The admin distance of each static route may optionally be set to a different value when a custom value is entered for Admin Distance.

For example, if a simple (destination only) static route (for example, destination = `14.1.1.0/24`) is defined with a metric of 10 and the admin distance set to its default of **Auto**, that route is populated into NSM with an admin distance and metric of 10.

Now assume the same `14.1.1.0/24` route is received from both RIP and OSPF. RIP routes have a default admin distance of 120 and OSPF routes 110, so the static route, with a default admin distance (== the metric) of 10 would be preferred over both routes, and NSM would not populate either the OSPF or RIP route into PBR. If the admin distance of the static route had been set to 115 (keeping the metric at 10), however, then the OSPF route (at 110) would be preferred over the static route, but the RIP route would not. If the OSPF route disappeared, NSM would withdraw the OSPF route and would not populate the RIP route as its 120 AD is greater than the static route's 115 AD.

In either of the above cases, the static route is still preferred in PBR because all non-default routes populated into PBR from NSM are added with a 110 metric, which is greater than the metric of 10 for the static route.

If an admin distance of 110 and a metric > 110 are used for the static routes, the metric value passed to NSM would be used by OSPF when it compares the metric of the static route to the OSPF metric (or cost) of any competing OSPF route.

Route Advertisement

SonicWall Security Appliances use RIPv1 or RIPv2 to advertise its static and dynamic routes to other routers on the network. Changes in the status of VPN tunnels between the Security Appliance and remote VPN gateways are also reflected in the RIPv2 advertisements. Based on your router's capabilities or configuration, choose between:

- RIPv1, which is an earlier version of the protocol, has fewer features, and sends packets through broadcast instead of multicast.
- RIPv2, which is a later version of the protocol, includes subnet information when multicasting the routing table to adjacent routers and route tags for learning routes. RIPv2 packets are backwards compatible and

can be accepted by some RIPv1 implementations that provide an option of listening for multicast packets. The RIPv2 Enabled (broadcast) selection, which broadcasts packets instead of multicasting them, is for heterogeneous networks with a mixture of RIPv1 and RIPv2 routers.

ECMP Routing

SonicOS 6.5 supports equal-cost multi-path (ECMP) routing, a technique for routing packets along multiple paths of equal cost. The forwarding engine identifies paths by next-hop. When forwarding a packet, the router must decide which next-hop (path) to use. Multi-path routing can be used in conjunction with most routing protocols.

In SonicOS, you can use ECMP routing to specify multiple next hops for a given route's destination. In environments with substantial requirements, there are several reasons for doing this. A router could just use one ISP most of the time, and switch to the other when the first one fails for some reason. Another application of multi-path is to keep a path on standby and enable it only when bandwidth requirements surpass a predefined threshold. SonicOS supports up to four next-hop paths.

Various routing protocols, including Open Shortest Path First (OSPF) and Intermediate System to Intermediate System (ISIS), explicitly allow ECMP routing. Some router implementations also allow equal-cost multi-path usage with RIP and other routing protocols.

Policy-based Routing

A simple static routing entry specifies how to handle traffic that matches specific criteria, such as destination address, destination mask, gateway to forward traffic, the interface that gateway is located, and the route metric. This method of static routing satisfies most static requirements, but is limited to forwarding based only on destination addressing.

Policy-based Routing (PBR) allows you to create extended static routes to provide more flexible and granular traffic handling capabilities. SonicOS PBR allows for matching based upon source address, source netmask, destination address, destination netmask, service, interface, and metric. This method of routing allows for full control of forwarding based upon a large number of user defined variables.

A FQDN can be used as the source or destination of the PBR entry.

Policy-based TOS Routing

SonicOS supports policy-based TOS (type of service) routing when defining policy-based routing (PBR) policies by Type of Service (TOS) and TOS mask values. When defined, the TOS and mask values are compared against the associated IP packet's TOS/DSCP field in the IP header when finding a route match.

The TOS value is compared to an 8-bit field in the IP packet header (for information about this header, see *RFC 2474, Differentiated Services*, and *RFC 2168, Explicit Congestion Notification*). The TOS value can be used to define services relating to quantitative performance requirements (for example, peak bandwidth) and those based on relative performance (for example, class differentiation).

TOS routing differs from existing SonicOS QoS marking, which does not affect the routing of a packet and cannot forward packets differently based on an inbound packet's TOS field. TOS Routing provides this capability by allowing policy routes to define a TOS Value/TOS Mask pair to be compared to inbound packets for differential forwarding. TOS routing only applies to packets as they enter the Security Appliance.

With TOS routing, it is possible to define multiple policy routes with identical source IP, destination IP, and service values, but differing TOS/TOS mask values. This allows packets with marked TOS fields to be forwarded differently based on the value of the TOS field in the inbound packet.

Any PBR policy routes defined before SonicOS 6.5 have no values defined for the TOS/TOS mask. Likewise, the default values for TOS/TOS mask fields are zero (no values defined).

Policy routes with a TOS value other than zero are prioritized before all simple destination-only routes, but below any policy routes that define a source or service. When comparing two TOS Policy routes, and assuming both have the same set of source, destination, and service values either defined or not defined, the TOS route with the greater number of TOS mask bits set to 1 is prioritized before TOS routes with fewer TOS mask bits set.

The general prioritization (high to low) of PBR routes is as follows, based on the policy fields defined as anything other than **Any** or zero for TOS:

Destination, Source, Service, TOS
Destination, Source, Service
Destination, Source, TOS
Destination, Source
Destination, Service, TOS
Destination, Service
Destination, TOS
Destination
Source, Service, TOS
Source, Service
Source, TOS
Source
Service, TOS
Service
TOS

PBR Metric-based Priority

SonicOS supports a metric weighted cost assigned to a route policy for policy-based routing (PBR) that allows the configured metric to take precedence in route prioritization over the route specificity that used by default. Metrics have a value between 0 and 255. Lower metrics are considered better and take precedence over higher ones.

The general prioritization (high to low) of PBR routes is as follows, based on the policy fields defined as anything other than **Any**, or zero for TOS:

Destination, Source, Service, TOS
Destination, Source, Service
Destination, Source, TOS
Destination, Source
Destination, Service, TOS
Destination, Service
Destination, TOS
Destination
Source, Service, TOS
Source, Service

Source, TOS

Source

Service, TOS

Service

TOS

Within these 15 classifications, routes are further prioritized based on the cumulative specificity of the defined route entries. For the source and destination fields, specificity is measured by counting the number of IP addresses represented in the address object. For example, the network address object, `10.0.0.0/24`, would include 256 IP addresses, while the network address object, `10.0.0.0/20`, would represent 4096. The longer `/24` (24 bit) network prefix represents fewer host IP addresses and is more specific.

The new metric-weighted option allows the configured metric to take precedence in prioritization over the route specificity. With the option enabled, the precedence used during prioritization is as follows (high to low):

1. Route class (determined by the combination of source, destination, service, and TOS fields with values other than Any or zero)
2. The value of the Metric
3. The cumulative specificity of the source, destination, service, and TOS fields

Policy-based Routing and IPv6

For complete information on the SonicOS implementation of IPv6, see IPv6 on page 932.

Policy-based Routing is fully supported for IPv6 by selecting IPv6 address objects and gateways for route policies on **Network > Routing**. You can switch the entries in the **Route Policies** table between **IPv4** and **IPv6**.

Routing Information Protocol next generation (RIPng) is an information routing protocol for IPv6, which allows routers to exchange information for computing routes through an IPv6-based network.

For information on route advertisement, see Route Advertisement on page 452. For information on setting up Route Policies, see Route Advertisement on page 452.

OSPF and RIP Advanced Routing Services

In addition to Policy-based Routing and RIP advertising, SonicOS offers the option of enabling Advanced Routing Services (ARS). Advanced Routing Services provides full advertising and listening support for the Routing Information Protocol (RIPv1 - RFC1058) and (RIPv2 - RFC2453), and Open Shortest Path First (OSPFv2 – RFC2328). Advanced Routing Service should only be enabled by those environments requiring support for either or both of these dynamic routing protocols.

RIP and OSPF are Interior Gateway Protocols (IGP) that are both widely used by networks of various sizes to automate the process of route distribution. RIP is commonly used within smaller networks, while OSPF is used by larger networks, although network size should not be the only factor used to determine the appropriateness of one protocol over the other – network speed, interoperability requirements, and relative overall complexity, for example, should also be considered. RIPv1 and RIPv2 are both supported by ARS, the largest differences between the two being that RIPv2 supports VLSM (Variable Length Subnet Masks), authentication, and routing updates. Routing Information Protocol differences illustrates the major differences between RIPv1, RIPv2, and OSPFv2/OSPFv3:

Routing Information Protocol differences

	RIPv1	RIPv2	OSPFv2/OSPFv3
Protocol metrics	Distance Vector	Distance Vector	Link State
Maximum Hops	15	15	Unlimited
Routing table updates	Full table broadcast periodically, slower convergence	Full table broadcast or multicast periodically, slower convergence	Link state advertisement multicasts, triggered by changes, fast convergence
Subnet Sizes Supported	Only class-based (a/b/c) subnets support	Class-based only	VLSM
Autonomous system topology	Indivisible and flat	Indivisible and flat	Area-based, allowing for segmentation and aggregation

Topics:

- About Routing Services on page 456
- OSPF Terms on page 459

About Routing Services

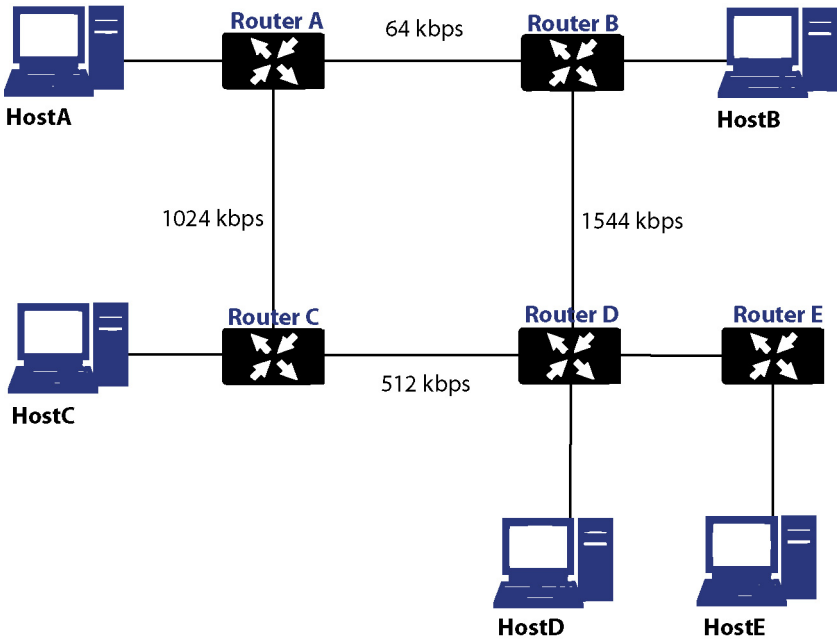
Topics:

- Protocol Type on page 456
- Maximum Hops on page 457
- Split-Horizon on page 458
- Poison Reverse on page 458
- Routing Table Updates on page 458
- Subnet Sizes Supported on page 458
- Autonomous System Topologies on page 459

Protocol Type

Distance Vector protocols such as RIP base routing metrics exclusively on hop counts, while Link state protocols such as OSPF consider the state of the link when determining metrics. For example, OSPF determines interface metrics by dividing its reference bandwidth (100mbits by default) by the interface speed – the faster the link, the lower the cost and the more preferable the path. Consider the example network shown in Example network for determining lowest cost route:

Example network for determining lowest cost route



In the sample network shown in Example network for determining lowest cost route, if Host A wanted to reach Host B, with RIP, the lowest cost route would be from Router A to Router B, across the relatively slow 64kbps link. With OSPF, the cost from Router A to Router B would be 1562, while the cost from Router A to Router C to Router D to Router B would be 364, making it the preferred route.

Maximum Hops

RIP imposes a hop count of 15 to help prevent routing loops which can occur when bad (for example, stale) routing information is broadcast and propagated through a network either due to misconfiguration, or slow convergence. Consider if the link between Router D and Router E failed in the example in Example network for determining lowest cost route, and there were no safeguards in place:

- Router A's routing information states that it can reach Network E through Router B or Router C with a metric of 3.
- When the link between Router D and Router E fail, and Router A broadcasts its routing information, Router B and Router C determine that they can reach Network E through Router A with a metric of 4.
- Router B and Router C broadcast this information, and it is received by Router D which then determines it can reach Network E through Router B or Router C with a metric of 5.
- This loop continues until the hop count of 16 (infinity) is reached.

Other measures against this sort of situation are also commonly employed by RIP, including:

- Split-Horizon on page 458
- Poison Reverse on page 458
- Routing Table Updates on page 458
- Subnet Sizes Supported on page 458
- Autonomous System Topologies on page 459

Split-Horizon

A preventative mechanism where routing information learned through an interface is not sent back out the same interface. This generally works well on broadcast links, but not on non-broadcast links such as Frame Relay, where a single link can commonly be used to reach two separate autonomous systems.

Poison Reverse

Also known as route poisoning, an extension of split-horizon where a network is advertised with a metric of 16 (unreachable), helping to ensure that incorrect alternative routes are not propagated.

OSPF does not have to impose a hop count limit because it does not advertise entire routing tables, rather it generally only sends link state updates when changes occur. This is a significant advantage in larger networks in that it converges more quickly, produces less update traffic, and supports an unlimited number of hops.

Routing Table Updates

As mentioned above, the practice of sending an entire routing table introduces the problems of slower convergences, higher bandwidth utilization, and increased potential for stale routing information. RIPv1 broadcasts its entire routing table at a prescribed interval (usually every 30 seconds), RIPv2 can either broadcast or multicast, and OSPF multicasts only link state updates whenever a change to the network fabric occurs. OSPF has a further advantage of using designated routers (DR) in forming adjacencies in multiple-access networks (more on these concepts later) so that updates do not have to be sent to the entire network.

Subnet Sizes Supported

RIPv1 was first implemented when networks were strictly class A, class B, and class C (and later D and E):

Class A `1.0.0.0` to `126.0.0.0` (`0.0.0.0` and `127.0.0.0` are reserved)

- Left most bit 0; 7 network bits; 24 host bits
- 0nnnnnnn hhhhhhhh hhhhhhhh hhhhhhhh (8-bit classful netmask)
- 126 Class A networks, 16,777,214 hosts each

Class B `128.0.0.0` to `191.255.0.0`

- Left most bits 10; 14 network bits; 16 host bits
- `10nnnnnn nnnnnnnn hhhhhhhh hhhhhhhh` (16-bit classful netmask)
- 16,384 Class B networks, 65,532 hosts each

Class C `192.0.0.0` to `223.255.255.0`

- Left most bits 110; 21 network bits; 8 host bits
- 110nnnnn nnnnnnnn nnnnnnnn hhhhhhhh (24-bit classful netmask)
- 2,097,152 Class Cs networks, 254 hosts each

Class D `225.0.0.0` to `239.255.255.255` (multicast)

- Left most bits 1110; 28 multicast address bits
- 1110mmmm mmmmmmmm mmmmmmmm mmmmmmmm

Class E `240.0.0.0` to `255.255.255.255` (reserved)

- Left most bits 1111; 28 reserved address bits
- 1111rrrr rrrrrrrr rrrrrrrr rrrrrrrr

This method of address allocation proved to be very inefficient because it provided no flexibility, neither in the way of segmentation (subnetting) or aggregation (supernetting, or CIDR – classless inter-domain routing) by means of VLSM – variable length subnet masks.

VLSM, supported by RIPv2 and OSPF, allows for classless representation of networks to break larger networks into smaller networks:

> For example, take the classful `10.0.0.0/8` network, and assign it a `/24` netmask. This subnetting allocates an additional 16-bits from the host range to the network range (24-8=16). To calculate the number of additional networks this subnetting provides, raise 2 to the number of additional bits: 2^16=65,536. Thus, rather than having a single network with 16.7 million hosts (usually more than most LAN's require) it is possible to have 65,536 networks, each with 254 usable hosts.

VLSM also allows for route aggregation (CIDR):

> For example, if you had 8 class C networks: `192.168.0.0/24` through `192.168.7.0/24`, rather than having to have a separate route statement to each of them, it would be possible to provide a single route to `192.168.0.0/21` which would encompass them all.

This ability, in addition to providing more efficient and flexible allocation of IP address space, also allows routing tables and routing updates to be kept smaller.

Autonomous System Topologies

An autonomous system (AS) is a collection of routers that are under common administrative control and that share the same routing characteristics. When a group of autonomous systems share routing information, they are commonly referred to as a confederation of autonomous systems. (RFC1930 and RFC975 address these concepts in much greater detail). In simple terms, an AS is a logical distinction that encompasses physical network elements based on the commonness of their configurations.

With regard to RIP and OSPF, RIP autonomous systems cannot be segmented, and all routing information must be advertised (broadcast) through the entire AS. This can become difficult to manage and can result in excessive routing information traffic. OSPF, on the other hand, employs the concept of Areas, and allows for logically, manageable segmentation to control the sharing of information within an AS. An Area ID is an administrative identifier. OSPF areas begin with the backbone area (area 0 or `0.0.0.0`), and all other areas must connect to this backbone area (although there are exceptions). This ability to segment the routing AS helps to ensure that it never becomes too large to manage, or too computationally intensive for the routers to handle.

OSPF Terms

OSPF is substantially more complicated to configure and maintain than RIP. The following concepts are critical to understanding an OSPF routing environment:

- **Link state** – As it pertains to OSPF, a link is an egress interface on a router, and the state describes characteristics of that interface, such as its cost. Link states are sent in the form of Link State Advertisements (*LSA*) which are contained within Link State Update (*LSU*) packets, one of five types of OSPF packets.
- **Cost** – A quantification of the overhead required to send a packet along a particular link. Cost is calculated by dividing a reference bandwidth (usually 100mbit, or 10^8 bit) by an interface's speed. The lower the cost, the more preferable the link. Some common path costs are shown in Cost calculation for different interfaces.

Cost calculation for different interfaces

Interface	Divided by 10^8 (100mbit) = OSPF Cost
Fast Ethernet	1
Ethernet	10
T1 (1.544mbit)	64
DSL (1mbit)	100
DSL (512kbps)	200

Cost calculation for different interfaces

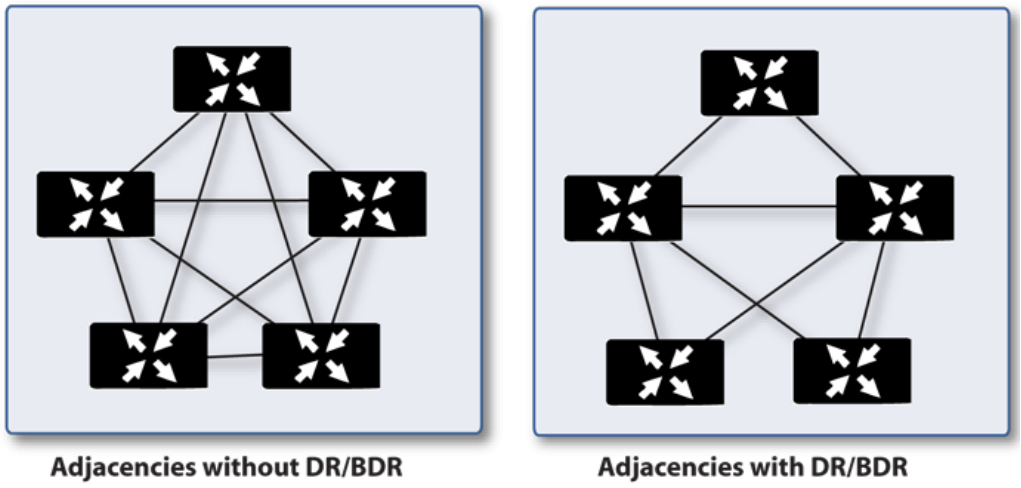
Interface	Divided by 10^8 (100mbit) = OSPF Cost
64kbps	1562
56kbps	1785

- **Area** – The network comprising the group of OSPF routers intended to share a common Link State Database. OSPF networks are built around the backbone area (area 0, or 0.0.0.0) and all other areas must connect to the backbone area (unless virtual links are used, which is generally discouraged). Area assignment is interface specific on an OSPF router; in other words, a router with multiple interfaces can have those interfaces configured for the same or different areas.
- **Neighbors** – OSPF routers on a common network segment have the potential to become neighbors by means of sending Hello packets. Hello packets act as a form of advertisement and identification, and if two OSPF routers share a common set of certain characteristics, they become neighbors upon seeing their own router ID in the other router's Hello packet. Hello packets are also used in the *DR* (Designated Router) and *BDR* (Backup Designated Router) election process. For two routers to become neighbors, the characteristics that they must have in common are:
  - **Area-ID** – An area ID identifies an OSPF *area* with a 32-bit value, and is generally represented in an IP address format. OSPF requires at a minimum the backbone area, area 0 (or 0.0.0.0) for operation.
  - **Authentication** – Authentication types can generally be set to none, simple text, or MD5. When using simple text, authentication should be used only for identification, as it is sent in the clear. For security, MD5 should be used.
  - **Timer intervals** – Hello and Dead intervals must be the same. The Hello interval specifies the number of seconds between Hello packets (as a Keepalive function), and the Dead interval specifies the number of seconds after which a router is considered unavailable if a Hello is not received.
  - **Stub area flag** – A *Stub area* is an area that only requires a single point of egress, and therefore does not require a full list of external link advertisements. The stub area flag on two potential neighbors must be the same to avoid inappropriate link state exchanges. Another factor that affects neighboring is the kind of network. OSPF recognizes three network types:
    - **Broadcast** – For example, Ethernet. In broadcast networks, neighboring can be established with all other routers in the broadcast domain.
    - **Point to Point** – For example, serial links. In point to point (or point to multipoint) networks, neighboring can be established with the router at the other end of the link.
    - **NBMA** (non-broadcast multiple access) – For example, frame relay. In NBMA networks, neighbors must be explicitly declared.
- **Link State Database** – The Link State Database is composed of the LSA's sent and received by *neighboring* OSPF routers that have created *adjacencies* within an *area*. The database, after complete, contains all the link state information for a given area, at which time the Shortest Path First (SPF) algorithm is applied to determine the optimal route to all connected networks based on cost. The SPF algorithm employs the Dijkstra pathfinding algorithm, which essentially regards all routers as vertices in a graph, and computes the cost between each vertex.
- **Adjacencies** – OSPF routers exchange LSA's with adjacent routers to create the LSDB. Adjacencies are created in different fashions depending on the network type (see *Neighbors* above). Generally, the network type is broadcast (for example, Ethernet) so adjacencies are formed by the exchanging OSPF packets in a handshake-like fashion (see OSPF Packet types below). To minimize the amount of information exchanged between adjacent routers, segments (broadcast domains) with multiple OSPF routers elect a Designated Router (DR) and a Backup Designated Router (BDR) using Hello packets.

- **DR** (Designated Router) – On multi-access segments, OSPF routers elect a DR and a BDR, and all other routers on the segment create adjacencies with the DR and the BDR. DR election is based on a router's OSPF Priority, which is a configurable value from 0 (not eligible for DR) to 255. The router with the highest priority becomes the DR. In the event of a priority tie, the router with the highest Router ID (based on interface addressing) wins. When a router is the DR, its role is uncontested until it becomes unavailable.

  LSA's are then exchanged within LSUs across these adjacencies rather than between each possible pairing combination of routers on the segment; see Routing adjacencies: Designated Router (DR). Link state updates are sent by non-DR routers to the multicast address `225.0.0.6`, the RFC1583 assigned 'OSPFIGP Designated Routers' address. They are also flooded by DR routers to the multicast address `225.0.0.5` 'OSPFIGP All Routers' for all routers to receives the LSA's.

Routing adjacencies: Designated Router (DR)



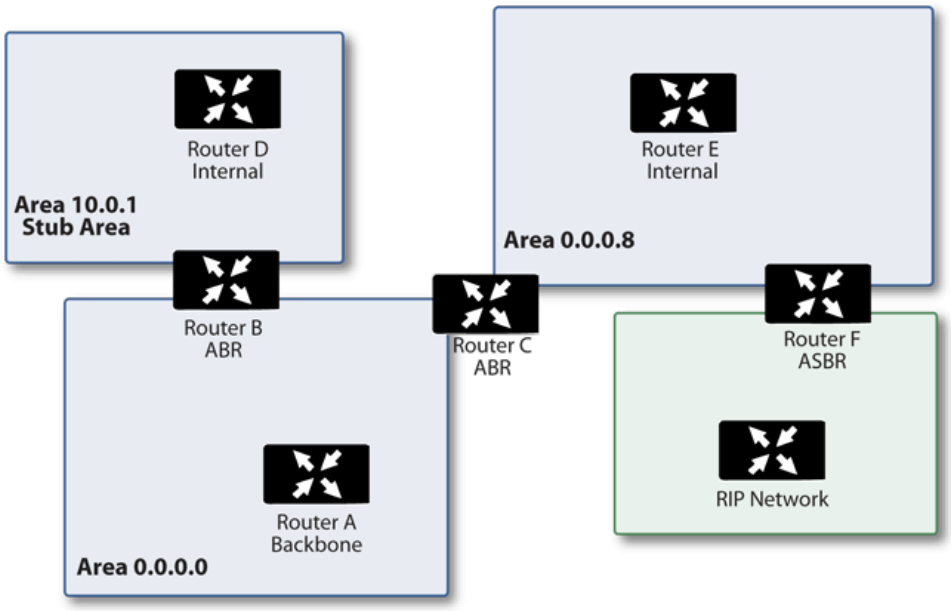
- **OSPF Packet types** – The five types of OSPF packets are:
  - **Hello** (OSPF type 1) – Sent at a certain interval to establish and maintain relationships with neighboring OSPF routers, and elect Designated Routers. (*Sent during the initialization and the 2-WAY phases on LSDB synchronization*).
  - **Database Description** (OSPF type 2) – Sent between OSPF routers during the creation of an adjacency. *During the Exstart phase of LSDB synchronization*, DD packets establish an ISN (initial sequence number) used to track LSA's, and they establish a master/slave relationship between neighboring OSPF routers. *In the Exchange phase of LSDB synchronization*, they contain short versions of Link State Advertisements. Because DD exchanges can span multiple packets, they are exchanged in a poll (master) and response (slave) fashion to ensure completeness.
  - **Link State Request** (OSPF type 3) – *During the Loading phase of LSDB synchronization*, LSR packets are sent to request database updates from a neighbor. This is the final step in the establishment of an adjacency.
  - **Link State Update** (OSPF type 4) – Sent in response to Link State Requests, LSU packets flood adjacencies with Link State Advertisements to achieve LSDB synchronization.
  - **Link State Acknowledgement** (OSPF type 5) – To ensure reliability of LSA flooding, all updates are acknowledged.
- **Link State Advertisements** (**LSA**) – There are 7 types of LSA's:
  - Type 1 (Router Link Advertisements) - Sent by an OSPF router to describe the links to each area to which it belongs. Type 1 LSA's are only flooded into a router's area.
  - **Type 2** (Network Links Advertisements) – Sent by the DR for an area describing the set of routers within the network. Type 2 LSA's are only flooded into a router's area.

- **Type 3** (Summary Link Advertisements) – Sent across areas by ABRs (Area Border Routers) to describe the networks within an area. Type 3 LSA's are also used for route aggregation purposes, and are not sent to Totally Stubby Areas.
- **Type 4** (AS Summary Link Advertisements) – Sent across areas by ABRs to describe networks within a different AS. Type 4 LSA's are not sent to Stub Areas.
- **Type 5** (AS External Link Advertisements) – Sent by ASBR (Autonomous System Boundary Routers) to describe routes to networks in a different AS. Type 5 LSA's are net sent to Stub Areas. There are two types of External Link Advertisements:
  - **External Type 1** - Type 1 packets add the internal link cost to the external link cost when calculating a link's metric. A Type 1 route is always preferred over a Type 2 route to the same destination.
  - **External Type 2** - Type 2 packets only use the external link cost to determine the metric. Type 2 is generally used when there is only one path to an external AS.
- **Type 6** (Multicast OSPF or MOSPF) - Called source/destination routing, this is in contrast to most unicast datagram forwarding algorithms (like OSPF) that route based solely on destination. For more information about MOSPF, see *RFC1584 – Multicast Extensions to OSPF*.
- **Type 7** (NSSA AS External Link Advertisements) – Sent by ASBRs that are part of an NSSA (see 'Stub Area').

- **Stub Area** – A stub area is an area that only requires one path, rather than an optimal path. This can be an area with only a single point of egress, or it can be an area where SPF optimization is not necessary. All routers in a stub area must be configured as stub routers, and rather than receiving the full state database, and computing the SPF tree, they receive only summary link information.

  There are different type of stub area:
  - **Stub area** – The standard stub area receives all LSA's except for LSA type 5 (AS External Link advertisement). This helps to keep the LSDB smaller, and reduces the computational overhead on the router.
  - **Totally Stubby Area** – A special type of stub area into which LSA types 3 (Summary Links), 4 (AS Summary Links) and 5 are not passed. Only intra-area routes, and a default route are advertised into totally stubby areas.
  - **NSSA** (Not So Stubby Area) – Described by RFC3101, NSSA is a hybrid stub area that allows external routes to be flooded within the NSSA area using type 7 LSA's (NSSA AS External Routes), but does not accept type 5 LSA's from other areas. NSSAs are useful when connecting a remote site running a different IGP (such as RIP) to an OSPF site, where the remote site's routes do not need to be distributed back to the main OSPF site. An NSSA *ABR* (Area Border Router) also has the ability to translate type 7 to type 5 LSA's (this is possible only from the SonicOS CLI).
- **Router Types** – OSPF recognizes 4 types of routers, based on their roles; see OSPF-recognized router types example.

OSPF-recognized router types example



- **IR** (Internal Router) - A router whose interfaces are all contained within the same area. An internal router's LSDB only contains information about its own area.
- **ABR** (Area Border Router) – A router with interfaces in multiple areas. An ABR maintains LSDBs for each area to which it is connected, one of which is typically the backbone.
- **Backbone Router** – A router with an interface connected to area 0, the backbone.
- **ASBR** (Autonomous System Boundary Router) – A router with an interface connected to a non-OSPF AS (such as a RIP network) which advertises external routing information from that AS into the OSPF AS.

Drop Tunnel Interface

A drop tunnel interface prevents traffic from being sent out using an incorrect route when the configured route is down. Traffic sent to a drop tunnel interface does not leave the Security Appliance, but is ostensibly dropped.

A drop tunnel interface should be used in conjunction with a VPN tunnel interface, although a drop tunnel interface can be used standalone. If a static route is bound to a tunnel interface, SonicWall recommends configuring a static route bound to a drop tunnel interface for the same network traffic. That way, if the tunnel interface goes down, the second static route is used and the traffic is effectively dropped. This prevents the data from being forwarded in the clear over another route.

When configuring a route over a VPN tunnel interface, if the tunnel is temporarily down, the corresponding route entry is disabled as well. SonicOS looks up a new route entry for the connections destined for the VPN protected network. In deployments that do not have a backup link for a remote VPN network, no other correct route entry is available. Traffic is sent to a wrong route entry, generally the default route, which causes security issues such as internal data sent without encryption.

For deployments without a backup link, consider configuring the route table as in this example:

```
route n:   local VPN network(source), remote VPN network(destination), VPN TI(egress_if)
route n+1: local VPN network(source), remote VPN network(destination), Drop If(egress_if)
```

When the VPN tunnel interface configured as in this example, the traffic matches the drop interface and is not sent out. When the VPN tunnel interface resumes, traffic resumes also.

App-Based Routing

App-Based Routing is a kind of PBF (policy-based forwarding) rule that allows traffic to take an alternative path from the next hop specified in the route table and is typically used to specify an egress interface for security or performance reasons.

When an App-Based Route entry is created, at the beginning the Security Appliance does not have enough information to identify the application and, therefore, cannot enforce the route entry. As more packets arrive, the Security Appliance determines the application and creates an internal entry in the App-ID cache, which is retained for the session. When a new session is created with the same destination IP address, destination port, and protocol ID, the Security Appliance could identify the application as the same from the initial session and apply the App-Based Route. Therefore, a session that is not an exact match and is not the same application, cannot be forwarded based on the App-Based Route.

This feature is available only when Gateway AV/Anti-Spyware/Intrusion Prevention/App Control/App Visualization is licensed and App Control is enabled in **MANAGE | Policies > Rules > App Control**.

Network > Routing

If you have routers on your interfaces, you configure static routes on the SonicWall Security Appliance on the **MANAGE | System Setup > Network > Routing** page. You can create static routing policies that create static routing entries that make decisions based upon source address, source netmask, destination address, destination netmask, service, interface, gateway and metric. This feature allows for full control of forwarding based upon a large number of user-defined variables.

Topics:

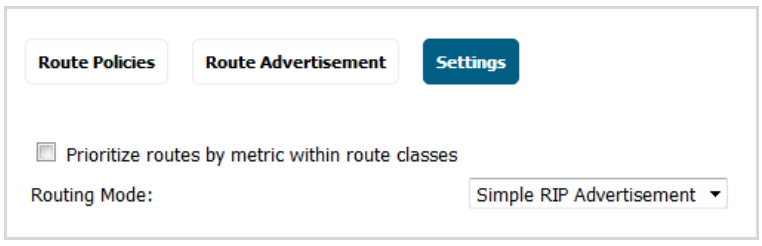
- **Network > Routing > Settings** on page 464

Network > Routing > Settings

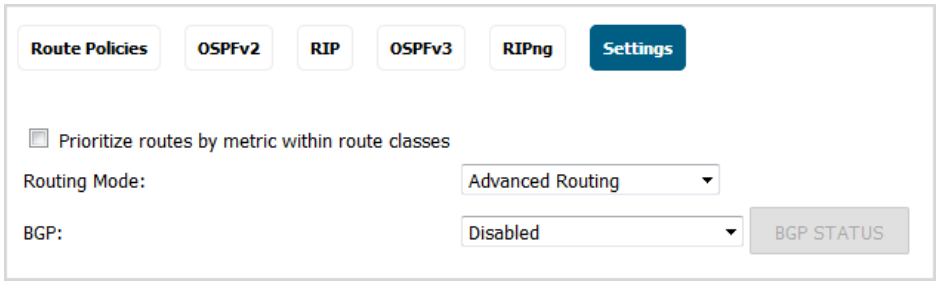
The look of **MANAGE | System Setup > Network > Routing > Settings** changes depending on the routing mode you select:

- **Simple RIP Advertisement**
- **Advanced Routing**

Simple RIP Advertisement



Advanced Routing



Network > Routing > Route Policies

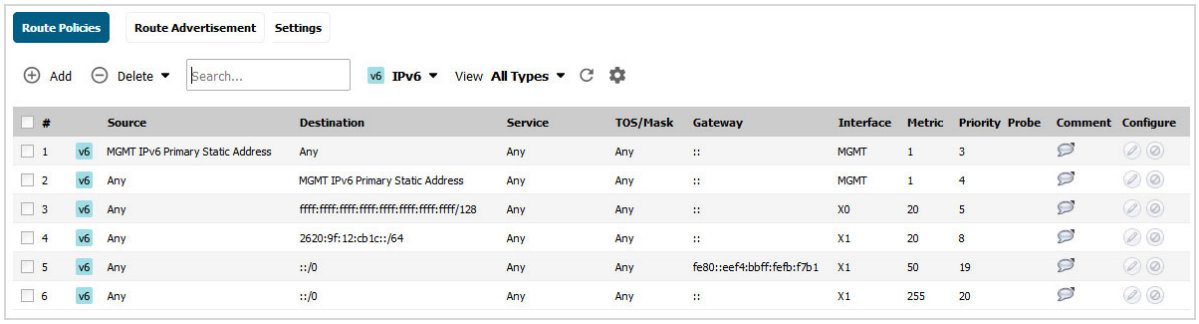
Network > Routing > Route Policies displays all the default and/or custom routes for either IPv4 or IPv6. The display is the same for either IP version except the IPv6 display shows the IPv6 link-local address instead of the IP address.

You can change the view of the route policies in the **Route Policies** table by selecting:

- **IPv4** or **IPv6**
- One of the view settings in **View**:

All Types	All routing policies including **Custom Policies** and **Default Policies**. Initially, only **Default Policies** are displayed in the **Route Policies** table when you select **All Types**.
Custom Polices	Ones you created.
Default Policies	Ones created by SonicOS.

You can filter the display by entering the source, destination, or interface in the **Search** field.

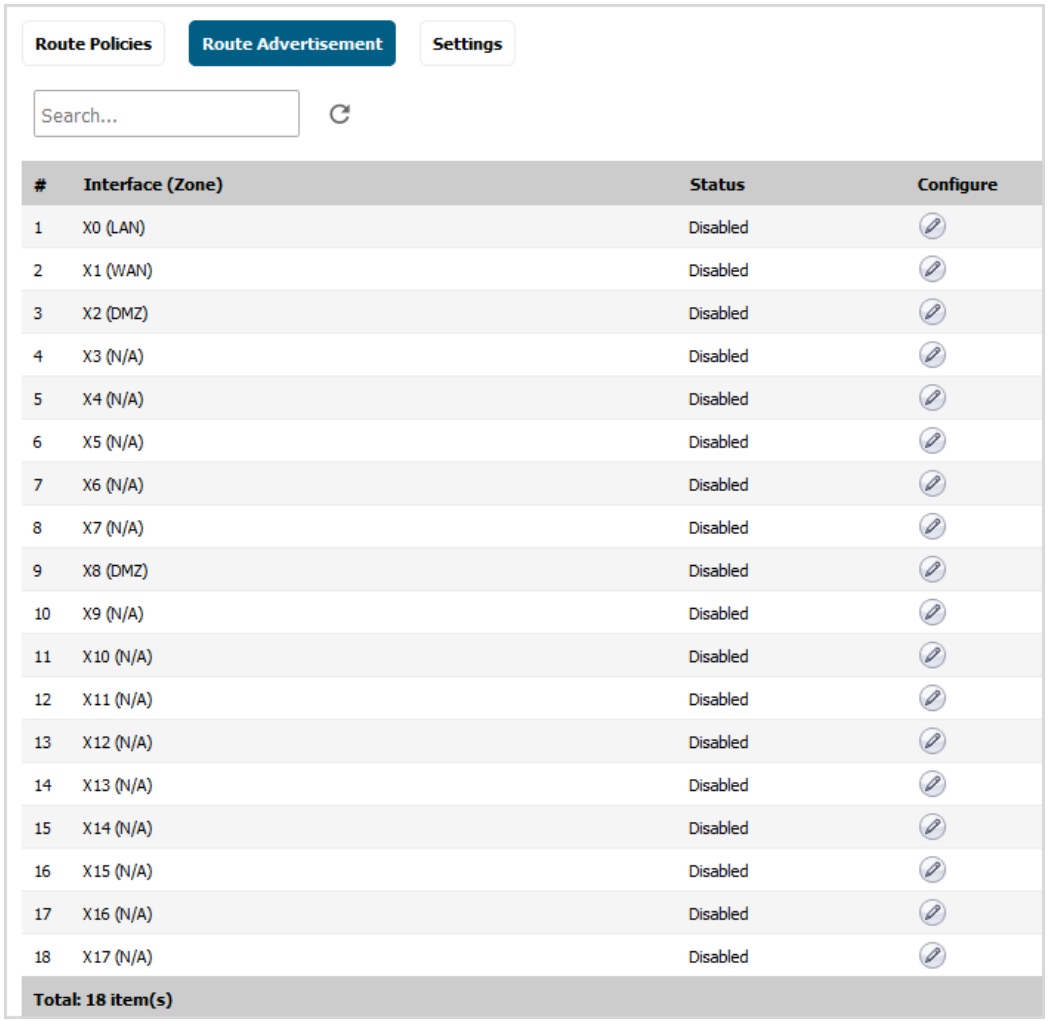


Column	Route policy configuration
Source	IP version icon and name for the source.
Destination	Destination IP address (IPv4) or MAC address (IPv6).
Service	Service object configured for the route policy.
TOS/Mask	TOS and TOS Mask configured for the route.
Gateway	Gateway IP address (IPv4) or MAC address (IPv6).
Interface	Interface configured for the route policy.
Metric	Metric configured for route priority.
Priority	Priority of the route policy.
Probe	Whether probe is configured.

Column	Route policy configuration
Comment	**Comment** icon containing the comment entered when the custom route was configured; `Auto-added Route Policy` for default policies.
Configure	**Edit** and **Delete** icons; icons for default policies are dimmed.

Network > Routing > Route Advertisement

Network > Routing > Route Advertisement displays only when **Simple RIP Advertisement** is chosen for **Routing Mode**.



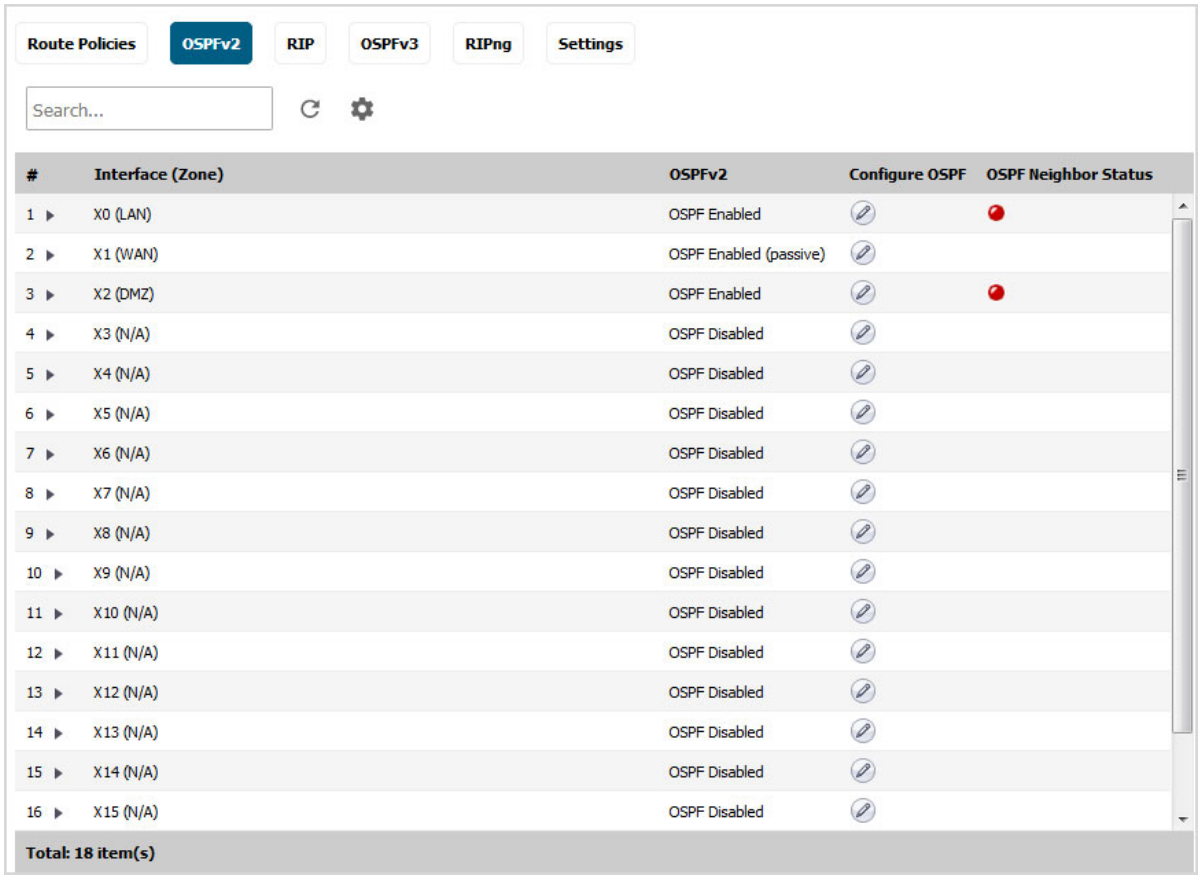
#	Interface (Zone)	Status	Configure
1	X0 (LAN)	Disabled	
2	X1 (WAN)	Disabled	
3	X2 (DMZ)	Disabled	
4	X3 (N/A)	Disabled	
5	X4 (N/A)	Disabled	
6	X5 (N/A)	Disabled	
7	X6 (N/A)	Disabled	
8	X7 (N/A)	Disabled	
9	X8 (DMZ)	Disabled	
10	X9 (N/A)	Disabled	
11	X10 (N/A)	Disabled	
12	X11 (N/A)	Disabled	
13	X12 (N/A)	Disabled	
14	X13 (N/A)	Disabled	
15	X14 (N/A)	Disabled	
16	X15 (N/A)	Disabled	
17	X16 (N/A)	Disabled	
18	X17 (N/A)	Disabled	

Total: 18 item(s)

Interface (Zone)	Interfaces configured for route advertisement. If a zone has not been configured for an interface, the **(Zone)** designation is **(N/A)**.
Status	Either **Enabled** or **Disabled**.
Configure	Contains the **Edit** icon.

Network > Routing > OSPFv2

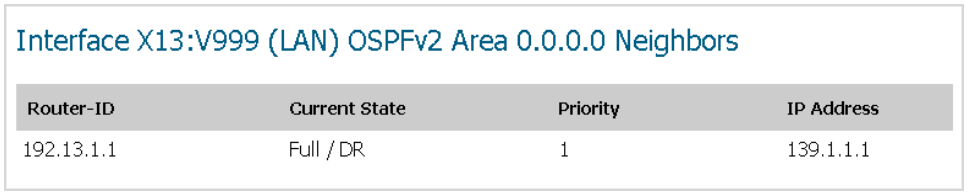
Network > Routing > OSPFv2, which displays only when **Advanced Routing** is chosen for **Routing Mode**, shows the status of OSPFv2 and allows you to configure OSPFv2 for an interface.



Settings	Icon that displays the **Settings** popup for configuring the metrics for default routes.
Interface (Zone)	Interfaces and their zone configured for OSPFv2. If a zone has not been configured for an interface, the **(Zone)** designation is **(N/A)**.
OSPFv2	Indicates whether OSPF is enabled on an interface: • **OSPF Enabled** • **OSPF Enabled (passive)** • **OSPF Disabled**
Configure OSPF	Displays the **Edit** icon for the interface.
OSPF Neighbor Status	Displays the **Status** icon, which indicates whether there are active or inactive neighbors; clicking the icon displays the **Interface OSPFv2 Neighbors** popup for more detail about the interface's neighbors. See **Network > Routing > OSPFv2 > Interface OSPFv2 Neighbors** on page 468.

Network > Routing > OSPFv2 > Interface OSPFv2 Neighbors

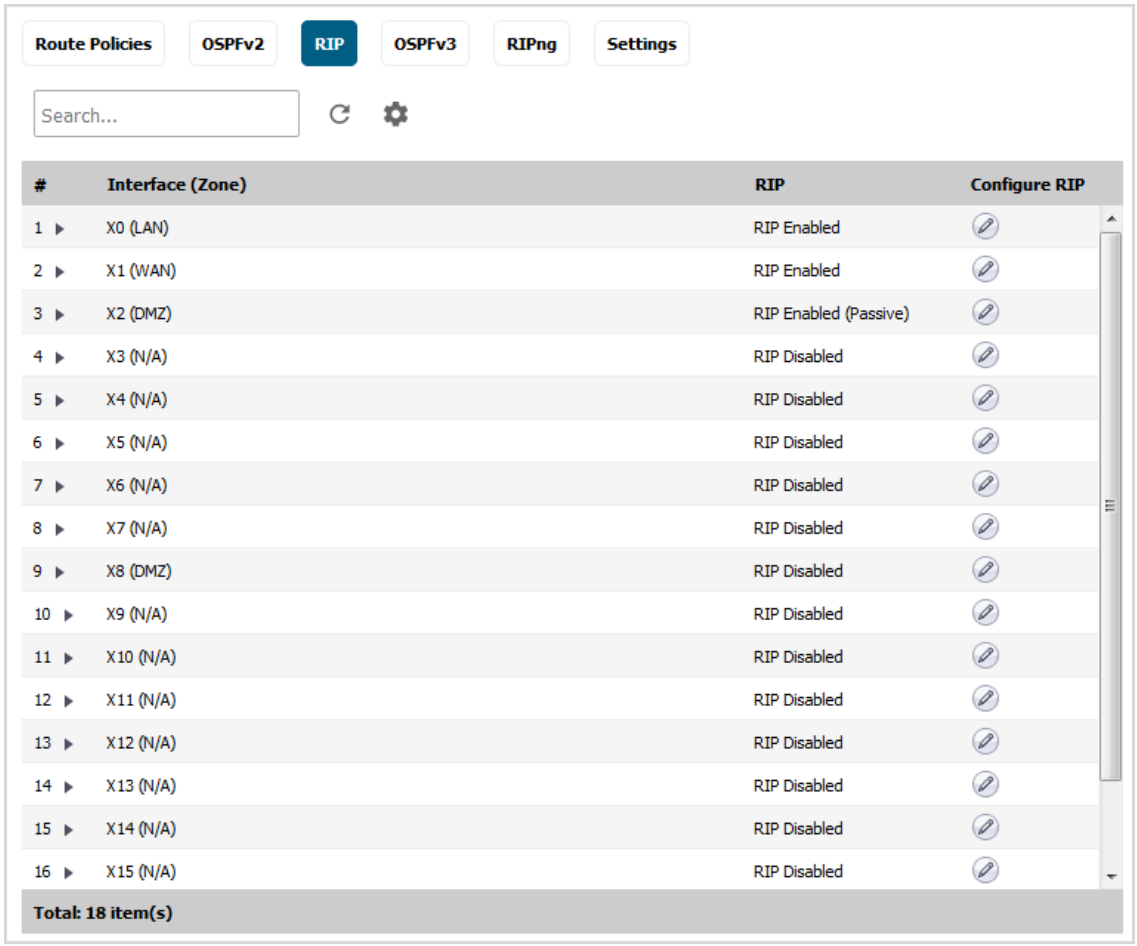
Display this popup by clicking the **Status** icon for the interface.



Router ID	Neighbor's router ID.
Current State	State of the OSPFv2 neighborhood when it is established: • **Init** • **2-way** • **ExStart** • **Exchange** • **Loading** • **Full**
Priority	Neighbor's router's priority.
IP Address	IP address of neighbor's router.

Network > Routing > RIP

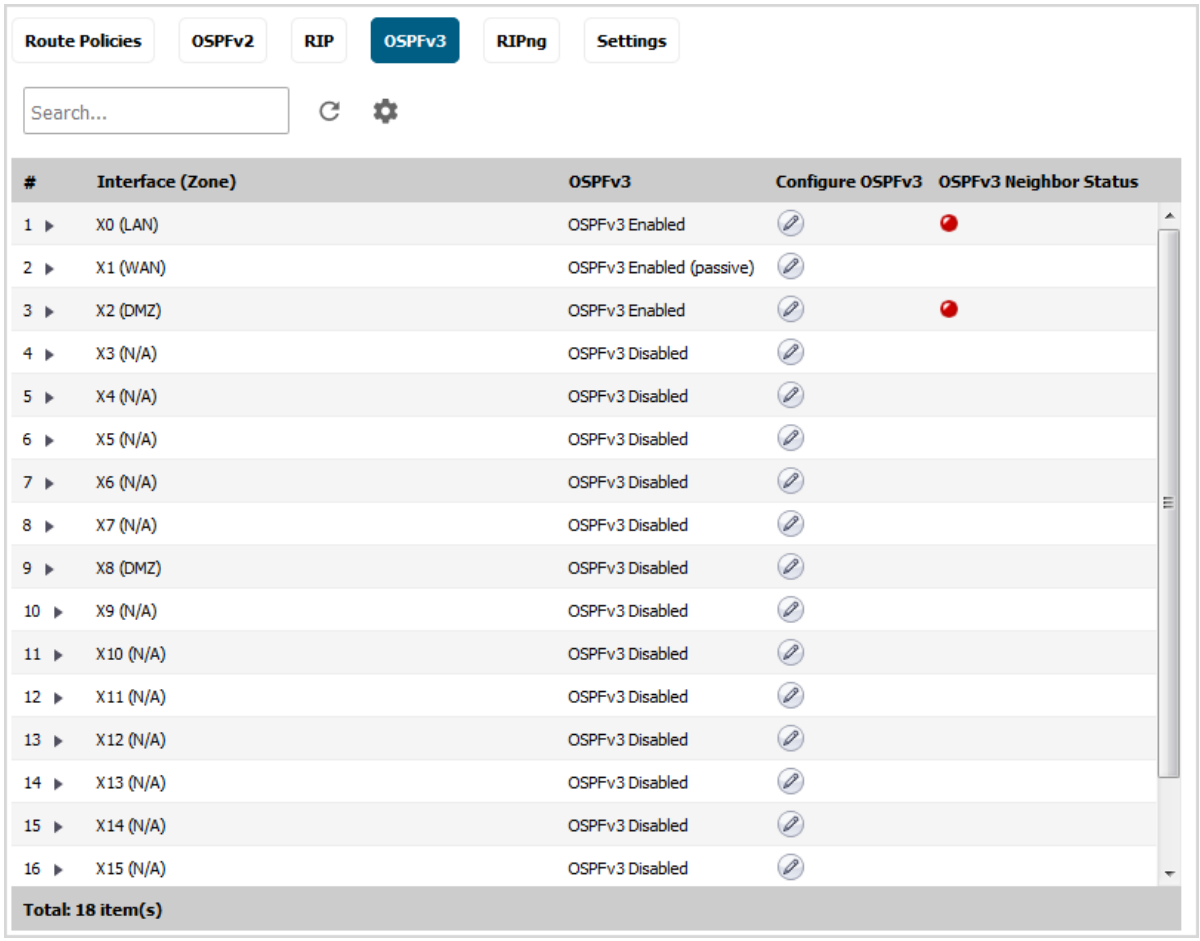
Network > Routing > RIP, which displays only when **Advanced Routing** is chosen for **Routing Mode**, shows the status of RIP and allows you to configure RIP for an interface.



Settings	Icon that displays the **Settings** popup for configuring the metrics for default routes.
Interface (Zone)	Interfaces and their zone configured for RIP. If a zone has not been configured for an interface, the **(Zone)** designation is **(N/A)**.
RIP	Indicates whether RIP is enabled on an interface: • **RIP Enabled** • **RIP Enabled (passive)** • **RIP Disabled**
Configure RIP	Displays the **Edit** icon for the interface.

Network > Routing > OSPFv3

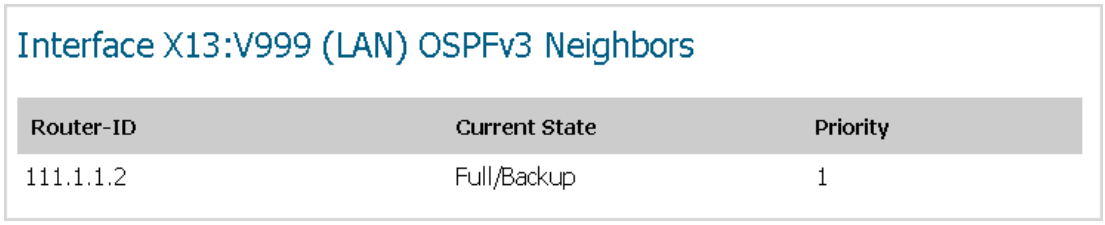
Network > Routing > OSPFv3, which displays only when **Advanced Routing** is chosen for **Routing Mode**, shows the status of OSPFv3 and allows you to configure OSPFv3 for an interface.



Settings	Icon that displays the **Settings** popup for configuring the metrics for default routes.
Interface (Zone)	Interfaces and their zone configured for OSPFv3. If a zone has not been configured for an interface, the **(Zone)** designation is **(N/A)**.
OSPFv3	Indicates whether OSPF is enabled on an interface: • **OSPFv3 Enabled** • **OSPFv3 Enabled (passive)** • **OSPFv3 Disabled**
Configure OSPFv3	Displays the **Edit** icon for the interface.
OSPFv3 Neighbor Status	Displays the **Status** icon, which indicates whether there are active or inactive neighbors; clicking the icon displays the **Interface OSPFv3 Neighbors** popup for more detail about the interface's neighbors. See **Network > Routing > OSPFv3 > Interface OSPFv3 Neighbors** on page 471.

Network > Routing > OSPFv3 > Interface OSPFv3 Neighbors

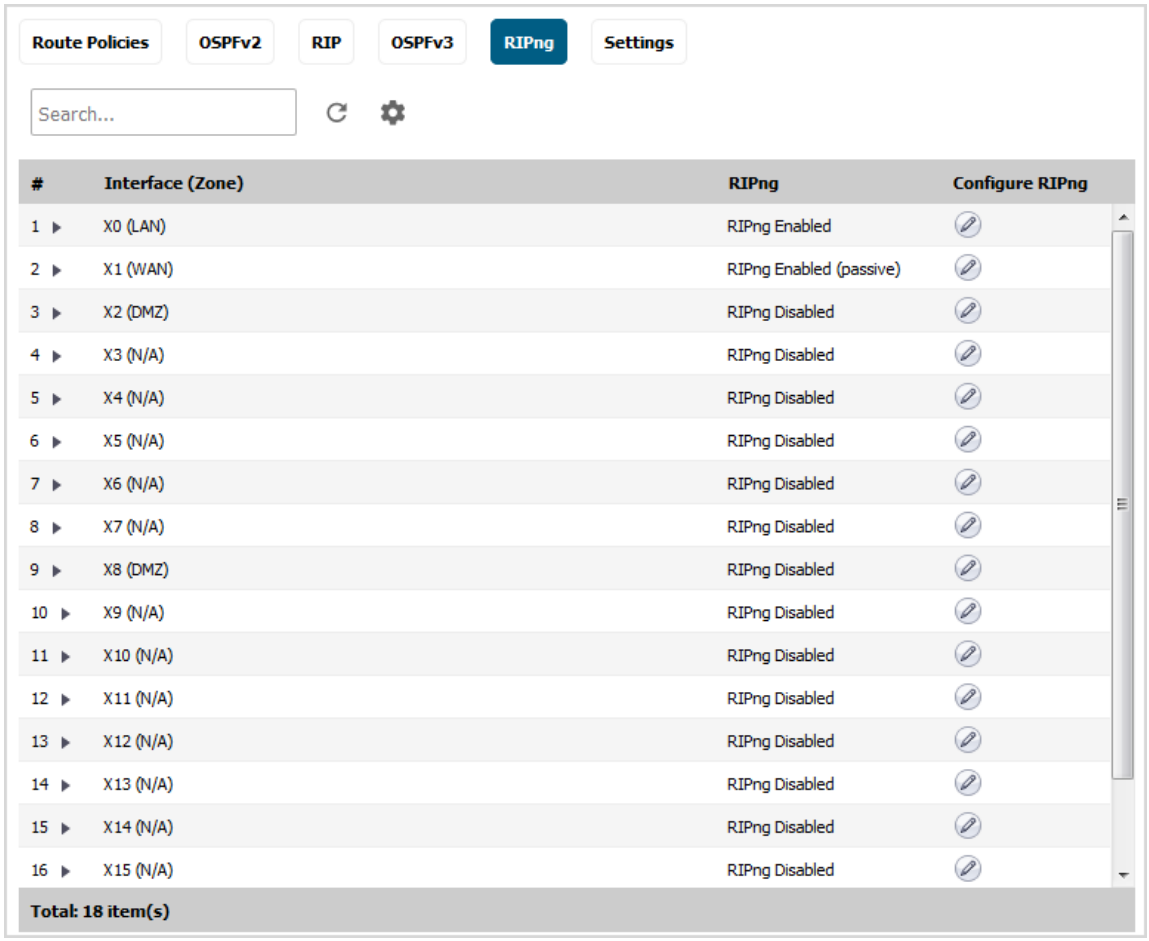
Display this popup by clicking the **Status** icon for the interface.



Router ID	Neighbor's router' ID.
Current State	State of the OSPFv3 neighborhood when it is established: • **Init** • **2-way** • **ExStart** • **Exchange** • **Loading** • **Full**
Priority	Neighbor's router's priority.

Network > Routing > RIPng

Network > Routing > RIPng, which displays only when **Advanced Routing** is chosen for **Routing Mode**, shows the status of RIPng and allows you to configure RIPng for an interface.



Settings	Icon that displays the **Settings** popup for configuring the metrics for default routes.
Interface (Zone)	Interfaces and their zone configured for RIPng. If a zone has not been configured for an interface, the **(Zone)** designation is **(N/A)**.
RIPng	Indicates whether RIPng is enabled on an interface: • **RIP Enabled** • **RIP Enabled (passive)** • **RIP Disabled**
Configure RIPng	Displays the **Edit** icon for the interface.

Configuring Routing

Topics:

- Prioritizing Routes by Metric on page 473
- Configuring Metrics for Default Routes Learned through Router Advertisement on page 473

- Configuring Static and Policy-based Routes on page 475
- Configuring a Static Route for a Drop Tunnel Interface on page 479
- Configuring OSPF and RIP Advanced Routing Services on page 481
- Configuring BGP Advanced Routing on page 490

Prioritizing Routes by Metric

IMPORTANT: Changing to metric-weighted route prioritization requires restarting the SonicWall Security Appliance.

The Metric-Weighted option allows the metric to take precedence in prioritization over route specificity. The precedence (high to low) used during prioritization when the metric option is:

- Not selected (default):
    a Route class (determined by the combination of source, destination, service, and TOS fields with values other than **ANY**.
    b The cumulative specificity of the source, destination, service, and TOS fields.
    c The metric.
- Selected:
    a Route class.
    b The metric.
    c The cumulative specificity of the source, destination, service, and TOS fields

To change to metric-weighted route prioritization:

1 Navigate to **MANAGE | System Setup > Network > Routing > Settings**.
2 Select **Prioritize routes by metric within route classes**. A confirmation message displays.



3 Click **OK**.
4 Navigate to **MANAGE | Updates > Restart** to manually restart SonicOS.

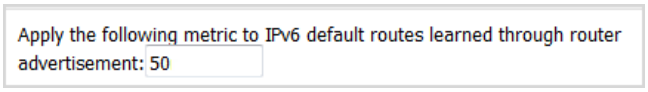
Configuring Metrics for Default Routes Learned through Router Advertisement

NOTE: This setting takes effect only on IPv6 default routes learned from Router Advertisement.

To configure metrics for default routes learned through router advertisement:

1 Navigate to **MANAGE | Network > Routing**.
2 Click **Route Policies**.

3 Click the **Settings** icon. The **Settings** dialog displays.



4 This route metric applies to default routes learned through router advertisement. Enter the metric in the **Apply the following metric to IPv6 default routes learned through router advertisement** field. The minimum is 1, the maximum is 255, and the default is **50**.

TIP: Lower metrics are considered better and take precedence over higher ones.

5 Click **ACCEPT**.

Configuring Route Advertisement

To enable Route Advertisement for a network interface:

1 Navigate to **MANAGE | Network > Routing**.

2 Click **Route Advertisement**.

3 Click the **Edit** icon in the **Configure** column for the interface. The **Interface Route Advertisement Configuration** dialog displays.

4 Select one of the following types from the **RIP Advertisements** drop-down menu:

- **Disabled** (default) - Disables RIP advertisements.
- **RIPv1 Enabled** - RIPv1 is the first version of Routing Information Protocol.
- **RIPv2 Enabled (multicast)** - To send route advertisements using multicasting (a single data packet to specific notes on the network).
- **RIPv2 Enabled (broadcast)** - To send route advertisements using broadcasting (a single data packet to all nodes on the network).

By selecting a type other than **Disable**, the other options become available.

5 From the **Advertise Default Route drop-down** menu, select:

- **Never** (default)
- **When WAN is up** (not available for a WAN interface)
- **Always**

6 Enable **Advertise Static Routes** if you have static routes configured on the Security Appliance, enable this feature to exclude them from Route Advertisement.

7 Enable **Advertise Remote VPN Networks** if you want to advertise VPN networks.

8 Enter a value in seconds between advertisements broadcast over a network in the **Route Change Damp Time (seconds)** field. The default value is **30** seconds, the minimum is 1 second, and the maximum is 99 seconds. A lower value corresponds with a higher volume of broadcast traffic over the network. The **Route Change Damp Time (seconds)** setting defines the delay between the time a VPN tunnel changes state (up or down) and the time the change is advertised with RIP. The delay, in seconds, prevents ambiguous route advertisements sent as a result of a temporary change in the VPN tunnel status.

9 Enter the number of advertisements that a deleted route broadcasts until it stops in the **Deleted Route Advertisements (0-99)** field. The default value is **1**.

10 Enter a value from **1** (default) to 15 in the **Route Metric (1-15)** field. This is the number of times a packet touches a router from the source IP address to the destination IP address.

> **NOTE:** The following options are available only if a RIPv2 advertisement option is selected in the **RIP Advertisements** drop-down menu. If you selected **RIPv1 Enabled**, go to Step 13.

11 You can enter a value for the route tag in the **RIPv2 Route Tag (4 HEX Digits)** field. This value is implementation-dependent and provides a mechanism for routers to classify the originators of RIPv2 advertisements. The default value is **0**.

12 If you want to enable RIPv2 authentication, select one of the following options from the **RIPv2 Authentication** drop-down menu (the default is **Disabled**):

- **User defined** - Two fields display:
  - **Authentication Type (4 Hex Digits)** – Enter 4 hex digits in the field. The default is **0**.
  - **Authentication Data (32 Hex Digits)** – Enter 32 hex digits in the field.
- **Cleartext Password** - The **Authentication Password** field displays. Enter a password of up to 16 characters in the field.
- **MD5 Digest** - Enter a numerical value from 0-255 in the Authentication **Key-Id (0-255)** field. Enter a 32 hex digit value for the **Authentication Key (32 hex digits)** field, or use the generated key.
  - **Authentication Key-Id (0-255)** – Enter up to 255 characters in the field. The default is **1**.
  - **Authentication Key** – Enter up to 32 characters in the field.

13 Click **OK**.

Configuring Static and Policy-based Routes

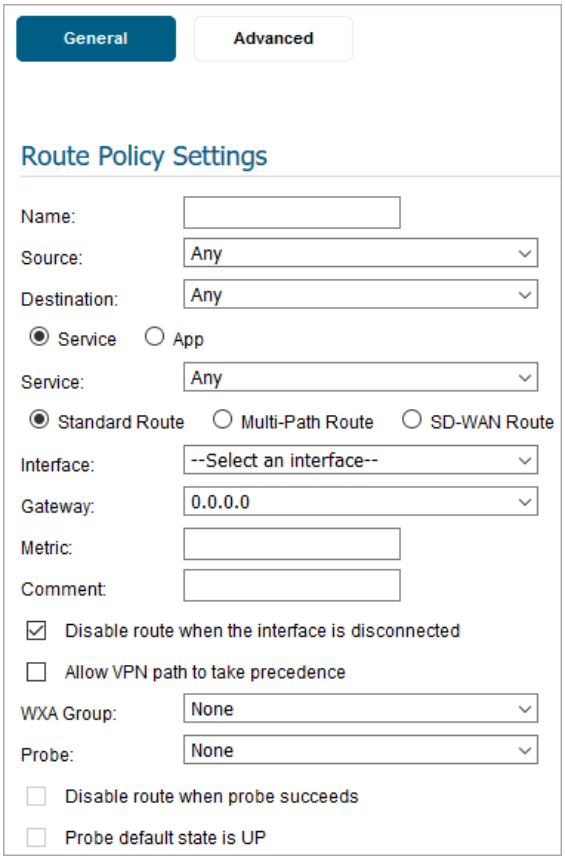
In SonicOS, a static route is configured through a basic route policy. For the maximum number of routes per Security Appliance, see the description of configuring route policies in *SonicOS 6.5 Policies*.

When configuring a static route, you can optionally configure a Network Monitor policy for the route. When a Network Monitor policy is used, the static route is dynamically disabled or enabled, based on the state of the probe for the policy.

To configure a static or policy-based route:

1 Navigate to **MANAGE | System Setup > Network > Routing > Route Policies**.

2 Click the **Add** icon. The **Add Route Policy** dialog displays.



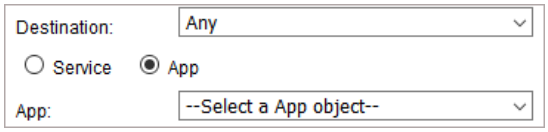
3 Enter a meaningful name in the **Name** field.

4 From **Source**, select the source address object for the static route or select **Create new address object** to dynamically create a new address object. The default is **Any**.

5 From **Destination**, select the destination address object or select **Create new address object** to dynamically create a new address object. The default is **Any**.

6 Choose the type of route policy:

- **Service** (default)
- **App** – Service changes to **App**:



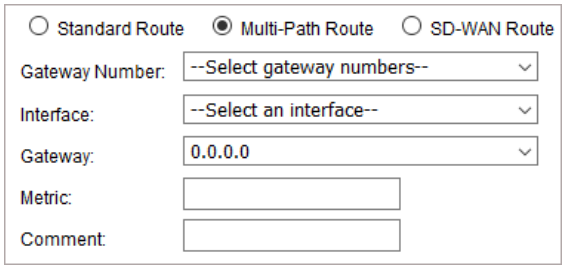
7 From **Service**, select a service object. For a generic static route that allows all traffic types, simply select **Any** (the default).

8 Go to Step 14.

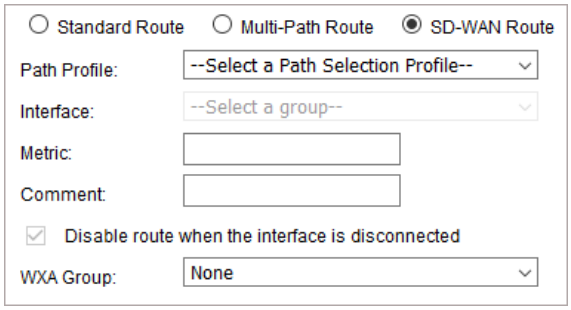
9 From **App**, select an App Object.

10 Choose the type of route to use:

- **Standard Route** (the default) – Go to Step 14.
- **Multi-Path Route** – The **Gateway Number** option displays:



- **SD-WAN Route** – The options change:



> **NOTE:** The **Interface** and **Disable route when the interface is disconnected** options are dimmed because these options cannot be edited in SD-WAN policies. The **Interface** option is populated with the SD-WAN group name in the associated Path Selection Profile (PSP) and cannot be changed. The interface for the SD-WAN route is selected from the SD-WAN group that is part of the PSP associated with the SD-WAN route and, therefore, cannot be configured.

Go to Step 13.

11 From **Gateway Number**, select the maximum number of gateways:

- **2**
- **3**
- **4**

12 Go to Step 16.

13 From **Path Profile**, select a Path Selection Profile.

14 From **Interface**, select the interface to be used for the route or select **Create VPN Tunnel interface** to dynamically create a new VPN policy. For information about creating a VPN policy, see *SonicOS 6.5 Connectivity*.

15 From **Gateway**, select the gateway address object to be used for the route or select **Create new address object** to dynamically create a new address object. The default is **0 . 0 . 0 . 0**. For information about creating address objects, see *SonicOS 6.5 Policies*.

16 Enter the **Metric** (weighted cost) for the route. The minimum is 1, and the maximum is 254. The default metric for:

- Static routes is **1**
- Dynamic routes learned from:
  - RIP/RIPng is **120**
  - OSPFv2/OSPFv3 is **110**
  - BGP is **20**

For more information on metrics, see About Metrics and Administrative Distance on page 451 and Policy-based Routing on page 453.

TIP: Lower metrics are considered better and take precedence over higher metrics (costs). SonicOS adheres to Cisco-defined metric values for directly connected interfaces, statically encoded routes, and all dynamic IP routing protocols.

17 Optionally, enter a **Comment** for the route. This field allows you to enter a descriptive comment for the new static route policy.

18 To have the route automatically disabled when the interface is disconnected, select **Disable route when the interface is disconnected**. This option is selected by default.

NOTE: If you are configuring a SD-WAN route, this option is both selected and dimmed; you cannot change it.

19 If you are configuring an SD-WAN route, go to Step 21.

20 Optionally, to create a backup route for a VPN tunnel, select **Allow VPN path to take precedence**. This option is not selected by default.

By default, a user-configured VPN tunnel static route has a metric of 1 and takes precedence over VPN traffic. The **Allow VPN path to take precedence** option gives precedence over the static route to VPN traffic to the same destination address object. This results in the following behavior when a VPN tunnel:

- **Is active**: Static routes matching the destination address object of the VPN tunnel are automatically disabled if the **Allow VPN path to take precedence** option is enabled. All traffic is routed over the VPN tunnel to the destination address object.
- **Goes down**: Static routes matching the destination address object of the VPN tunnel are automatically enabled. All traffic to the destination address object is routed over the static routes.

21 If WXA is licensed, select the WXA group from **WXA Group**. The default is **None**.

22 If you are configuring an SD-WAN route, go to Step 28.

23 To:

- Use probe-enabled, policy-based routing, go to Step 24.
- Ignore probe-enabled routing and configure TOS and administration distance values, go to Step 28.
- Apply the configuration, go to Step 32.

24 From **Probe**, select:

- **None** (default). Go to Step 27.
- A Network Monitor object; the following two options become available for configuring Probe-Enabled Policy-based Routing.
- **Create new Network Monitor object**. The **Add Policy** dialog displays. For how to create a Network Monitor object, see the procedure in *SonicOS 6.5 Investigate*.

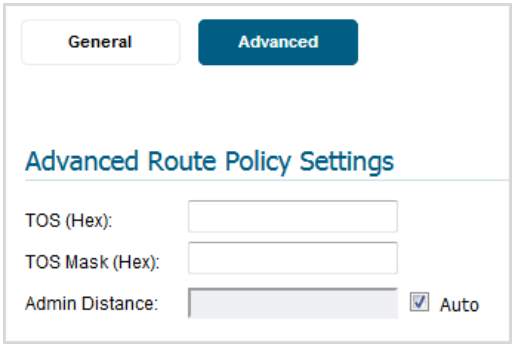
25 To have the route disabled when a probe succeeds, select **Disable route when probe succeeds**. This option is not selected by default.

IMPORTANT: Typical configurations do not check the **Disable route when probe succeeds** checkbox, because typically administrators want to disable a route when a probe to the route's destination fails. This option gives you added flexibility for defining routes and probes.

26 To have the route consider the probe to be successful (that is, in the UP state) when the attached Network Monitor policy is in the UNKNOWN state, select the **Probe default state is UP**. This is useful to control the probe-based behavior when a unit of a High Availability pair transitions from IDLE to ACTIVE, because this transition sets all Network Monitor policy states to UNKNOWN.

27 To use default TOS and admin distance values, go to Step 32.

28 Click **Advanced**.



29 Enter a TOS value in the **TOS (Hex)** field. The maximum value is FF. If the **TOS** and **TOS Mask** fields are not configured, a value of 0 is used. For further information about TOS and TOS Mask values, see Policy-based TOS Routing on page 453.

30 Enter the same value in the **TOS Mask (Hex)** field.

31 To manually specify an administration distance:

   a Deselect **Auto**. The **Admin Distance** field becomes available. This option is selected by default. For information about administration distance, see About Metrics and Administrative Distance on page 451.

   b Enter the administration distance in the **Admin Distance** field.

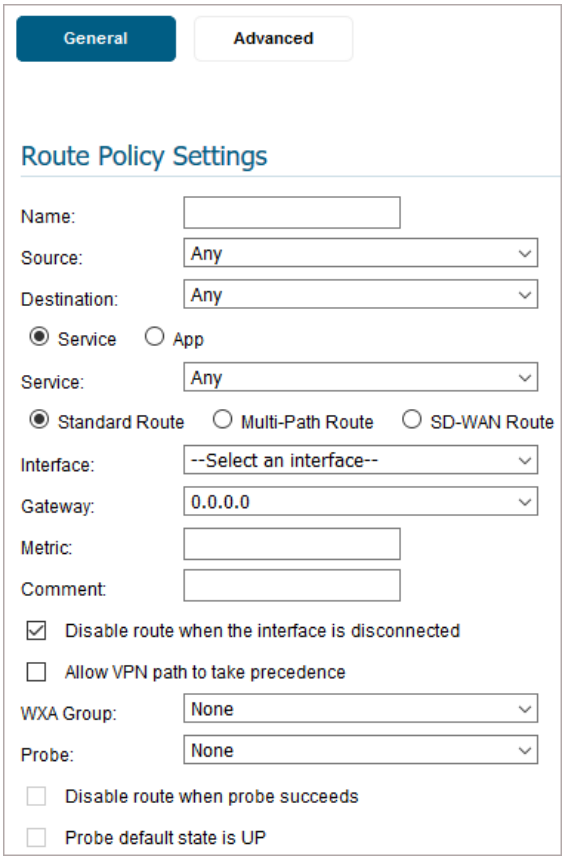
32 Click **OK**.

Configuring a Static Route for a Drop Tunnel Interface

To add a static route for a drop tunnel interface:

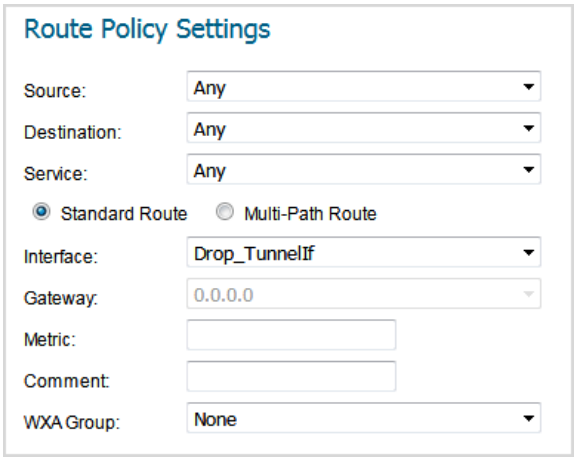
1 Navigate to **MANAGE | System Setup | Network > Routing > Route Policies**.

2 Click the **Add** icon. The **Add Route Policy** dialog displays.



3 Configure the values for **Source**, **Destination**, **Service** and **Route** options as described in Configuring Static and Policy-based Routes on page 475.

4 For **Interface**, select **Drop_TunnelIf**. The options change.



5 Finish configuring the options as in Configuring Static and Policy-based Routes on page 475.

6 Click **OK**. The route is enabled and displayed in the **Route Polices** table.

Configuring OSPF and RIP Advanced Routing Services

NOTE: ARS is a fully featured multi-protocol routing suite. The sheer number of configurable options and parameters provided is incongruous with the simplicity of a graphical user interface. Rather than limiting the functionality of ARS, an abbreviated representation of its capabilities has been rendered in the SonicOS Management Interface, providing control over the most germane routing features, while the full command suite is available through the CLI. The ARS CLI can be accessed from an authenticated CLI session and contains 3 modules:

- **route ars-nsm** – The Advanced Routing Services Network Services Module. This component provides control over core router functionality, such as interface bindings and redistributable routes.
- **route ars-rip** – The RIP module. Provides control over the RIP router.
- **route ars-ospf** – The OSPF module. Provides control over the OSPF router.

In general, all of the functionality needed to integrate the Security Appliance into most RIP and OSPF environments is available through the Web-based GUI. The additional capabilities of the CLI make more advanced configurations possible.

By default, Advanced Routing Services are disabled, and must be enabled to be made available.

The operation of the RIP and OSPF routing protocols is interface dependent. Each interface and virtual subinterface can have RIP and OSPF settings configured separately, and each interface can run both RIP and OSPF routers.

Topics:

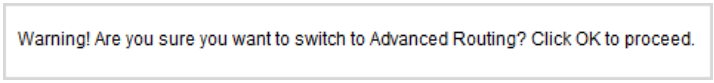
- Enabling Advanced Routing Services and BGP on page 481
- Configuring OSPF on page 482
- Configuring RIP and RIPng on page 486
- Configuring Advanced Routing for Tunnel Interfaces on page 489

Enabling Advanced Routing Services and BGP

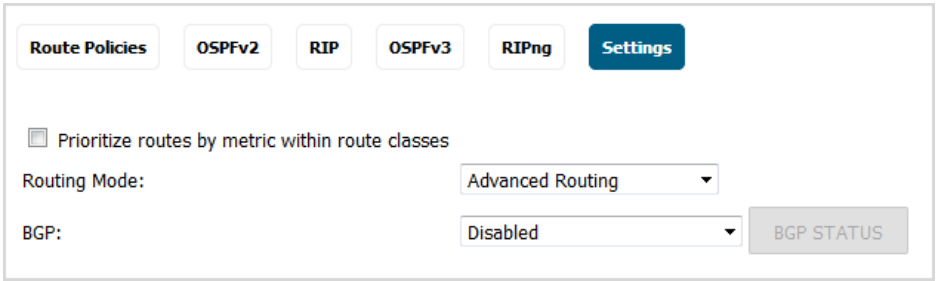
To enable advanced routing services:

1 Navigate to **MANAGE | System Setup | Network > Routing > Settings**.

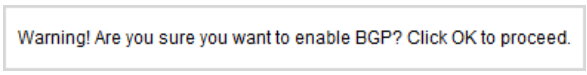
2 From **Routing mode**, select **Advanced Routing**. A confirmation message displays.



3 Click **OK**. The options on **Network > Routing** change:



4 To enable BGP, select **Enabled (Configure with CLI)** from **BGP**. The default is **Disabled**. A confirmation message displays.



5 Click **OK**. **BGP STATUS** becomes available.

Configuring OSPF

NOTE: OSPF design concepts are beyond the scope of this document. This section describes how to configure a SonicWall Security Appliance to integrate into an OSPF network, be it existing or newly implemented, but it does not offer design guidelines. For terms used throughout this section, refer to OSPF Terms on page 459.

Topics:

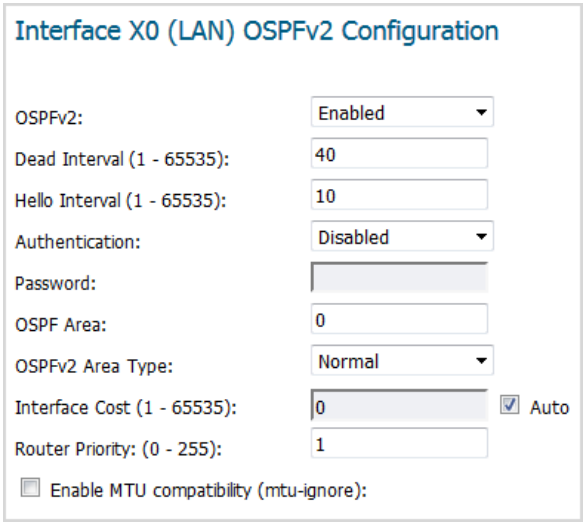
- Configuring OSPFv2 on page 482
- Configuring OSPFv3 on page 484

Configuring OSPFv2

To configure an interface for OSPFv2:

1 Navigate to **MANAGE | System Setup | Network > Routing > OSPFv2**.

2 Click the **Edit** icon for the interface. The **Interface OSPFv2 Configuration** dialog displays.



3 From **OSPFv2**, select:

Disabled (default)	OSPF Router is disabled on this interface. Go to Step 13.
Enabled	OSPF Router is enabled on this interface.
Passive	The OSPF router is enabled on this interface, but only advertises connected networks using type 1 LSAs (Router Link Advertisements) into the local area. All the options except **OSPF Area** are dimmed; go to Step 9.

4 To specify the period after which an entry in the LSDB is removed if Hello is not received, enter the time, in seconds, in the **Dead Interval (1 - 65535)** field. The default is **40** seconds, with a minimum of 1 and a maximum on 65,535.

IMPORTANT: Be sure this value agrees with the other OSPF routers on the segment for successful neighbor establishment.

5 To specify the period of time between Hello packets, enter the time, in seconds, in the **Hello Interval (1 - 65535)** field. The default is **10** seconds, with a minimum of 1 and a maximum on 65,535.

IMPORTANT: Be sure this value agrees with the other OSPF routers on the segment for successful neighbor establishment.

6 From **Authentication**, select the type of authentication used on this interface:

Disabled	No authentication is used; go to Step 8.
Simple Password	A plain-text password is used for identification purposes by the OSPF router.
Message Digest	An MD5 hash is used to securely identify the OSPF route.

IMPORTANT: Be sure this setting agrees with the other OSPF routers on the segment for successful neighbor establishment.

7 If you specified:

Simple Password	Enter a 1- to 15-character alphanumeric password.
Message Digest	Enter a 1- to 15-character alphanumeric password.

8 Enter the Area ID in the **OSPF Area** field. The OSPF Area can be represented in either IP or decimal notation. For example, the area connected to `X4:100` as either `100.100.100.100` or `1684300900`. The default is **0**.

9 Select the OSPFv2 area type from **OSPFv2 Area Type** (for a detailed description of these settings, see OSPF Terms on page 459):

Normal	Default; receives and sends all applicable LSA types.
Stub Area	Does not receive type 5 LSAs (AS External Link Advertisements).
Totally Stubby Area	Does not receive LSA types 3, 4, or 5.
Not So Stubby Area	Receives type 7 LSAs (NSSA AS External Routes).
Totally Stubby NSSA	Receives type 1 and 2 LSAs.

10 To:

- Specify the overhead of sending packets across this interface, enter the overhead in the **Interface Cost (1 - 65535)** field. The default value is **0**, generally used to indicate an Ethernet interface. The minimum and default value is 0 (for example, Fast Ethernet) and the maximum value is 65,535 (for example, pudding).
- Have the cost determined automatically, select **Auto**, which dims the **Interface Cost** field. This option is selected by default.

11 To specify the router priority value is used in determining the Designated Router (DR) for a segment, enter the value in the **Router Priority (0-255)** field. The higher the value, the higher the priority. For a priority tie, the Router ID acts as the tie-breaker. Setting a value of 0 makes the OSPF router on this interface ineligible for DR status. The default value is **1**, and the maximum value is 255.

12 To enable MTU compatibility, select **Enable MTU compatibility (mtu-ignore)**. This option is not selected by default.

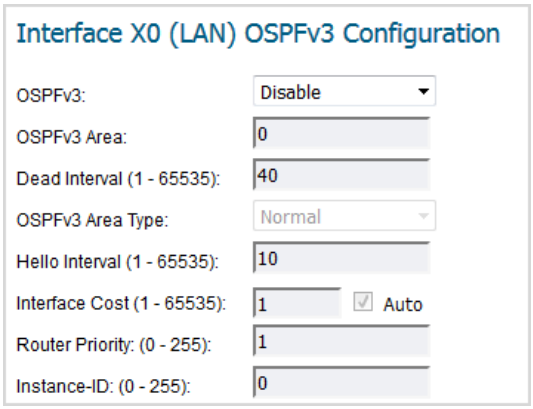
13 Click **OK**.

Configuring OSPFv3

To configure an interface for OSPFv3:

1 Navigate to **MANAGE | System Setup | Network > Routing > OSPFv3**.

2 Click the **Edit** icon for the interface. The **Interface OSPFv3 Configuration** dialog displays.



3 From **OSPFv3**, select:

Disabled (default)	OSPF Router is disabled on this interface. Go to Step 12.
Enabled	OSPF Router is enabled on this interface.
Passive	The OSPF router is enabled on this interface, but only advertises connected networks using type 1 LSAs (Router Link Advertisements) into the local area. All the options except **OSPFv3 Area** are dimmed.

4 Enter the Area ID in the **OSPF Area** field. The OSPF Area can be represented in either IP or decimal notation. For example, the area connected to `X4:100` as either `100.100.100.100` or `1684300900`. The default is **0**.

5 If you selected **Passive** for **OSPFv3**, go to Step 12.

6 To specify the period after which an entry in the LSDB is removed if Hello is not received, enter the time, in seconds, in the **Dead Interval (1 - 65535)** field. The default is **40** seconds, with a minimum of 1 and a maximum on 65,535.

IMPORTANT: Be sure this value agrees with the other OSPF routers on the segment for successful neighbor establishment.

7 Select the OSPFv3 area type from **OSPFv3 Area Type** (for a detailed description of these settings, see OSPF Terms on page 459):

Normal	Default; receives and sends all applicable LSA types.
Stub Area	Does not receive type 5 LSAs (AS External Link Advertisements).
Totally Stubby Area	Does not receive LSA types 3, 4, or 5.

8 To specify the period of time between Hello packets, enter the time, in seconds, in the **Hello Interval (1 - 65535)** field. The default is **10** seconds, with a minimum of 1 and a maximum on 65,535.

IMPORTANT: Be sure this value agrees with the other OSPF routers on the segment for successful neighbor establishment.

9 To:

- Specify the overhead of sending packets across this interface, enter the overhead in the **Interface Cost (1 - 65535)** field. The default value is **0**, generally used to indicate an Ethernet interface. The minimum and default value is 0 (for example, Fast Ethernet) and the maximum value is 65,535 (for example, pudding).
- Have the cost determined automatically, select **Auto**, which dims the **Interface Cost** field. This option is selected by default.

10 To specify the router priority value is used in determining the Designated Router (DR) for a segment, enter the value in the **Router Priority (0-255)** field. The higher the value, the higher the priority. For a priority tie, the Router ID acts as the tie-breaker. Setting a value of 0 makes the OSPF router on this interface ineligible for DR status. The default value is **1**, and the maximum value is 255.

11 To configure the instance ID for the interface, enter a value in the **Instance-ID (0 - 255)** field. The minimum and default is 0, and the maximum is 255. This option is not selected by default.

> **IMPORTANT:** This option is normally dimmed and should be set only through the SonicOS command line interface (for information about the SonicOS CLI, see *SonicOS Command Line Interface*.

12 Click **OK**.

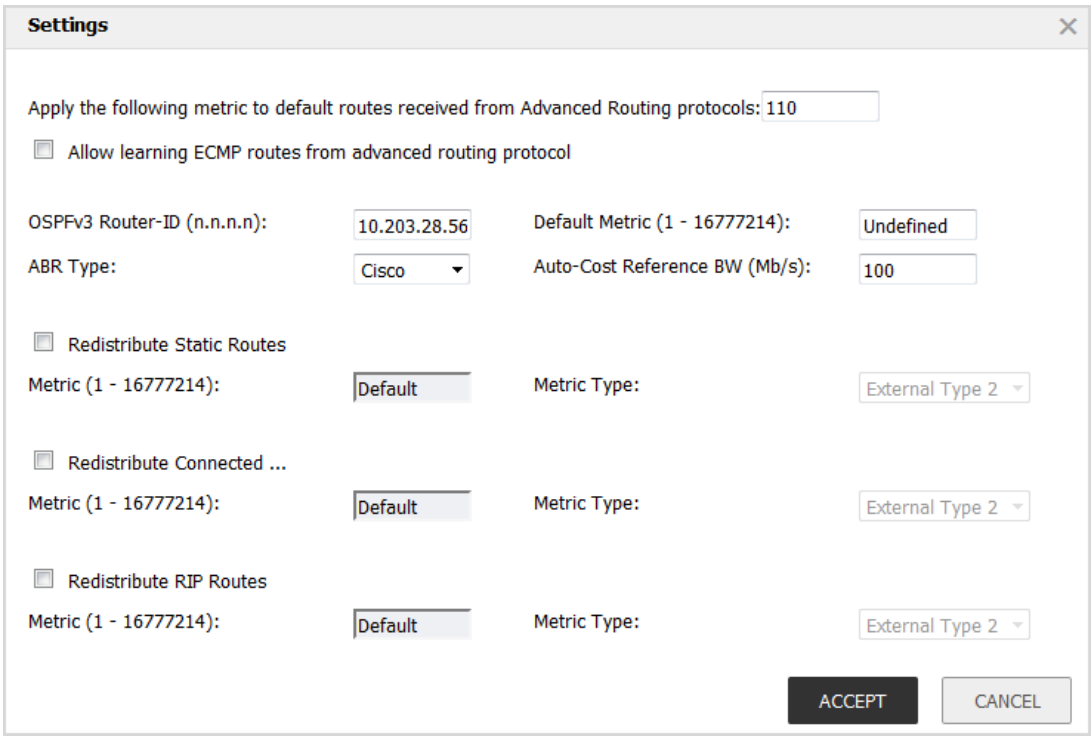
Global OSPFv3 Configuration

To configure global OSPFv3:

1 Navigate to **MANAGE | System Setup | Network > Routing**.

2 Click **OSPFv3**.

3 Click the **Setting** icon. The **Settings** popup dialog displays:



4 Configure these options:

- **OSPFv3 Router ID (n.n.n.n)** – The Router ID can be any value, represented in IP address notation. It is unrelated to the any of the IP addresses on the Security Appliance, and can be set to any *unique* value within your OSPF network.
- **ABR Type** – Allows for the specification of the topology with which this OSPF router is participating, for the sake of compatibility. The options are:
  - **Standard** – Full RFC2328 compliant ABR OSPF operation.
  - **Cisco** – For interoperating with Cisco's ABR behavior, which expects the backbone to be configured and active before setting the ABR flag.
  - **IBM** – For interoperating with IBM's ABR behavior, which expects the backbone to be configured before settings the ABR flag.
  - **Shortcut** – A shortcut area enables traffic to go through the non-backbone area with a lower metric whether or not the ABR router is attached to area 0.
- **Default Metric (1-16777214)** – Specifies the metric used when redistributing routes from other (Default, Static, Connected, RIP, or VPN) routing information sources. The default value (**Undefined**) is **1**, and the maximum is 16,777,214.
- **Auto-Cost Reference B@ (Mb/s)** – The default is 100.
- **Redistribute Static Routes** – Enables or disables the advertising of static (Policy-based Routing) routes into the OSPF system. This option is not selected by default.

NOTE: The following applies to all Redistributed routes:

- **Metric** – Can be explicitly set for this redistribution, or it can use the value (**Default**) specified in the **Default Metric** option.
- **Metric Type** – The redistributed route advertisement i an LSA Type 5, and the type may be selected as either **External Type 1** (adds the internal link cost) or **External Type 2** (only uses the external link cost).

NOTE: These fields are dimmed unless the Redistributed route option is selected.

- **Redistribute Connected Networks** - Enables or disables the advertising of locally connected networks into the OSPF system. This option is not selected by default.
- **Redistribute RIP Routes** - Enables or disables the advertising of routes learned through RIP into the OSPF system. This option is not selected by default.

5 Click **ACCEPT**.

The **Routing Protocols** section shows the status of all active OSPF routers by interface.

The **Routing Policies** section shows routes learned by OSPF as **OSPF** or **RIP Routes**.

Status becomes available.

Configuring RIP and RIPng

Topics:

- Configuring RIP on page 486

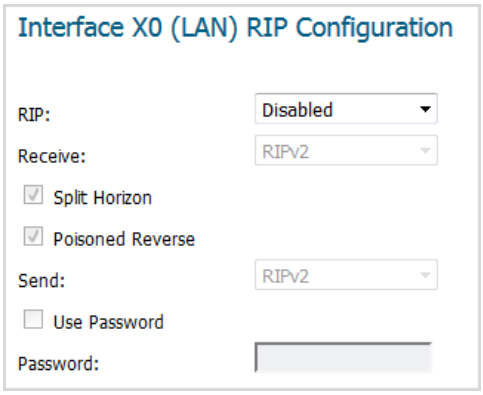
Configuring RIP

To configure RIP routing on an interface:

1 Navigate to **MANAGE | Network > Routing**.

2 Click **RIP**.

3 Click the **Edit** icon for the interface. The **Interface RIP Configuration** dialog displays.



4 From **RIP**, select a mode:

Disabled (default)	RIP is disabled on this interface; go to Step 12.
Send and Receive	The RIP router on this interface sends updates and processes received updates.
Send Only	The RIP router on this interface only sends updates and does not process received updates. This is similar to the basic routing implementation.
Receive Only	The RIP router on this interface only processes received updates.
Passive	The RIP router on this interface does not process received updates and only sends updates to neighboring RIP routers specified with the CLI `neighbor` command. **IMPORTANT:** This mode should be used only when configuring advanced RIP options from the ARS-RIP CLI. When selected, all other options are dimmed.

5 If you specified:

- **Send Only**, go to Step 8.
- **Passive**, go to Step 12.

6 From **Receive**, select the RIP version for receiving RIP packets:

RIPv1	Receive only *broadcast* RIPv1 packets.
RIPv2 (default)	Receive only *multicast* RIPv2 packets. RIPv2 packets are sent by multicast, although some implementations of RIP routers (including basic routing on SonicWall devices) have the ability to send RIPv2 in either broadcast or multicast formats. **IMPORTANT:** Be sure the device sending RIPv2 updates uses multicast mode, or the updates are not processed by the ars-rip router.

7 If you selected **Receive Only** for **RIP**, go to Step 11.

8 To suppress the inclusion of routes sent in updates to routers from which they were learned, select **Split Horizon**. This is a common RIP mechanism for preventing routing loops; see Maximum Hops on page 457. This option is selected by default.

9 To specify an optional mode of Split Horizon operation, select **Poisoned Reverse**. Rather than suppressing the inclusion of learned routes, the routes are sent with a metric of infinity (16), thus indicating that they are unreachable; see Maximum Hops on page 457. This option is selected by default.

10 From **Send**, select the RIP version for sending packets:

RIPv1	Send *broadcast* RIPv1 packets.

RIPv2 - v1 compatible	Send *multicast* RIPv2 packets that are compatible with RIPv1.
RIPv2 (default)	Send *multicast* RIPv2 packets.

11 To enforce use of a password, select **Use Password**. The **Password** field becomes available. This option is not selected by default.

   a Enter the password in the **Password** field.

12 Click **OK**.

Configuring RIPng

To configure RIPng routing on an interface:

1 Navigate to **MANAGE | System Setup | Network > Routing > RIPng**.

2 Click the **Edit** icon for the interface. The **Interface RIPng Configuration** dialog displays.



3 From **RIPng**, select a mode:

Disabled (default)	RIPng is disabled on this interface; go to Step 6.
Enable	The RIPng router on this interface sends updates and processes received updates.
Passive	The RIPng router on this interface does not process received updates and only sends updates to neighboring RIPng routers specified with the CLI `neighbor` command. **IMPORTANT:** This mode should be used only when configuring advanced RIPng options from the ARS-RIP CLI.

4 To suppress the inclusion of routes sent in updates to routers from which they were learned, select **Split Horizon**. This is a common RIP mechanism for preventing routing loops; see Maximum Hops on page 457. This option is selected by default.

5 To specify an optional mode of Split Horizon operation, select **Poisoned Reverse**. Rather than suppressing the inclusion of learned routes, the routes are sent with a metric of infinity (16), thus indicating that they are unreachable; see Maximum Hops on page 457. This option is selected by default.

6 Click **OK**.

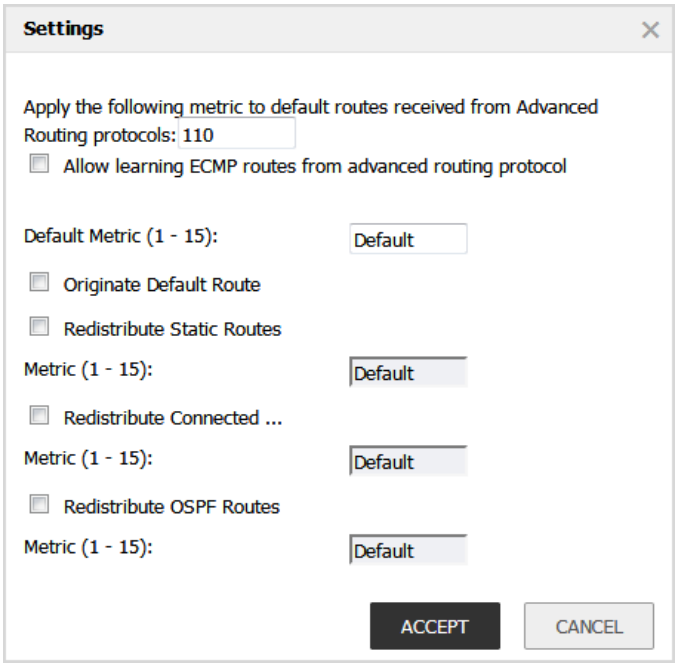
Global RIPng Configuration

To configure global OSPFv3:

1 Navigate to **MANAGE | System Setup | Network > Routing**.

2 Click **OSPFv3**.

3 Click the **Setting** icon. The **Settings** popup dialog displays:



4 Configure these options:

- **Default Metric** – Used to specify the metric that is used when redistributing routes from other (Default, Static, Connected, OSPF, or VPN) routing information sources. The default value (undefined) is **1** and the maximum is 15.
- **Originate Default Route** – This checkbox enables or disables the advertising of the Security Appliance's default route into the RIP system.
- **Redistribute Static Routes** – Enables or disables the advertising of static (Policy-based Routing) routes into the RIP system. The metric can be explicitly set for this redistribution, or it can use the value (default) specified in the **Default Metric** setting.
- **Redistribute Connected Networks** - Enables or disables the advertising of locally connected networks into the RIP system. The metric can be explicitly set for this redistribution, or it can use the value (default) specified in the **Default Metric** setting.
- **Redistribute OSPF Routes** - Enables or disables the advertising of routes learned through OSPF into the RIP system. The metric can be explicitly set for this redistribution, or it can use the value (default) specified in the **Default Metric** setting.

5 Click **ACCEPT**.

Configuring Advanced Routing for Tunnel Interfaces

VPN Tunnel Interfaces can be configured for advanced routing. To do so, you must enable advanced routing for the tunnel interface on the **Advanced** tab of its configuration.

After you have enabled advanced routing for a Tunnel Interface, it is displayed in the tables with the other interfaces in the various views on **Network > Routing**.

To configure Advanced Routing options:

1 Click the **Edit** icon in the **Configure RIP/RIPng** or **Configure OSPF/OSPFv3** column for the Tunnel Interface you wish to configure. The RIP and OSPF configurations for Tunnel Interfaces are very similar to the configurations for traditional interfaces.

Global Unnumbered Configuration

Because unnumbered Tunnel Interfaces are not physical interfaces and have no inherent IP address, they must "borrow" the IP address of another interface. Therefore, the advanced routing configuration for a Tunnel Interface includes the following options for specifying the source and destination IP addresses for the tunnel:

- **IP Address Borrowed From** - The interface whose IP address is used as the source IP address for the Tunnel Interface.

  NOTE: The borrowed IP address must be a static IP address.

- **Remote IP Address** - The IP address of the remote peer to which the Tunnel Interface is connected. In the case of a SonicWall-to-SonicWall configuration with another Tunnel Interface, this should be the IP address of the borrowed interface of the Tunnel Interface on the remote peer.

Guidelines for Configuring Tunnel Interfaces for Advanced Routing

The following guidelines ensure success when configuring Tunnel Interfaces for advanced routing:

- The borrowed interface must have a static IP address assignment.
- The borrowed interface cannot have RIP or OSPF enabled on its configuration.

  TIP: SonicWall recommends creating a VLAN interface that is dedicated solely for use as the borrowed interface. This avoids conflicts when using wired connected interfaces.

- The IP address of the borrowed interface should be from a private address space, and should have a unique IP address in respect to any remote Tunnel Interface endpoints.
- The Remote IP Address of the endpoint of the Tunnel Interface should be in the same network subnet as the borrowed interface.
- The same borrowed interface may be used for multiple Tunnel Interfaces, provided that the Tunnel interfaces are all connected to different remote devices.
- When more than one Tunnel Interface on an appliance is connected to the same remote device, each Tunnel Interface must use a unique borrowed interface.

Depending on the specific circumstances of your network configuration, these guidelines may not be essential to ensure that the Tunnel Interface functions properly. But these guidelines are SonicWall best practices that avoid potential network connectivity issues.

Configuring BGP Advanced Routing

NOTE: BGP is supported on the:

- NSA 2600 and above Security Appliances.
- TZ400 series, TZ500 series, and TZ600 Security Appliances with the purchase of a SonicOS Expanded License.

BGP is not supported on the TZ300 series or SOHO Wireless Security Appliances.

Border Gateway protocol (BGP) is a large-scale routing protocol used to communicate routing information between Autonomous Systems (ASs), which are well-defined, separately administered network domains. BGP support allows for Security Appliances to replace a traditional BGP router on the edge of a network's AS. The current SonicWall implementation of BGP is most appropriate for single-provider/single-homed environments, where the network uses one ISP as their Internet provider and has a single connection to that provider. SonicWall BGP is also capable of supporting single-provider/multi-homed environments, where the network uses a single ISP but has a small number of separate routes to the provider. BGP is enabled on the **Network > Routing** page of the SonicOS Management Interface, and then it is fully configured through the SonicOS Command Line Interface (CLI).

For complete information on SonicWall's implementation of BGP, see BGP Advanced Routing on page 964.

Configuring an IPSec Tunnel for BGP Sessions

BGP transmits packets in the clear. For strong security, therefore, SonicWall recommends configuring an IPSec tunnel to use for BGP sessions. For how to configure an IPSec Tunnel for BGP and to enable BGP, see BGP Advanced Routing on page 964.

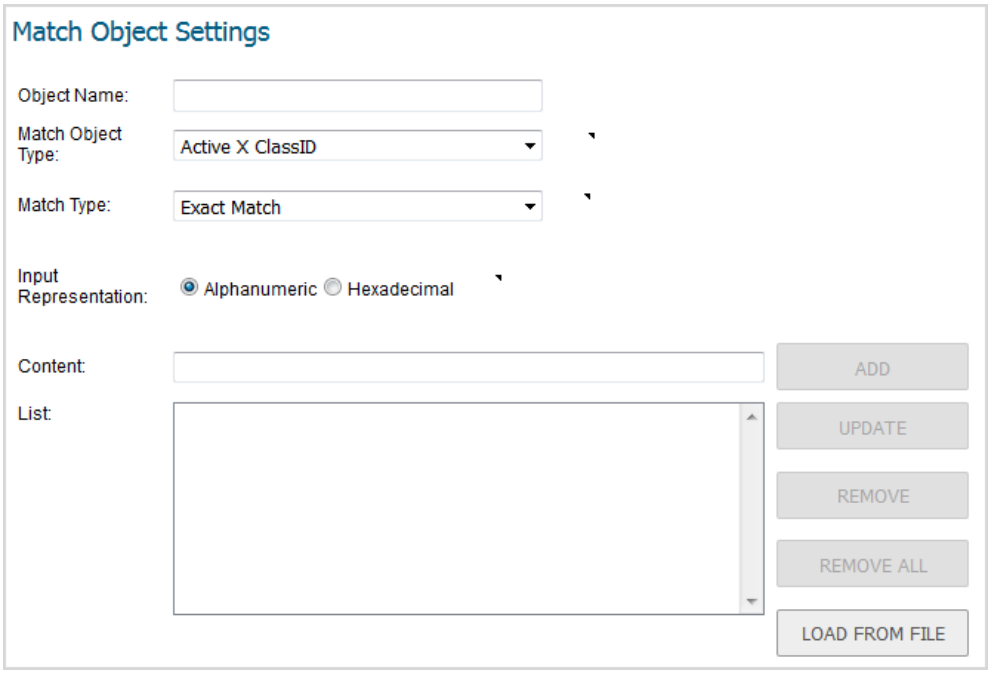
After BGP has been enabled through the Management Interface, the specifics of the BGP configuration are performed using the SonicOS command line interface (CLI). For complete information on the implementation of BGP on a SonicWall Security Appliance, see BGP Advanced Routing on page 964.

Configuring App-based Routes

To configure an app-based route entry:

1 Navigate to **MANAGE | Policies > Objects > Match Objects**.

2 From **Add**, select **Match Object**. The **Create Match Object** dialog displays.



3 From **Match Object Type**, select

- **Application Category List**
- **Application List**

4 Complete the configuration of the match object.

5 Click **OK**.

6 Navigate to **MANAGE | System Setup > Network > Routing**.

7 Create a route policy.

8 Click **OK**.

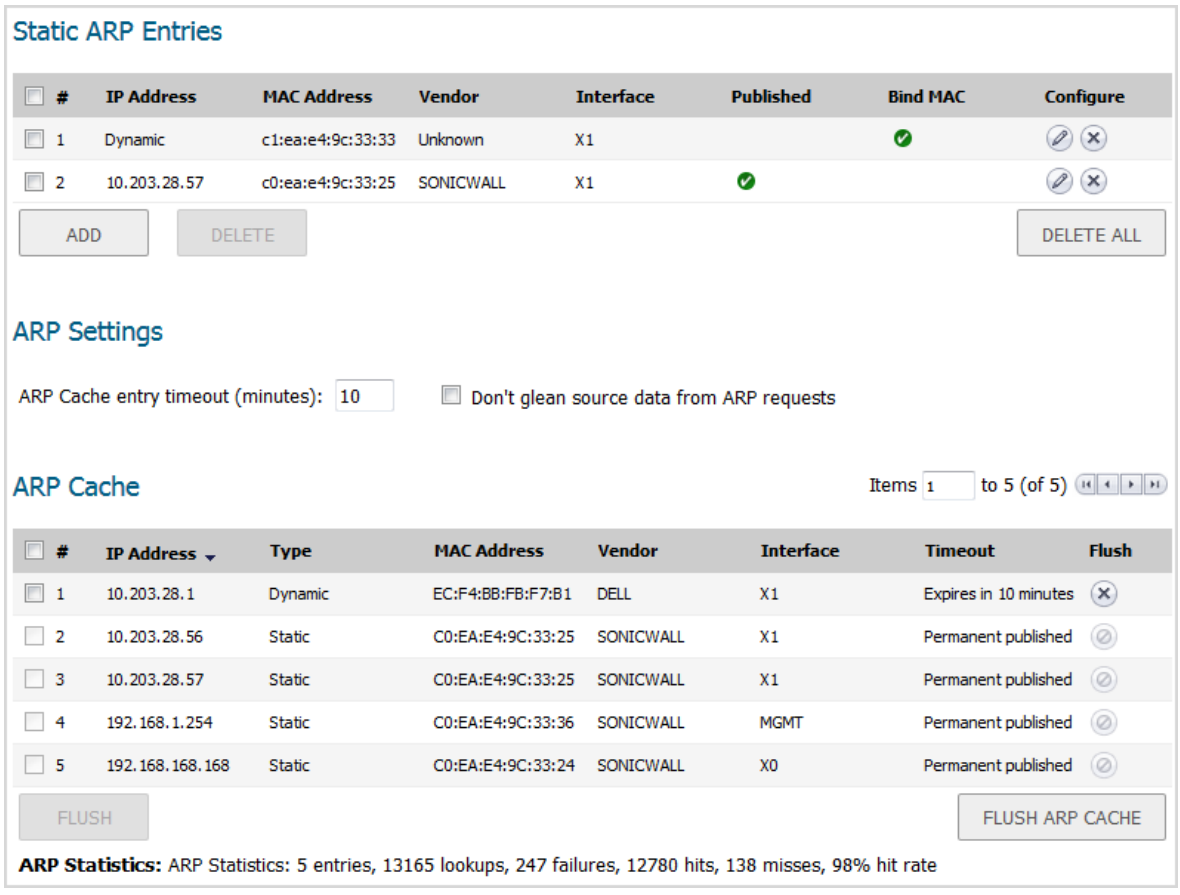
23

Managing ARP Traffic

Topics:

- **Network > ARP** on page 492
  - Static ARP Entries on page 493
  - ARP Settings on page 496
  - ARP Cache on page 496

Network > ARP



ARP (Address Resolution Protocol) maps layer 3 (IP addresses) to layer 2 (physical or MAC addresses) to enable communications between hosts residing on the same subnet. ARP is a broadcast protocol that can create excessive amounts of network traffic on your network. To minimize the broadcast traffic, an ARP cache is maintained to store and reuse previously learned ARP information.

Topics:

- Static ARP Entries on page 493
- ARP Settings on page 496
- ARP Cache on page 496

Static ARP Entries

The Static ARP feature allows for static mappings to be created between layer 2 MAC addresses and layer 3 IP addresses.

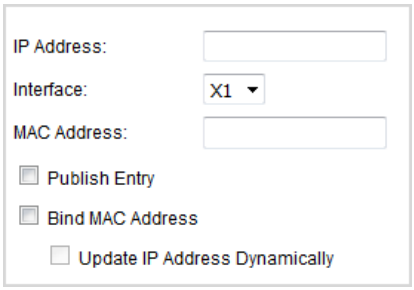
Topics:

- Configuring a Static ARP on page 493
- Editing a Static ARP Entry on page 494
- Secondary Subnets with Static ARP on page 494
- Viewing Static ARP Entries on page 496

Configuring a Static ARP

To configure a Static ARP:

1. Navigate to **Network > ARP**.
2. Under the **Static ARP Entries** table, click **Add**. The **Add Static ARP** dialog displays.



3. In the **IP Address** field, enter the IP address of the SonicWall Security Appliance.
4. From **Interface**, select the LAN interface on the Security Appliance to be associated with this static ARP entry.
5. In the **MAC Address** field, enter the MAC address of the Security Appliance.
6. To cause the Security Appliance to respond to ARP queries for the specified IP address with the specified MAC address, select the **Publish Entry** option. This option is not selected by default.

   This option can be used, for example, to have the Security Appliance reply for a secondary IP address on a particular interface by adding the MAC address of the Security Appliance. See Secondary Subnets with Static ARP on page 494. Selecting this option dims the **MAC Address** field and **Bind MAC Address** option.
7. If you selected **Publish Entry**, go to Step 10.
8. To bind the MAC address specified to the designated IP address and interface, select **Bind MAC Address**. This option is not selected by default.

This option ensures that a particular workstation (as recognized by the network card's unique MAC address) can only the used on a specified interface on the Security Appliance. After the MAC address is bound to an interface, the Security Appliance:

- Does not respond to that MAC address on any other interface.
- Removes any dynamically cached references to that MAC address that might have been present.
- Prohibits additional (non-unique) static mappings of that MAC address.

When Bind MAC Address is selected, **Update IP Address Dynamically** becomes available.

9 To allow a MAC address to be bound to an interface when DHCP is being used to dynamically allocate IP addressing, select **Update IP Address Dynamically**, which is a sub-feature of the **Bind MAC Address** option.

Enabling this option dims the **IP Address** field and sets it to `0.0.0.0`, makes the **MAC Address** field available, and populates the ARP Cache with the IP address allocated by either the Security Appliance's internal DHCP server or, if IP Helper is in use, by the external DHCP server.

10 Click **OK**.

Editing a Static ARP Entry

To edit a Static ARP entry:

1 Navigate to **Network > ARP**.

2 In the **Static ARP Entries** table, click the entry's **Edit** icon in the **Configure** column. The **Edit Static ARP** dialog displays.



3 Make the changes.

4 Click **OK**. The entry is updated.

Secondary Subnets with Static ARP

The Static ARP feature allows for secondary subnets to be added on other interfaces without the addition of automatic NAT rules.

Topics:

- Adding a Secondary Subnet on page 495
- An Example on page 495

Adding a Secondary Subnet

To add a Secondary Subnet using the Static ARP Method:

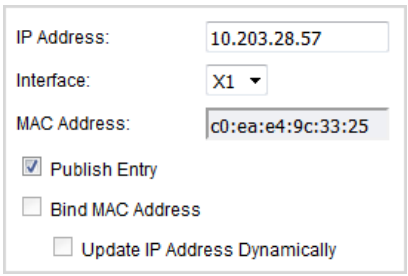
1. Add a 'published' static ARP entry for the gateway address that is used for the secondary subnet, assigning it the MAC address of the Security Appliance interface to which it is connected.
2. Add a static route for that subnet, so that the Security Appliance regards it as valid traffic, and knows to which interface to route that subnet's traffic.
3. Add Access Rules to allow traffic destined for that subnet to traverse the correct network interface.
4. Optional: Add a static route on upstream device(s) so that they know which gateway IP to use to reach the secondary subnet.

An Example

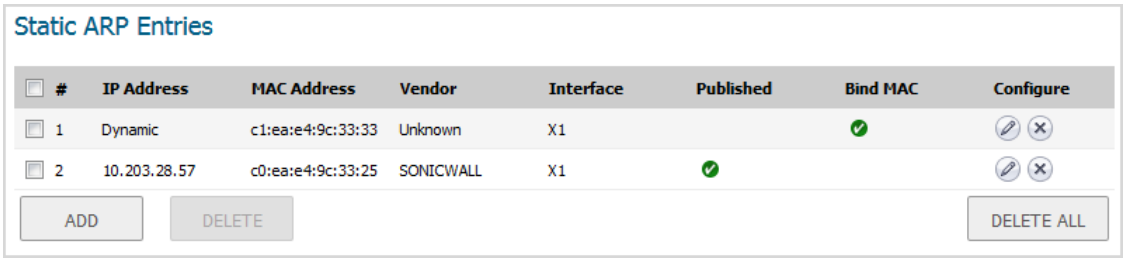
Consider the following network example (see Adding a Secondary Subnet on page 495).

To support the added configuration:

1. Create a published static ARP entry for `192.168.50.1`, the address that serves as the gateway for the secondary subnet.
2. Associate it with the appropriate LAN interface. From **Network > ARP**, click **Add** under the **Static ARP Entries** table.
3. Add this entry:

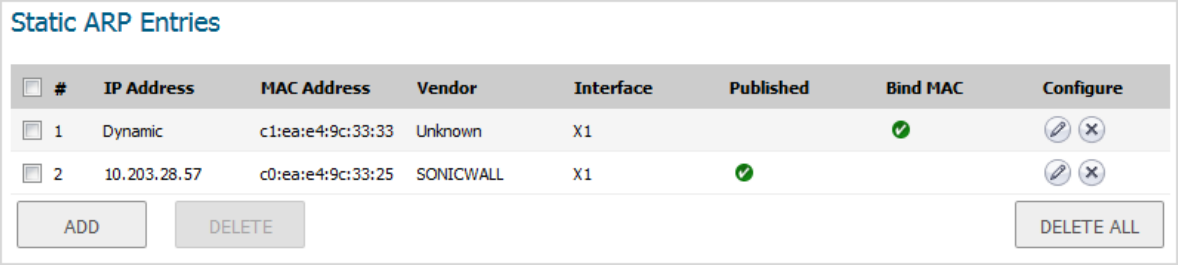


4. Click **OK**. The entry appears in the **Static ARP Entries** table.



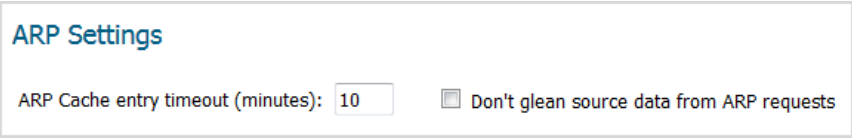
5. Navigate to **Network > Routing**.
6. Add a static route for the `192.168.50.0/24` network, with the `255.255.255.0` subnet mask on the X3 Interface. For information about adding static routes, see Configuring Route Advertisements and Route Policies on page 450.
7. To allow traffic to reach the `192.168.50.0/24` subnet and to allow the `192.168.50.0/24` subnet to reach the hosts on the LAN, navigate to **Policies | Rules > Access Rules** page.
8. Add appropriate Access Rules to allow traffic to pass. For information about adding Access Rules, see *SonicOS Polices*.

Viewing Static ARP Entries



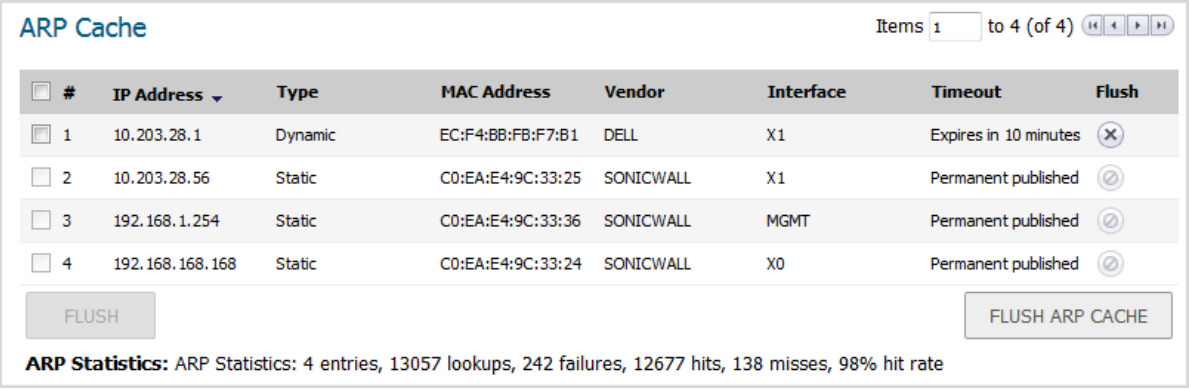
IP Address	IP address of the Security Appliance serving as the gateway.
MAC Address	MAC address of the Security Appliance serving as the gateway.
Vendor	Name of the Security Appliance's manufacturer.
Interface	LAN interface associated with this entry.
Published	Indicates with a green checkmark whether the Security Appliance responds to ARP queries for the specified IP address with the specified MAC address.
Bind MAC	Indicates with a green checkmark whether the MAC address is bound to the designated IP address and interface.
Configure	Displays the **Edit** and **Delete** icons for the entry.

ARP Settings



ARP Cache entry timeout (minutes)	Specify a length of time for the entries to time out and be flushed from the cache. The minimum time is two minutes, the maximum is 600 minutes (10 hours), and the default is **10** minutes.
Don't glean source data from ARP requests	Select to prevent source data from being obtained from ARP requests. This option is not selected by default.

ARP Cache



IP Address	IP Address of the Security Appliance.
Type	Indicates whether the ARP is **Static** or **Dynamic**.
MAC Address	MAC address associated with the IP Address.
Vendor	Name of the Security Appliance's manufacturer.
Interface	LAN interface associated with this ARP entry.
Timeout	Indicates the time left in cache for this entry. If the entry was published when configured, **Timeout** displays `Permanent published`.
Flush	Displays the **Delete** icon for flushing the entry from ARP cache. **NOTE:** Only **Dynamic** entries have the **Delete** icon.

Flushing the ARP Cache

It is sometimes necessary to flush the ARP cache if the IP address has changed for a device on the network. As the IP address is linked to a physical address, the IP address can change, but still be associated with the physical address in the ARP Cache. Flushing the ARP Cache allows new information to be gathered and stored in the ARP Cache.

TIP: To configure a specific length of time for an entry to time out, enter a value in minutes in the **ARP Cache entry time out (minutes)** field; see ARP Settings on page 496.

To flush a dynamic entry in the ARP Cache table:

1 Click its **Delete** icon in the **Flush** column.

To flush one or more dynamic entries in the ARP Cache table:

1 Select the checkbox(es) of one or more entries to be flushed. **Flush** becomes active.

2 Click **Flush**.

To flush all the dynamic entries in the ARP Cache table:

1 Click **Flush ARP Cache**.

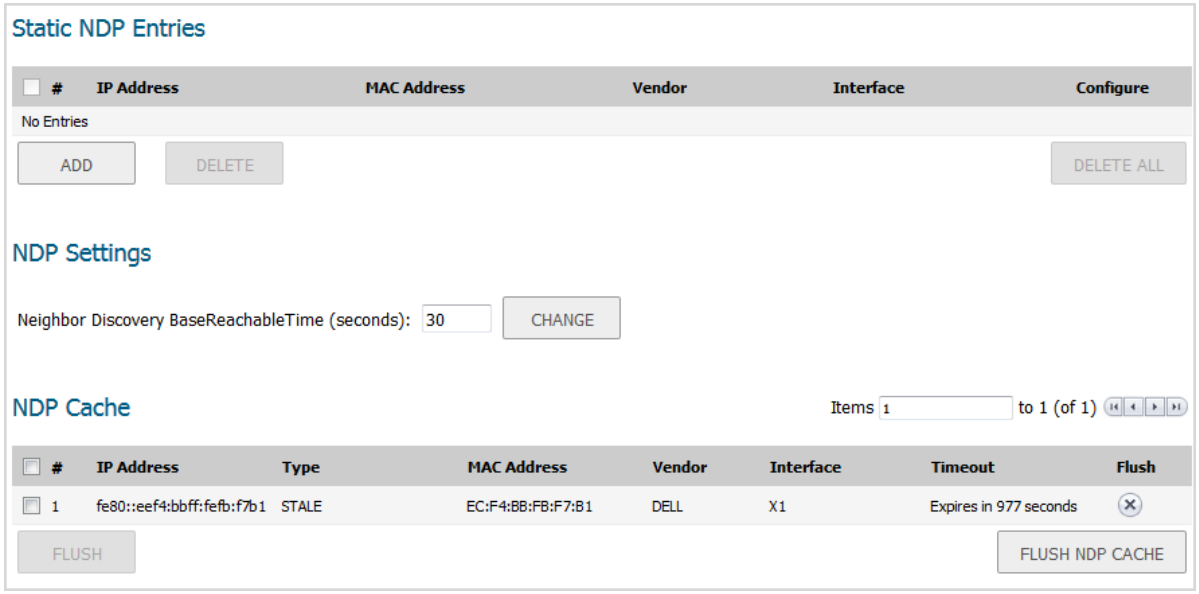
24

Configuring Neighbor Discovery Protocol

Topics:

- **Network > Neighbor Discovery (IPv6 Only)** on page 498
  - Static NDP Entries on page 499
  - NDP Settings on page 500
  - NDP Cache on page 500
  - Configuring a Static NDP Entry on page 501
  - Editing a Static NDP Entry on page 501
  - Flushing the NDP Cache on page 502

Network > Neighbor Discovery (IPv6 Only)



The Neighbor Discovery Protocol (NDP) is a new messaging protocol that was created as part of IPv6 to perform a number of the tasks that ICMP and ARP accomplish in IPv4. Just like ARP, Neighbor Discovery builds a cache of dynamic entries, and you can configure static Neighbor Discovery entries. the IPv4/IPv6 neighbor messages and functions table shows the IPv6 neighbor messages and functions that are analogous to the traditional IPv4 neighbor messages.

IPv4/IPv6 neighbor messages and functions

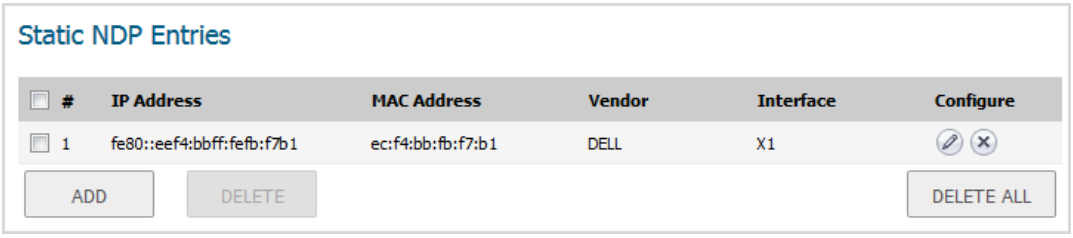
IPv4 Neighbor message	IPv6 Neighbor message
ARP request message	Neighbor solicitation message
ARP reply message	Neighbor advertisement message
ARP cache	Neighbor cache
Gratuitous ARP	Duplicate address detection
Router solicitation message (optional)	Router solicitation (required)
Router advertisement message (optional)	Router advertisement (required)
Redirect message	Redirect Message

The Static NDP feature allows for static mappings to be created between a Layer 3 IPv6 address and a Layer 2 MAC address.

Topics:

- Static NDP Entries on page 499
- NDP Settings on page 500
- NDP Cache on page 500
- Configuring a Static NDP Entry on page 501
- Editing a Static NDP Entry on page 501
- Flushing the NDP Cache on page 502

Static NDP Entries



IP Address	IPv6 IP address for the remote device.
MAC Address	MAC address for the remote device.
Vendor	Name of the remote device's manufacturer.
Interface	Interface associated with the remote device.
Configure	Contains the **Edit** and **Delete** icons for the entry.

NDP Settings



You specify the maximum time to reach a neighbor in **NDP Settings**.

NOTE: For IPv6, this value also can be set for each interface on the **Network > Interfaces > Edit Interface > Advanced** dialog. If Router Advertisement is enabled on an interface, the value set for the interface is used for that interface only. For more information, see Configuring Interfaces on page 272.

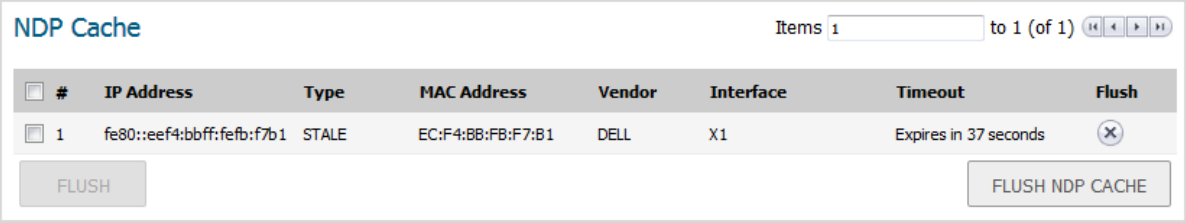
To specify the maximum time:

1 Enter a number in the **Neighbor Discover BaseReachableTime (seconds)** field. The minimum is 0 seconds, the maximum is 3600 seconds, and the default is **20** seconds.

TIP: When this option's value is set to 0, the global value of NDP settings is used.

2 Click **Change**.

NDP Cache



The **NDP Cache** table displays all current IPv6 neighbors.

IP Address	IPv6 IP Address of the neighbor device.
Type	Type of neighbor: • **REACHABLE** - The neighbor is known to have been reachable within 30 seconds. • **STALE** - The neighbor is no longer known to be reachable, and traffic has been sent to the neighbor within 1200 seconds. • **STATIC** - The neighbor was manually configured as a static neighbor.,
MAC Address	IPv6 MAC Address of the neighbor device.
Vendor	Name of the neighbor device's manufacturer.
Interface	Interface associated with this neighbor device.
Timeout	The length of inactivity time until the user times out.
Flush	Contains the **Delete** icon for the entry.

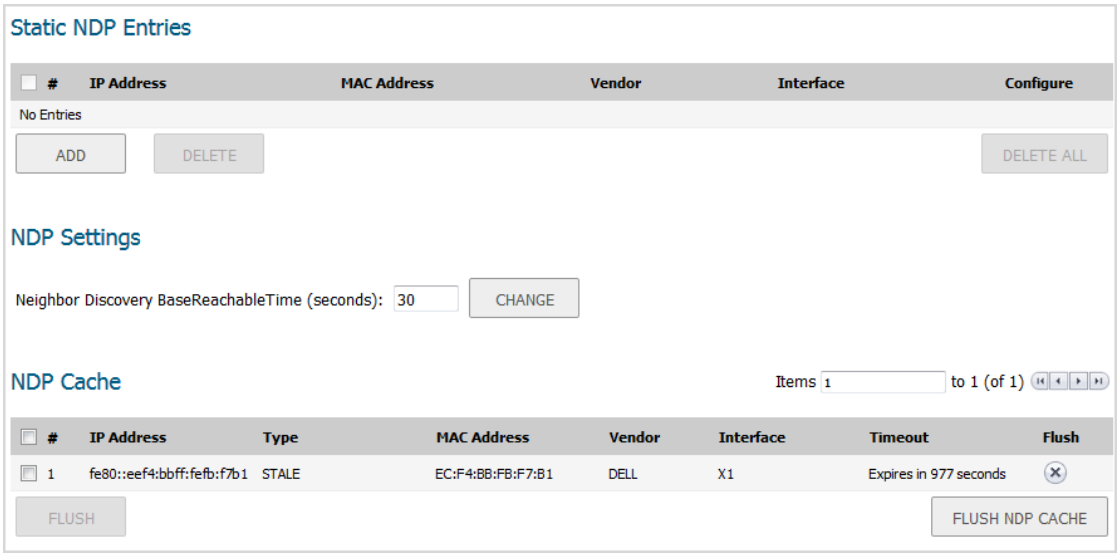
These types of neighbors are displayed:

- **REACHABLE** - The neighbor is known to have been reachable within 30 seconds.
- **STALE** - The neighbor is no longer known to be reachable, and traffic has been sent to the neighbor within 1200 seconds.
- **STATIC** - The neighbor was manually configured as a static neighbor.

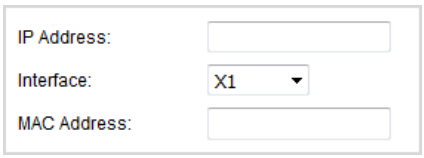
Configuring a Static NDP Entry

To configure a Static NDP entry:

1 Navigate to the **Network > Neighbor Discovery** page.



2 Under the **Static NDP Entries** table, click **Add**. The **Add Static NDP** dialog displays.



3 In the **IP Address** field, enter the IPv6 address for the remote device.

4 From **Interface**, select the interface on the SonicWall Security Appliance that is used for the entry.

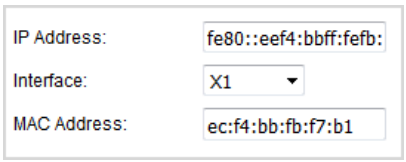
5 In the **MAC Address** field, enter the MAC address of the remote device.

6 Click **OK**. The static NDP entry is added.

Editing a Static NDP Entry

To edit a Static NDP entry:

1 In the **Static NDP Entries** table, click the entry's **Edit** icon in the **Configure** column. The **Edit Static NDP** dialog displays.



2 Make the changes.

3 Click **OK**. The entry is updated.

Flushing the NDP Cache

It is sometimes necessary to flush the NDP cache if the IP address has changed for a device on the network. As the IP address is linked to a physical address, the IP address can change, but still be associated with the physical address in the NDP Cache. Flushing the NDP Cache allows new information to be gathered and stored in the NDP Cache.

TIP: To configure a specific length of time for an entry to time out, enter a value in minutes in the **NDP Cache entry time out (minutes)** field; see NDP Settings on page 500.

To flush an entry in the NDP Cache table:

1 Click its **Delete** icon in the **Flush** column.

To flush one or more entries in the NDP Cache table:

1 Select the checkbox of one or more entries to be flushed. The two flush buttons become active.

2 Click either **Flush** or **Flush NDP Cache**.

To flush all the entries in the NDP Cache table:

1 Select the checkbox in the **NDP Cache** table header. The two flush buttons become active.

2 Click either **Flush** or **Flush NDP Cache**.

25

Configuring MAC-IP Anti-spoof

Topics:

- About MAC-IP Anti-spoof Protection on page 503
  - Extension to IP Helper on page 504
- **Network > MAC-IP Anti-spoof** on page 504
  - Settings for Interface(s) on page 505
  - Anti-Spoof Cache on page 507
  - Spoof Detected List on page 508
- Configuring MAC-IP Anti-spoof Protection on page 508
  - Displaying Traffic Statistics on page 509
  - Editing MAC-IP Anti-spoof Settings for an IPv6 Interface on page 509
  - Editing MAC-IP Anti-spoof Settings for an IPv4 Interface on page 510
  - Adding Devices to Anti-Spoof Cache on page 512
  - Deleting Anti-Spoof Cache Entries on page 512
  - Filtering What Is Displayed on page 513
  - Adding Static Entries from Spoof Detected List on page 513

About MAC-IP Anti-spoof Protection

MAC and IP address-based attacks are increasingly common in today's network security environment. These types of attacks often target a Local Area Network (LAN) and can originate from either outside or inside a network. In fact, anywhere internal LANs are somewhat exposed, such as in office conference rooms, schools, or libraries, could provide an opening to these types of attacks. These attacks also go by various names: man-in-the-middle attacks, ARP poisoning, SPITS. The MAC-IP Anti-spoof feature lowers the risk of these attacks by providing you with different ways to control access to a network and by eliminating spoofing attacks at OSI Layer 2/3.

The effectiveness of the MAC-IP Anti-spoof feature focuses on two areas:

- Admission control, which gives you the ability to select which devices gain access to the network.
- Elimination of spoofing attacks, such as denial-of-service attacks, at Layer 2.

To achieve these goals, two caches of information must be built: the MAC-IP Anti-spoof Cache, and the ARP Cache.

The MAC-IP Anti-spoof cache validates incoming packets and determines whether they are to be allowed inside the network. An incoming packet's source MAC and IP addresses are looked up in this cache. If they are found, the packet is allowed through. The MAC-IP Anti-spoof cache is built through one or more of these sub-systems:

- DHCP Server-based leases (SonicWall's - DHCP Server; IPv4 only)

- DHCP relay-based leases (SonicWall's - IP Helper; IPv4 only)
- Static ARP entries; IPv4 only
- User-created static entries

The ARP Cache is built through these subsystems:

- ARP packets; both ARP requests and responses; IPv4 only
- Static ARP entries from user-created entries; IPv4 only
- MAC-IP Anti-spoof Cache

The MAC-IP Anti-spoof subsystem achieves egress control by locking the ARP cache, so egress packets (those exiting the network) are not spoofed by a bad device or by unwanted ARP packets. This prevents a SonicWall Security Appliance from routing a packet to the unintended device, based on mapping. This also prevents man-in-the-middle attacks by refreshing a client's own MAC address inside its ARP cache.

Extension to IP Helper

To support leases from the DHCP relay subsystem of IP Helper (**Network > IP Helper**):

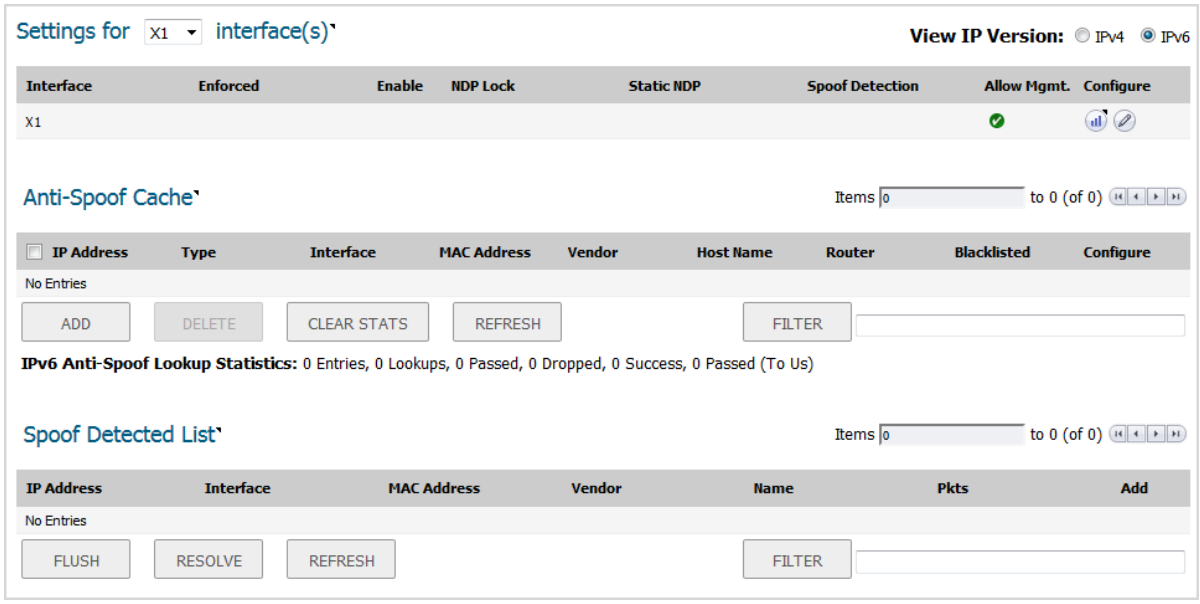
- As part of the DHCP relay logic, IP Helper learns leases exchanged between clients and the DHCP server, then saves them into flash memory.
- These learned leases are synchronized to the idle SonicWall Security Appliance, as part of the IP Helper state sync messages.

MAC and IP address bindings from the leases are transferred into the MAC-IP Anti-spoof cache.

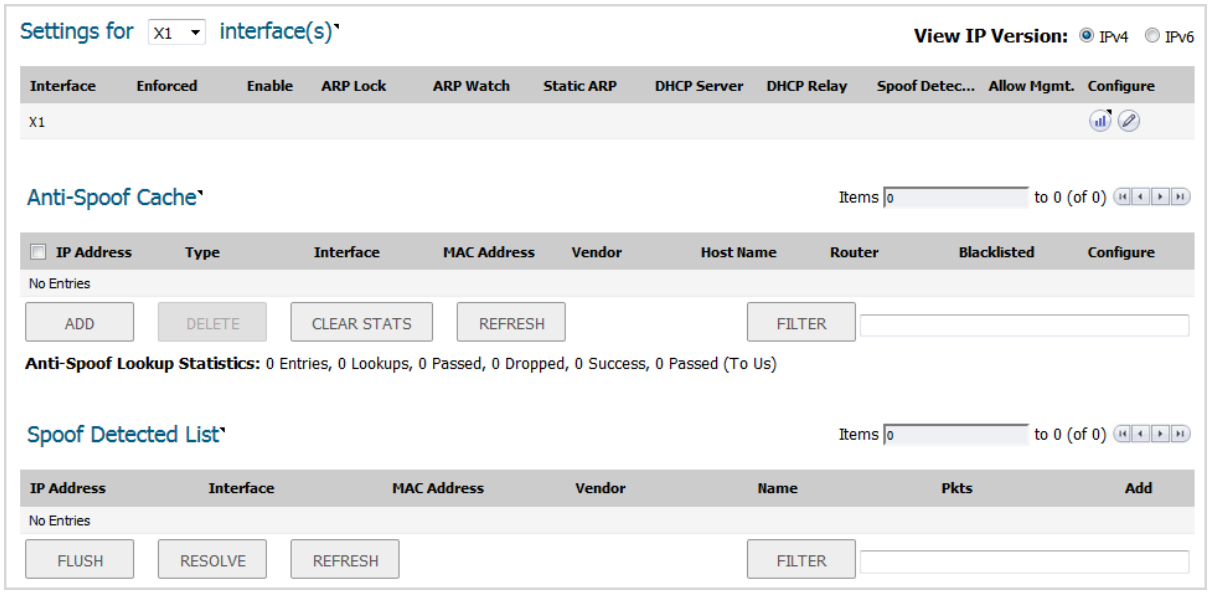
For more information about IP Helper, see Using IP Helper on page 541.

Network > MAC-IP Anti-spoof

IPv6



IPv4



This section describes how to plan, design, and implement MAC-IP Anti-spoof protection in SonicWall SonicOS.

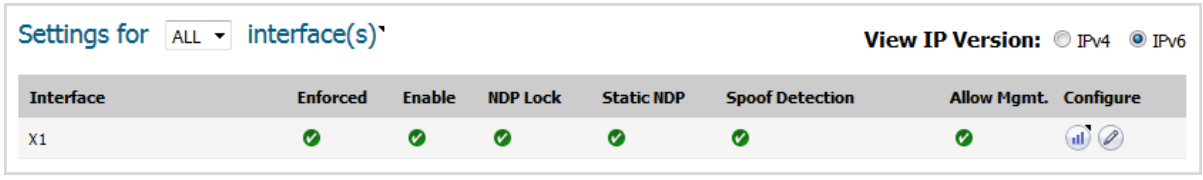
Topics:

- Settings for Interface(s) on page 505
- Anti-Spoof Cache on page 507
- Spoof Detected List on page 508

Settings for Interface(s)

NOTE: The green checkmark icons denote which settings have been enabled.

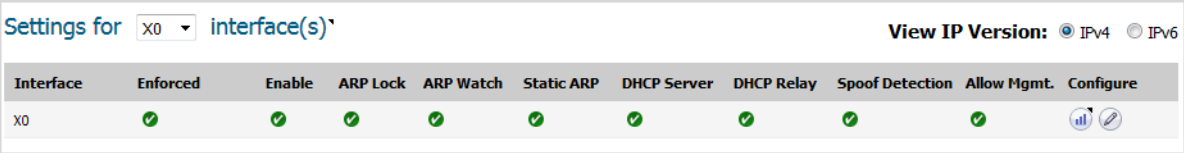
IPv6



Settings for interface(s)	Lists all interfaces on which MAC-IP Anti-spoof settings can be applied. The default for display is **All**.
Interface	Interface selected from **Settings for interface(s)**.
Enforced	Indicates whether ingress anti-spoof is enforced on this interface.
Enable	Indicates whether MAC-IP Anti-spoof is enabled on this interface. b
NDP Lock	Indicates whether MAC-IP Anti-spoof check is enabled for every transmit packet on this interface.
Static NDP	Indicates whether a corresponding MAC-IP Anti-spoof table entry is created for every static NDP entry.

Spoof Detection	Indicates whether a MAC-IP Anti-spoof-detected list is created for packets failing to match the anti-spoof cache. **NOTE:** These interfaces are excluded from the MAC-IP Anti-spoof list: • Non-ethernet interfaces • Port-shield member interfaces • Layer 2 bridge pair interfaces • High availability interfaces • High availability data interfaces
Allow Mgmt.	Indicates whether all traffic destined to the Security Appliance is allowed without a valid MAC-IP Anti-spoof cache.
Configure	Contains the **Statistics** and **Edit** icons for the entry.

IPv4



Settings for interface(s)	Lists all interfaces on which MAC-IP Anti-spoof settings can be applied. The default for display is **All**.
Interface	Interface selected from **Settings for interface(s)**.
Enforced	Indicates whether ingress anti-spoof is enforced on this interface.
Enable	Indicates whether MAC-IP Anti-spoof is enabled on this interface. b
ARP Lock	Indicates whether MAC-IP Anti-spoof check is enabled for every transmit packet on this interface.
ARP Watch	Indicates whether prevention of ARP poisoning of connected machines is enabled.
Static ARP	Indicates whether a corresponding MAC-IP Anti-spoof table entry is created for every static ARP entry.
DHCP Server	Indicates whether the MAC-IP Anti-spoof entry is populated from the DHCP Lease (SonicWall's DHCP server).
DHCP Relay	Indicates whether the MAC-IP Anti-spoof entry is populated from the DHCP Lease (DHCP relay - IP Helper).
Spoof Detection	Indicates whether a MAC-IP Anti-spoof-detected list is created for packets failing to match the anti-spoof cache. **NOTE:** These interfaces are excluded from the MAC-IP Anti-spoof list: • Non-ethernet interfaces • Port-shield member interfaces • Layer 2 bridge pair interfaces • High availability interfaces • High availability data interfaces
Allow Management	Indicates whether all traffic destined to the firewall is allowed without a valid MAC-IP Anti-spoof cache.
Configure	Contains the **Statistics** and **Edit** icons for the entry.

Anti-Spoof Cache

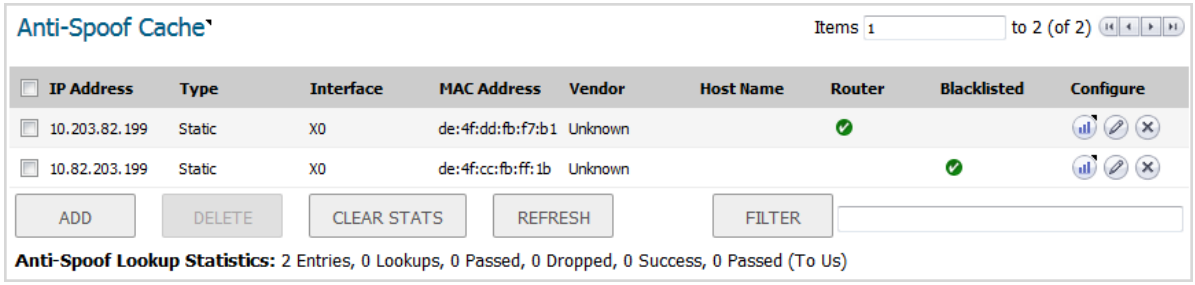
MAC-IP Anti-spoof Cache lists all MAC address-to-IP address bindings, which can include all the devices presently:

- Listed as “authorized” to access the network.
- Marked as a device that acts like a router with a network behind it.
- Marked as “blacklisted” (denied access) from the network.

Some packet types are bypassed even though the MAC-IP Anti-spoof feature is enabled:

- Non-IP packets.
- DHCP packets with source IP as 0.
- Packets from a VPN tunnel.
- Packets with invalid unicast IPs as their source IPs.
- Packets from interfaces where the Management status is not enabled under anti-spoof settings.

Anti-spoof lookup statistics are displayed at the bottom of the table.



IP Address	IP address of the device
Type	Type of entry: **Static** or **Lease**
Interface	Interface receiving incoming traffic
MAC Address	MAC address of the device
Vendor	Manufacturer of the device, if known
Host Name	Device’s host name, if known
Router	Device was designated as a possible router when configured
Blacklisted	Device was designated blacklisted when configured
Configure	Displays the **Statistics**, **Edit**, and **Delete** icons for each entry

To clear cache statistics on one or more devices:

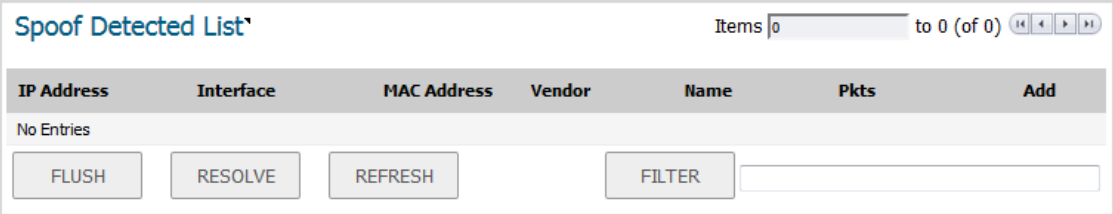
1. Navigate to **Network > MAC-IP Anti-spoof**.
2. Select one or more devices.
3. Click **CLEAR STATS**.

To view the most recent available cache information:

1. Navigate to **Network > MAC-IP Anti-spoof**.
2. Click **REFRESH** at the bottom of the **Anti-Spoof Cache** table.

Spoof Detected List

The **Spoof Detected List** displays devices that failed to pass the ingress anti-spoof cache check. Entries on this list can be added as a static anti-spoof entry in the **Anti-Spoof Cache** table.



IP Address	IP address of the device.
Interface	Interface receiving incoming traffic.
MAC Address	MAC address of the device.
Vendor	Manufacturer of the device, if known.
Name	Name of the device.
Pkts	Number of packets received.
Add	Displays the **Edit** icon.

To flush entries from the spoof-detected list:

1 Click **FLUSH**.

To resolve the name of each device using NetBios:

1 Click **RESOLVE**.

To view the most recent available cache information:

1 Click **REFRESH** at the bottom of the **Spoof Detected List** table.

Configuring MAC-IP Anti-spoof Protection

Topics:

- Displaying Traffic Statistics on page 509
- Editing MAC-IP Anti-spoof Settings for an IPv6 Interface on page 509
- Editing MAC-IP Anti-spoof Settings for an IPv4 Interface on page 510
- Adding Devices to Anti-Spoof Cache on page 512
- Deleting Anti-Spoof Cache Entries on page 512
- Filtering What Is Displayed on page 513
- Adding Static Entries from Spoof Detected List on page 513

Displaying Traffic Statistics

To display traffic statistics for an interface in the Settings or Anti-Spoof Cache tables:

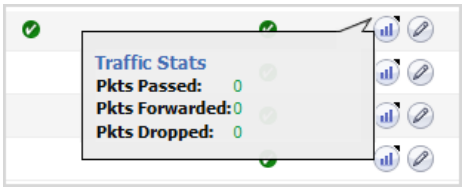
1 Navigate to **Network > MAC-IP Anti-spoof**.

2 To display traffic statistics for an interface in the **Settings** table, select an interface to display from **Settings for interface(s)**; the default is **All**.

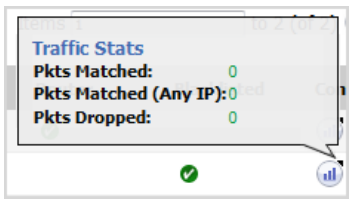
3 Mouse over the **Statistics** icon for the interface.

4 The **Traffic Stats** popup displays:

Settings table



Anti-Spoof Cache table

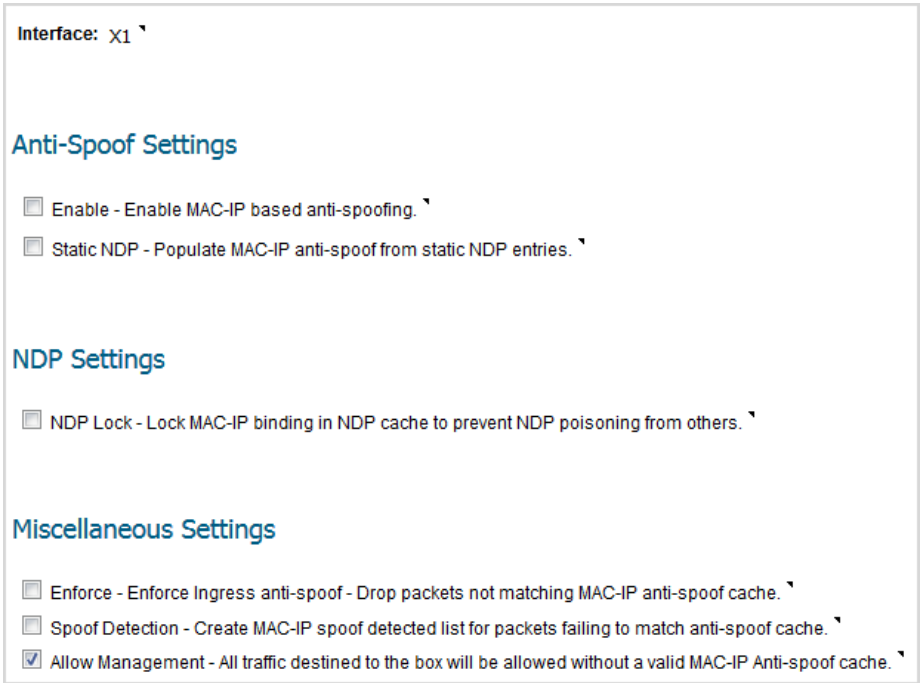


Editing MAC-IP Anti-spoof Settings for an IPv6 Interface

To configure MAC-IP Anti-spoof settings for a particular interface:

1 Navigate to **Network > MAC-IP Anti-spoof**.

2 In the **Settings for Interface(s)** table, click the **Configure** icon for the desired interface. The **Edit MAC-IP Anti-spoof Settings** dialog displays.



3 To enable MAC address and IP address traffic based on Anti-spoofing through this interface, select **Enable – Enable MAC-IP-based anti-spoofing** in the **Anti-Spoof Settings** section. This option is not selected by default.

4 To create for every static NDP entry a corresponding entry in the MAC-IP Anti-spoof table, select **Static NDP – Populate MAC-IP anti-spoof from static NDP entries**. This option is not selected by default.

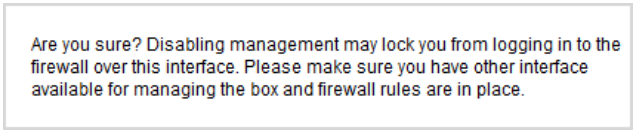
5 To add an NDP cache entry for every MAC-IP binding inside the anti-spoof cache, select **NDP Lock – Lock MAC-IP binding in NDP cache to prevent NDP poisoning from others** in the **NDP Settings** section. This option is not selected by default.

6 To enable a MAC-IP anti-spoof check for every transit packet, select **Enforce – Enforce Ingress anti-spoof - Drop packets not matching MAC-IP anti-spoof cache** in the **Miscellaneous Settings** section. This option is not selected by default.

7 To create a spoof-detected list for every device that fails a MAC-IP Anti-spoof cache check, select **Spoof Detection – Create MAC-IP spoof detected list for packets failing to match anti-spoof cache**. This option is not selected by default.

8 To allow all traffic destined to the Security Appliance, including without a valid MAC-IP Anti-spoof cache, select **Allow Management – All traffic destined to the box will be allowed without a valid MAC-IP Anti-spoof cache**. This option is selected by default.

CAUTION: **If you disable this option, you might be prevented from logging in to the SonicWall Security Appliance over this interface. Ensure you have other interfaces available for managing the Security Appliance and that proper rules and policies are in place. A warning message displaces if you disable the option:**



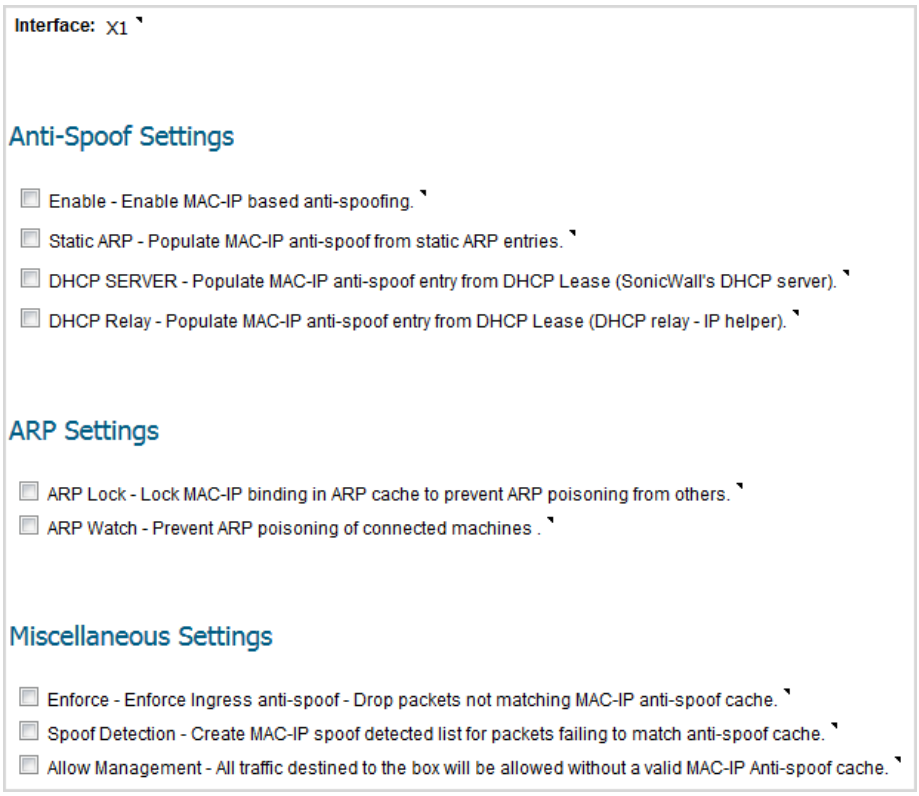
9 Click **OK**.

Editing MAC-IP Anti-spoof Settings for an IPv4 Interface

To configure MAC-IP Anti-spoof settings for a particular interface:

1 Navigate to **Network > MAC-IP Anti-spoof**.

2 In the **Settings for Interface(s)** table, click the **Configure** icon for the desired interface. The **Edit MAC-IP Anti-spoof Settings** dialog displays.



3 To enable MAC address and IP address traffic based on Anti-spoofing through this interface, select **Enable – Enable MAC-IP-based anti-spoofing** in the **Anti-Spoof Settings** section. This option is not selected by default.

4 To create for every static ARP entry a corresponding entry in the MAC-IP anti-spoof table, select **Static ARP – Populate MAC-IP anti-spoof from static ARP entries**. This option is not selected by default.

5 To create for every DHCP Lease allocated by a DHCP server a corresponding entry in the MAC-IP Anti-spoof table, select **DHCP SERVER – Populate MAC-IP anti-spoof entry from DHCP Lease (SonicWall's DHCP server)**. This option is not selected by default.

6 To create for every DHCP Lease allocated by a remote DHCP server a corresponding entry in the MAC-IP Anti-spoof table based on the DHCP relay configuration, select **DHCP Relay – Populate MAC-IP anti-spoof entry from DHCP Lease (DHCP relay - IP helper)**. This option is not selected by default.

7 To add an ARP cache entry for every MAC-IP binding inside the anti-spoof cache, select **ARP Lock – Lock MAC-IP binding in ARP cache to prevent ARP poisoning from others** in the **ARP Settings** section. This option is not selected by default.

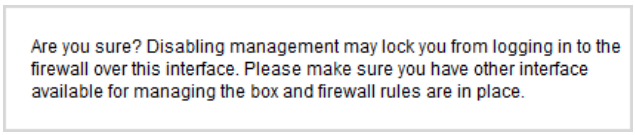
8 To prevent ARP poisoning of connected appliances and protect all client PCs from man-in-the-middle attacks, select **ARP Watch – Prevent ARP poisoning of connected machines**. This option is not selected by default.

9 To enable a MAC-IP anti-spoof check for every transit packet, select **Enforce – Enforce Ingress anti-spoof - Drop packets not matching MAC-IP anti-spoof cache** in the **Miscellaneous Settings** section. This option is not selected by default.

10 To create a spoof-detected list for every device that fails a MAC-IP Anti-spoof cache check, select **Spoof Detection – Create MAC-IP spoof detected list for packets failing to match anti-spoof cache**. This option is not selected by default.

11 To allow all traffic destined to the Security Appliance, including without a valid MAC-IP Anti-spoof cache, select **Allow Management – All traffic destined to the box is allowed without a valid MAC-IP Anti-spoof cache**. This option is selected by default.

CAUTION: If you disable this option, you could be prevented from logging in to the SonicWall Security Appliance over this interface. Ensure you have other interfaces available for managing the Security Appliance and that proper rules and policies are in place. A warning message displaces if you disable the option:



12 Click **OK**.

Adding Devices to Anti-Spoof Cache

To add a device to Anti-Spoof Cache:

1 Navigate to **Network > MAC-IP Anti-spoof**.

2 Click **ADD** below the **Anti-Spoof Cache** table. The **Add Static MAC-IP Anti-spoof** dialog displays.



3 From **Interface**, select the interface on which traffic from the device arrives.

4 In the **IP Address** field, type in the IP address of the device.

5 In the **MAC Address** field, type in the MAC address of the device.

6 To designate the device as a router that might have a network behind it, select **A Router**. This option is selected by default.

7 To put this device on the blacklist and block traffic from it, select **A blacklisted device**. This option is not selected by default.

Blacklisting the device causes packets to be blocked from this device regardless of its IP address.

8 Click **OK**.

Deleting Anti-Spoof Cache Entries

To delete a single static anti-spoof cache entry:

1 Navigate to **Network > MAC-IP Anti-spoof**.

2 Click the **Delete** icon of the entry.

To delete one or more static anti-spoof cache entries:

1 Navigate to **Network > MAC-IP Anti-spoof**.

2 Select the entries to be deleted. **DELETE** becomes active.

3 Click **DELETE**.

To delete all static anti-spoof cache entries

1 Navigate to **Network > MAC-IP Anti-spoof**.

2 Select the checkbox in the Anti-Spoof Cache table header. **DELETE** becomes active.

3 Click **DELETE**.

Filtering What Is Displayed

You can display only a specific device(s) in the **Anti-Spoof Cache** and **Spoof Detected List** tables by using the **Filter** function.

To filter the table display:

1 Navigate to **Network > MAC-IP Anti-spoof**.

2 In the Filter field below the table to be filtered, specify either the device's IP address, interface, MAC address, host name, or name. The field must be filled using the appropriate syntax for operators shown in Filter operator syntax options

Filter operator syntax options

Operator	Syntax Options
Value with a type	• Ip=1.1.1.1 or ip=1.1.1.0/24 • Mac=00:01:02:03:04:05 • Iface=x1
String	• X1 • 00:01 • Tst-mc • 1.1.
AND	• Ip=1.1.1.1;iface=x1 • Ip=1.1.1.0/24;iface=x1;just-string
OR	• Ip=1.1.1.1,2.2.2.2,3.3.3.0/24 • Iface=x1,x2,x3
Negative	• !ip=1.1.1.1;!just-string • !iface=x1,x2
Mixed	• Ip=1.1.1.1,2.2.2.2;mac=00:01:02:03:04:05; just-string;!iface=x1,x2

Adding Static Entries from Spoof Detected List

To add a static entry from the Spoof Detected List:

1 Navigate to **Network > MAC-IP Anti-spoof**.

2 In the **Spoof Detected List** table, click the **Edit** icon under the **Add** column for the desired device. An alert message displays, asking if you wish to add this static entry.

3 Click **OK**.

26

Setting Up the DHCP Server

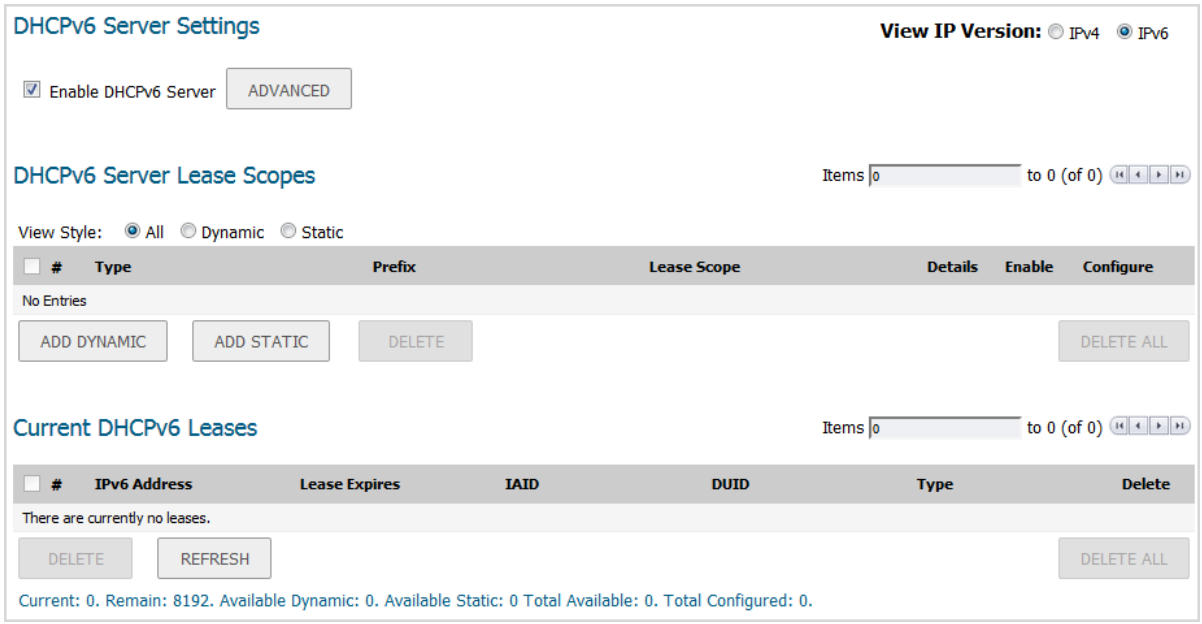
Topics:

- **Network > DHCP Server** on page 514
  - DHCP Server Options Feature on page 516
  - Multiple DHCP Scopes per Interface on page 517
  - About DHCP Server Persistence on page 519
  - Configuring the DHCP Server on page 519
  - DHCP Server Lease Scopes on page 521
  - Current DHCP Leases on page 521
  - DHCPv6 Relay on page 523
- Configuring Advanced Options on page 523
  - Configuring Advanced Options on page 523
  - Configuring DHCP Server for Dynamic Ranges on page 528
  - Configuring Static DHCP Entries on page 534
  - Configuring DHCP Generic Options for DHCP Lease Scopes on page 536
  - RFC-Defined DHCP Option Numbers on page 536
  - DHCP and IPv6 on page 540

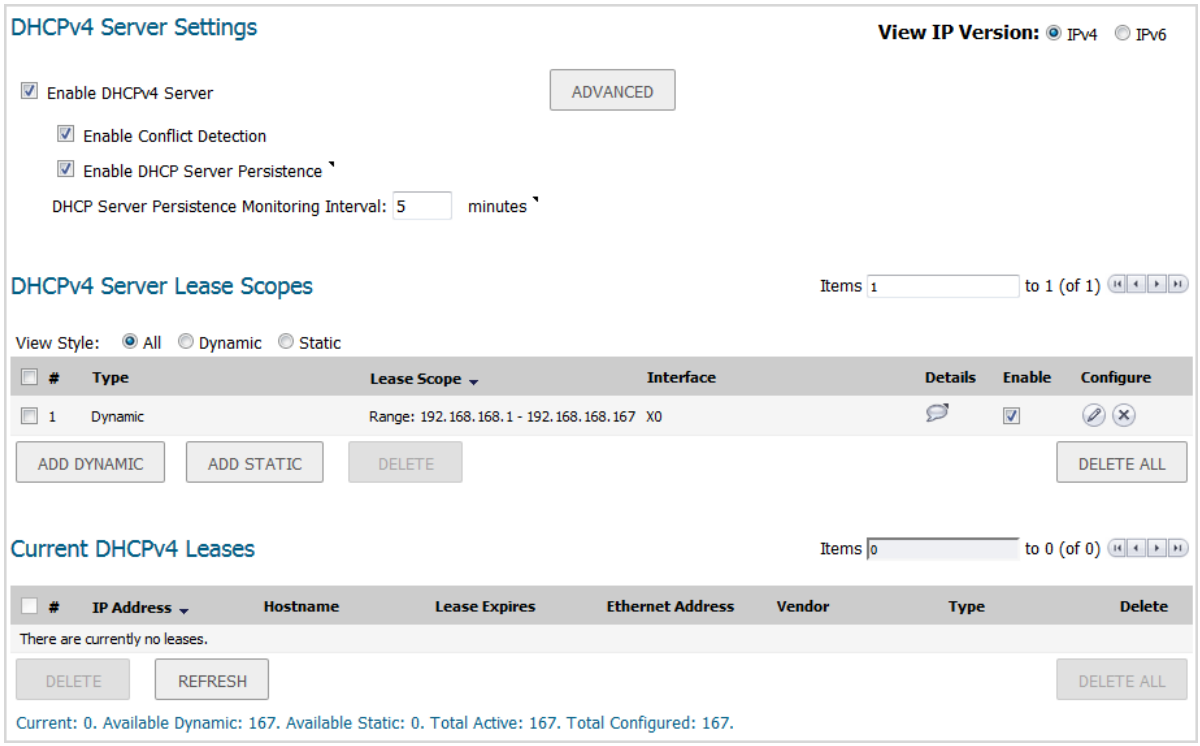
Network > DHCP Server

There are only minor differences between the IPv6 (IPv6 **Network > DHCP Server**) and IPv4 versions (IPv4 **Network > DHCP Server**) of **Network > DHCP Server**. Differences are noted within procedures.

IPv6 Network > DHCP Server



IPv4 Network > DHCP Server



The SonicWall Security Appliance includes a DHCP (Dynamic Host Configuration Protocol) server to distribute IP addresses, subnet masks, gateway addresses, and DNS server addresses to your network clients. **Network > DHCP Server** includes settings for configuring the Security Appliance's DHCP server.

IMPORTANT: You can use the Security Appliance's DHCP server or use existing DHCP servers on your network. If your network uses its own DHCP servers, make sure **Enable DHCP Server** is cleared.

The number of address ranges and IP addresses that the firewall's DHCP server can assign depends on the model, operating system, and licenses of the Security Appliance. Maximum DHCP Leases shows maximum allowed DHCP leases for SonicWall Security Appliances.

Maximum DHCP Leases

Platform	Maximum DHCP Leases	Platform	Maximum DHCP Leases	Platform	Maximum DHCP Leases
SM 9600	16384	NSA 6650	16384	TZ600	4096
SM 9400	16384	NSA 6600	16384	TZ500/TZ500 W	4096
SM 9200	16384	NSA 5650	8192	TZ400/TZ400 W	4096
		NSA 5600	8192	TX350/TZ350 W	4096
		NSA 4650	8192	TZ300/TZ300 W	4096
		NSA 4600	8192		
		NSA 3650	4096	SOHO 250/SOHO 250 W	4096
		NSA 3600	4096	SOHO W	4096
		NSA 2650	4096		
		NSA 2600	4096		

Topics:

- DHCP Server Options Feature on page 516
- Multiple DHCP Scopes per Interface on page 517
- About DHCP Server Persistence on page 519
- Configuring the DHCP Server on page 519
- DHCP Server Lease Scopes on page 521
- Current DHCP Leases on page 521
- Configuring Advanced Options on page 523
- Configuring DHCP Server for Dynamic Ranges on page 528
- Configuring Static DHCP Entries on page 534
- Configuring DHCP Generic Options for DHCP Lease Scopes on page 536
- RFC-Defined DHCP Option Numbers on page 536
- DHCP and IPv6 on page 540

DHCP Server Options Feature

The SonicWall DHCP server options feature provides support for DHCP options, also known as vendor extensions, as defined primarily in RFCs 2131 and 2132. DHCP options allow users to specify additional DHCP parameters in the form of predefined, vendor-specific information that is stored in the options field of a DHCP message. When the DHCP message is sent to clients on the network, it provides vendor-specific configuration and service information. The section RFC-Defined DHCP Option Numbers on page 536 lists DHCP options by RFC-assigned option number.

Topics:

- Benefits on page 517

- How the DHCP Server Options Feature Works on page 517
- Supported Standards on page 517

Benefits

The SonicWall DHCP server options feature provides a simple interface for selecting DHCP options by number or name, making the DHCP configuration process quick, easy, and compliant with RFC-defined DHCP standards.

How the DHCP Server Options Feature Works

The DHCP server options feature allows definition of DHCP options using a drop-down menu based on RFC-defined option numbers, allowing administrators to easily create DHCP objects and object groups, and configure DHCP generic options for dynamic and static DHCP lease scopes. After being defined, the DHCP option is included in the options field of the DHCP message, which is then passed to DHCP clients on the network, describing the network configuration and service(s) available.

Supported Standards

The DHCP server options feature supports the following standards:

- RFC 2131 - Dynamic Host Configuration Protocol
- RFC 2132 - DHCP Options and BOOTP Vendor Extensions

Multiple DHCP Scopes per Interface

Topics:

- What are Multiple DHCP Scopes per Interface? on page 517
- Benefits of Multiple DHCP Scopes on page 517
- How Do Multiple DHCP Scopes per Interface Work? on page 518

What are Multiple DHCP Scopes per Interface?

Often, DHCP clients and server(s) reside on the same IP network or subnet, but sometimes DHCP clients and their associated DHCP server(s) do not reside on the same subnet. The Multiple DHCP Scopes per Interface feature allows one DHCP server to manage different scopes for clients spanning multiple subnets.

Benefits of Multiple DHCP Scopes

Efficiency	A single DHCP server can provide IP addresses for clients spanning multiple subnets.
Compatible with DHCP over VPN	The processing of relayed DHCP messages is handled uniformly, regardless of whether it comes from a VPN tunnel or a DHCP relay agent.
Multiple Scopes for Site-to-Site VPN	When using an internal DHCP server, a remote subnet could be configured using scope ranges that differ from the LAN/DMZ subnet. The scope range for the remote subnet is decided by the “Relay IP Address” set in the remote gateway.

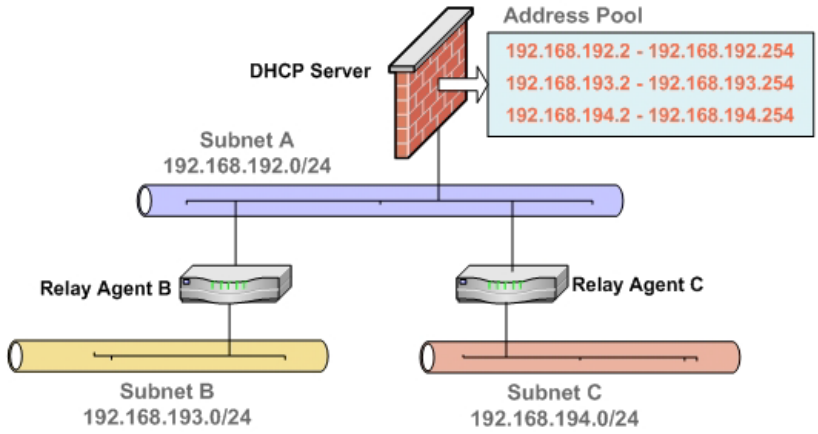
Multiple Scopes for Group VPN	When using an internal DHCP server, a SonicWall GVC client could be configured using scope ranges that differ from the LAN/DMZ subnet. The scope range for GVC client is decided by the Relay IP Address (Optional) option set in the central gateway.
Compatible with Conflict Detection	Currently, DHCP server performs server-side conflict detection when this feature is enabled. The advantage of server-side conflict detection is that it detects conflicts even when the DHCP client does not run client-side conflict detection. However, if there are a lot of DHCP clients on the network, server-side conflict detection can result in longer waits for a full IP address allocation to complete. Conflict Detection (and Network Pre-Discovery) are not performed for an IP address which belongs to a “relayed” subnet scope. The DHCP server only performs a conflict detection ICMP check for a subnet range attached to its interface.

How Do Multiple DHCP Scopes per Interface Work?

Normally, a DHCP client initiates an address allocating procedure by sending a Broadcast DHCP Discovery message. As most routes do not forward broadcast packets, this method requires DHCP clients and server(s) to reside on the same IP network or subnet.

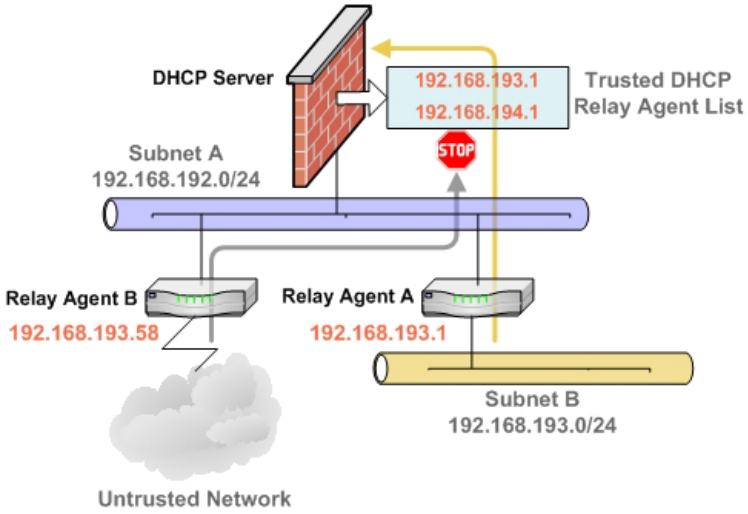
When DHCP clients and their associated DHCP server are not on the same subnet, some type of third-party agent (such as BOOTP relay agent, IP Helper) is required to transfer DHCP messages between clients and server; see Multiple subnets sharing one DHCP server. The DHCP relay agent populates the giaddr field with its ingress interface IP address and then forwards it to the configured DHCP server. When the DHCP server receives the message, it examines the giaddr field to determine if it has a DHCP scope that could be used to supply an IP address lease to the client.

Multiple subnets sharing one DHCP server



The Multiple DHCP Scopes per Interface feature provides security enhancements to protect against potential vulnerabilities inherent in allowing wider access to the DHCP server. The **DHCP Advanced Setting** dialog provides security with a tab for Trusted Agents for specifying trusted DHCP relay agents; see Trusted DHCP relay agents. The DHCP server discards any messages relayed by agents which are not in the list.

Trusted DHCP relay agents



About DHCP Server Persistence

DHCP server persistence is the ability of the Security Appliance save DHCP lease information and to provide the client with a predictable IP address that does not conflict with another use on the network, even after a client reboot.

DHCP server persistence works by storing DHCP lease information periodically to flash memory. This ensures that users have predicable IP addresses and minimizes the risk of IP addressing conflicts after a reboot.

DHCP server persistence provides a seamless experience when a user reboots a workstation. The DHCP lease information is saved, and the user retains the same workstation IP address. When a firewall is restarted, usually due to maintenance or an upgrade, DHCP server persistence provides these benefits:

- IP address uniqueness: Lease information is stored in flash memory, so the risk of assigning the same IP address to multiple users is nullified.
- Ease of use: By saving the lease information in the flash memory, the user's connections are automatically restored.

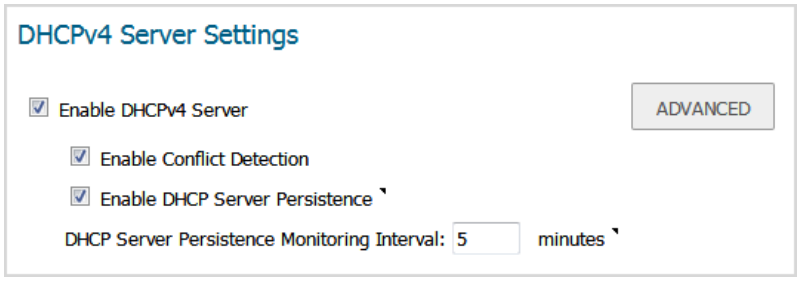
Configuring the DHCP Server

To use the SonicWall Security Appliance's DHCP server:

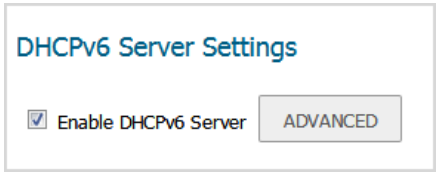
1 Navigate to **MANAGE | System Setup | Network > DHCP Server**.

2 Choose which IP version to use from **View IP Version**:

- **IPv4**



- **IPv6**



3 To distribute IP addresses, subnet masks, gateway addresses, and DNS server addresses to your network clients, select **Enable DHCPv4/6 Server**. This option is selected by default. **ADVANCED** and, for IPv4, the server settings options become available.

4 For configuring DHCPv6, go to Step 7.

5 To turn on automatic DHCP scope conflict detection on each zone when another DHCP server is present, select **Enable Conflict Detection**. This option is selected by default.

Currently, DHCP server performs server-side conflict detection when this feature is enabled. The advantage of server-side conflict detection is that it detects conflicts even when the DHCP client does not run client-side conflict detection. However, if there are a lot of DHCP clients on the network, server-side conflict detection can result in longer waits for a full IP address allocation to complete.

NOTE: Conflict detection is not performed for an IP address that belongs to a "relayed" subnet scope. The DHCP server only performs a conflict detection ICMP check for a subnet range attached to its interface.

6 To allow the current state of the DHCP leases in the network to be periodically written to Flash, select **Enable DHCP Server Persistence**. At reboot, the system restores the previous DHCP server network DHCP allocation knowledge based on the IP. Lease times stored in Flash. This option is selected by default. When this option is selected, the **DHCP Server Persistence Monitoring Interval** option is available.

- To control how often changes in the network are examined and, if necessary, written to Flash, enter the time, in minutes, in **DHCP Server Persistence Monitoring Interval**. The default is **5** minutes, the minimum is five minutes, and the maximum is 1440 minutes (24 hours).

7 To configure **Option Objects**, **Option Groups**, and **Trusted Agents**, click **ADVANCED**. For detailed information on configuring these features, see Configuring Advanced Options on page 523.

8 Click **ACCEPT**.

Topics:

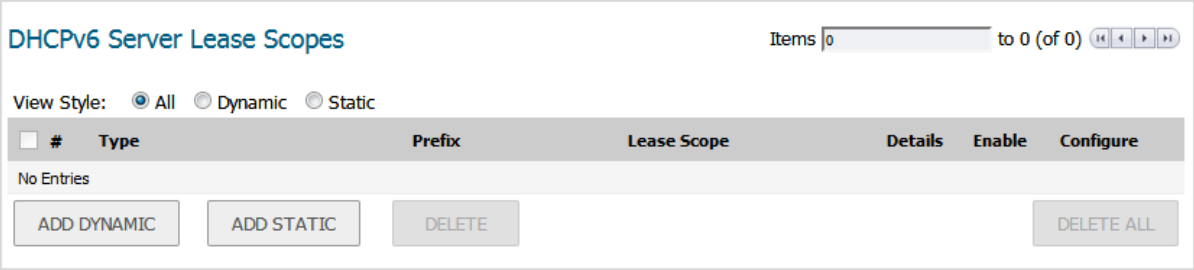
- Configuring the DHCP Server for DNS Proxy on page 521
- Current DHCPv4 Leases on page 522

Configuring the DHCP Server for DNS Proxy

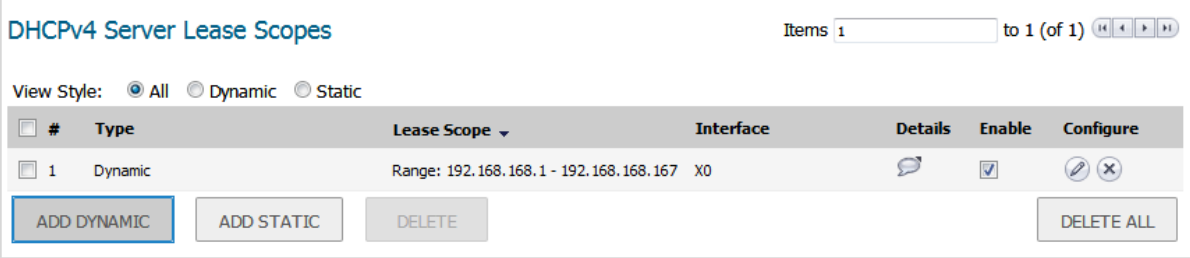
When DNS proxy is enabled on an interface, the device needs to push the interface IP as DNS server address to clients, so the you need to configure the DHCP server manually; use the interface address as the **DNS Server 1** address in the DHCP server settings on the **DNS/WINS** tab. The I**nterface Pre-populate** checkbox in the DHCP page makes this easy to configure; if the selected interface has enabled DNS proxy, the DNS server IP is auto-added into the **DNS/WINS** page.

DHCP Server Lease Scopes

DHCPv6 Server Lease Scopes



DHCPv4 Server Lease Scopes



The **DHCP Server Lease Scopes** table displays the currently configured DHCP IP ranges:

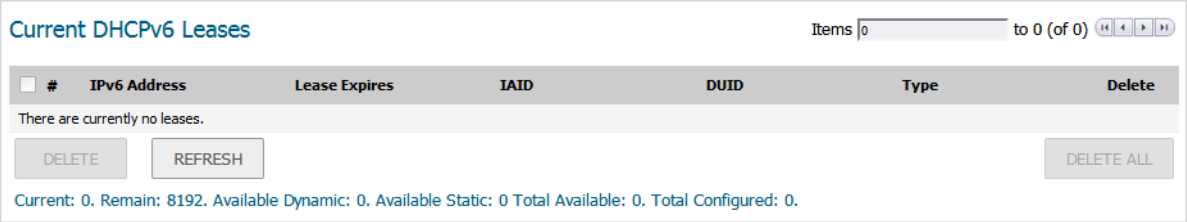
Type	**Dynamic** or **Static**.
Prefix	IPv6 only.
Lease Scope	The IP address range, for example, `172.16.31.2-172.16.31.254`.
Interface	IPv4 only. The Interface the range is assigned to.
Details	Detailed information about the lease, displayed as a tool tip when you hover the mouse pointer over the **Comment** icon.
Enable	Select the checkbox to enable the DHCP range. Clear it to disable the range.
Configure	Contains the **Configure** and **Delete** icons for the table entry.

Current DHCP Leases

Topics:

- Current DHCPv6 Leases on page 522
- Current DHCPv4 Leases on page 522

Current DHCPv6 Leases



The current DHCP lease information is displayed in the **Current DHCP Leases** table. Each binding entry displays the:

- **IPv6 Address**
- **Lease Expires**
- **IAID**
- **DUID**
- **Type** of binding (**Dynamic**, **Dynamic BOOTP**, or **Static BOOTP**)
- **Delete** icon

To delete a binding, which frees the IP address on the DHCP server:

1. Click the **Delete** icon next for the entry. For example, use the **Delete** icon to remove a host when it has been removed from the network and you need to reuse its IP address.
2. Click **Accept**.

Current DHCPv4 Leases



The current DHCP lease information is displayed in the **Current DHCP Leases** table. Each binding entry displays the:

- **IP Address**
- **Hostname**
- **Lease Expires**
- **Ethernet Address**
- **Vendor**
- **Type** of binding (**Dynamic**, **Dynamic BOOTP**, or **Static BOOTP**)
- **Delete** icon

To delete a binding, which frees the IP address on the DHCP server:

1. Click the **Delete** icon next for the entry. For example, use the **Delete** icon to remove a host when it has been removed from the network and you need to reuse its IP address.
2. Click **Accept**.

DHCPv6 Relay

SonicOS supports DHCPv6 Relay. For information about DHCPv6 relay in SonicOS, see DHCPv6 Relay on page 543.

Configuring Advanced Options

NOTE: Configuring DHCP server options is essentially the same for both IPv4 and IPv6. Exceptions are noted in the procedures.

Topics:

- Configuring DHCP Option Objects on page 523
- Configuring DHCP Option Groups on page 525
- Configuring a Trusted DHCP Relay Agent Address Group (IPv4 Only) on page 527
- Enabling Trusted DHCP Relay Agents on page 527

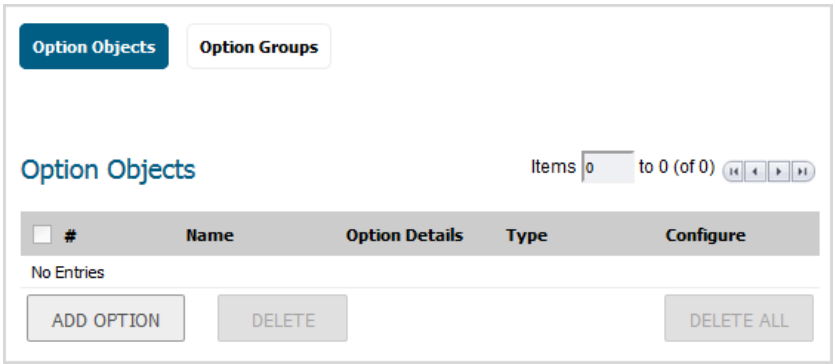
The RFC-Defined DHCP Option Numbers on page 536 provides a list of DHCP options by RFC-assigned option number.

Configuring DHCP Option Objects

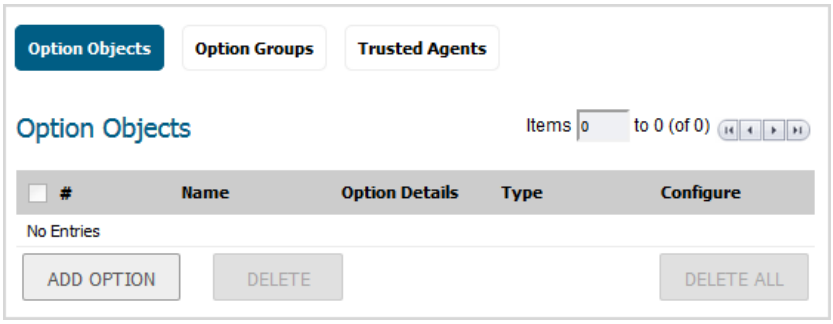
To configure DHCP option objects:

1. Navigate to **MANAGE | System Setup | Network > DHCP Server**.
2. Under **DHCPv4/6 Server Settings**, click **Advanced**. The **DHCP Advanced Settings** dialog displays. The dialogs for IPv4 and IPv6 are slightly different; see IPv6 DHCP Advanced Settings and IPv4 DHCP Advanced Settings.

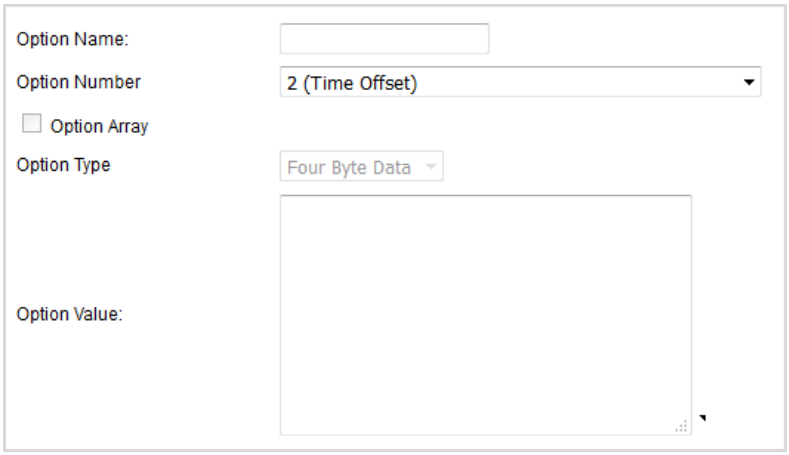
IPv6 DHCP Advanced Settings



IPv4 DHCP Advanced Settings



3 Click **Add Option**. The **Add DHCP Option Objects** dialog displays.



4 Type a name for the option in the **Option Name** field.

5 From **Option Number**, select the option number that corresponds to your DHCP option. For a list of option numbers, names, and descriptions, refer to RFC-Defined DHCP Option Numbers on page 536.

NOTE: Available options differ depending on whether you are configuring an IPv4 or IPv6 option.

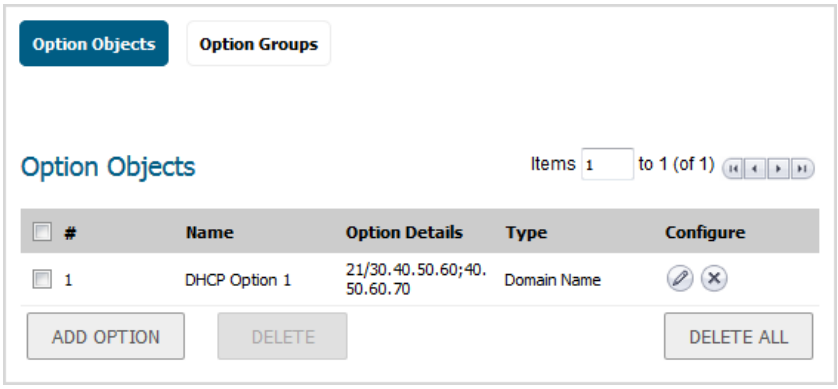
6 If:

- Only one option type is available, for example, for **Option Number 2 (Time Offset)**, **Option Array** is dimmed. Go to Step 7.
- There are multiple option types available, for example, for **77 (User Class Information)**, **Option Type** becomes available and lists allowable types of the option, such as **IP Address**, **Two-Byte Data**, **String**, or **Boolean**. Select the option type.

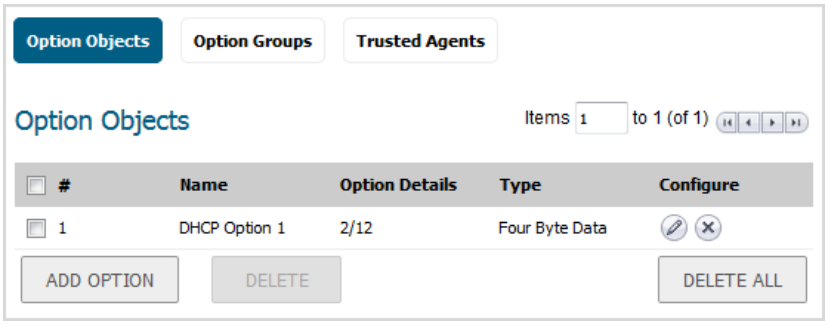
7 Type the option value, for example, an IP address, in the **Option Value** field. If **Option Array** is checked, multiple values may be entered, separated by a semi-colon (;).

8 Click **OK**. The object displays in the **Option Objects** table (see DHCPv6 Option Objects table and DHCPv4 Option Objects table).

DHCPv6 Option Objects table



DHCPv4 Option Objects table



Configuring DHCP Option Groups

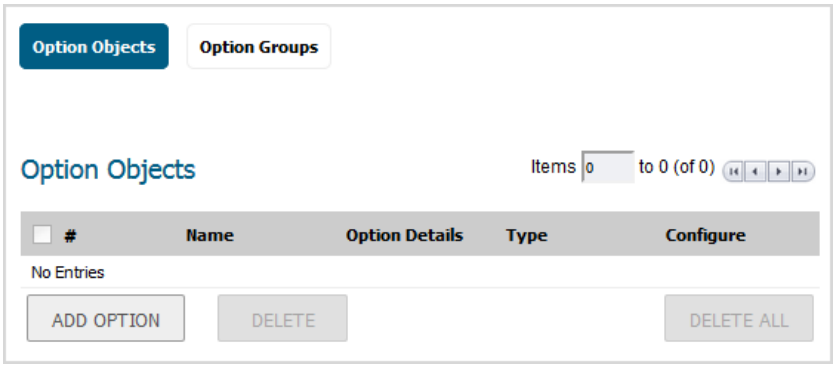
To configure DHCP option groups:

1 Navigate to **MANAGE | System Setup | Network > DHCP Server**.

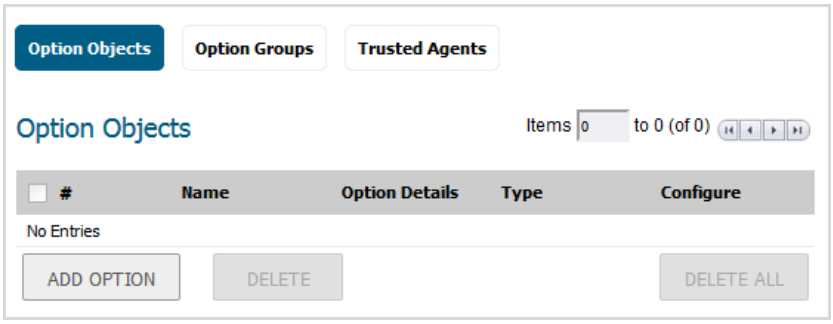
2 Under **DHCPv4/6 Server Settings**, click **Advanced**. The **DHCP Advanced Settings** dialog displays.

NOTE: Available options differ depending on whether you are configuring an IPv4 or IPv6 option (see IPv6 DHCP Advanced Settings or IPv4 DHCP Advanced Settings).

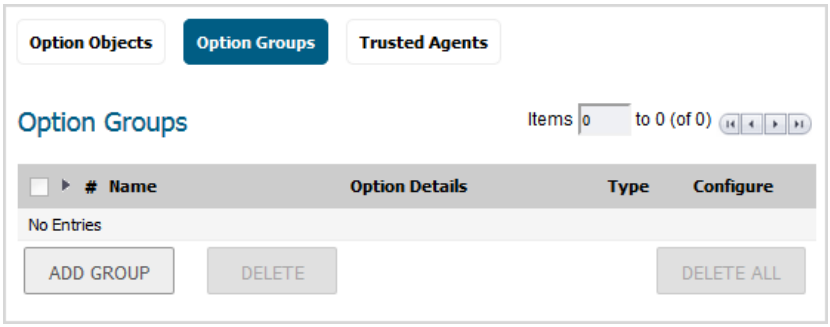
IPv6 DHCP Advanced Settings



IPv4 DHCP Advanced Settings



3 Click **Option Groups**.



4 Click **Add Group**. The **Add DHCP/v6 Option Group** dialog displays.

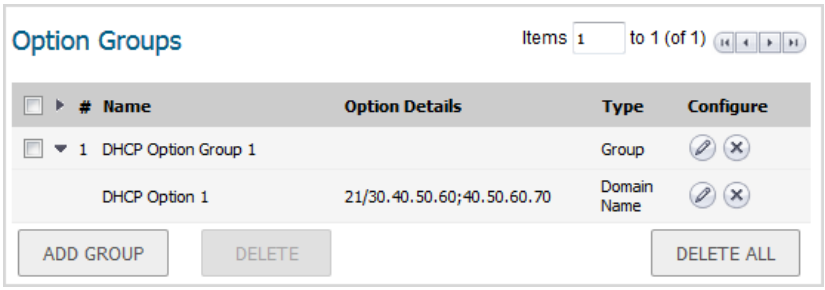


5 Enter a name for the group in the **Name** field.

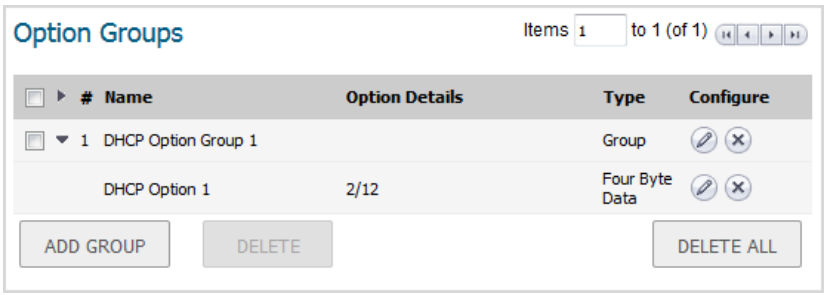
6 Select an option object from the left column and click **Right Arrow** to add the option object to the group. To select multiple option objects at the same time, hold the **Ctrl** key while selecting the option objects.

7 Click **OK**. The group displays in the **Option Groups** table.

DHCPv6 Option Groups table



DHCPv4 Option Groups table



Configuring a Trusted DHCP Relay Agent Address Group (IPv4 Only)

To configure the **Default Trusted Relay Agent List** Address Group, you must first configure an Address Object for each trusted relay agent, then add these Address Objects to the **Default Trusted Relay Agent List** Address Group or to a custom Address Group.

Address Objects and Address Groups are configured on **MANAGE | Policies | Objects > Address Objects**. For information on how to configure Address Objects and Address Groups, see *SonicOS 6.5 Policies*.

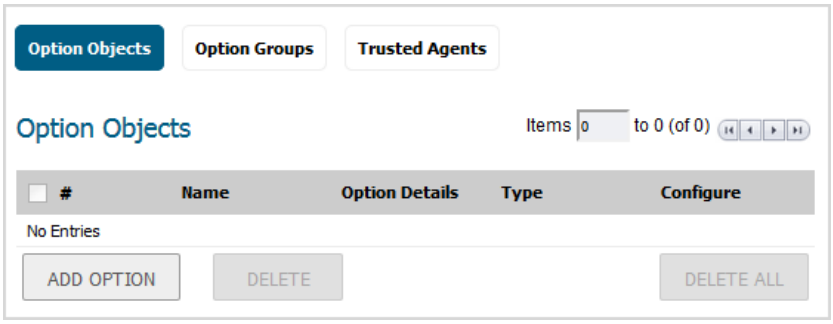
Enabling Trusted DHCP Relay Agents

In the **DHCP Advanced Settings** dialog, you can enable the **Trusted Relay Agent List** option using the **Default Trusted Relay Agent List** Address Group or create another Address Group using existing Address Objects.

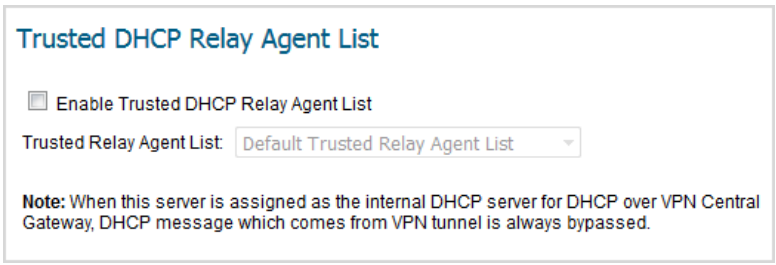
NOTE: When a server is assigned as the internal DHCP server for DHCP over VPN Central Gateway, DHCP messages that come from the VPN tunnel are always bypassed.

To enable the Trusted Relay Agent List option and select the desired Address Group:

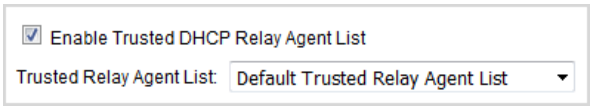
1. Navigate to **MANAGE | System Setup | Network > DHCP Server**.
2. Under **DHCPv4 Settings**, click **ADVANCED**. The **DHCP Advanced Settings** dialog displays.



3 Click **Trusted Agents**.



4 Select **Enable Trusted DHCP Relay Agent List**. This option is not selected by default. **Trusted Relay Agent List** becomes available.



5 Select the Address Group from **Default Trusted Relay Agent List**. This option includes all existing address groups as well as the **Create new Address Object Group** option.

NOTE: To create a custom Address Group for this option, select **Create new Address Object Group**. The **Add Address Object Group** dialog displays. For information on how to configure Address Groups, see *SonicOS Policies*.

6 Click **OK** to enable the **Trusted Relay Agent List** option with the selected Address Group.

Configuring DHCP Server for Dynamic Ranges

Because SonicOS allows multiple DHCP scopes per interface, there is no requirement that the subnet range is attached to the interface when configuring DHCP scopes.

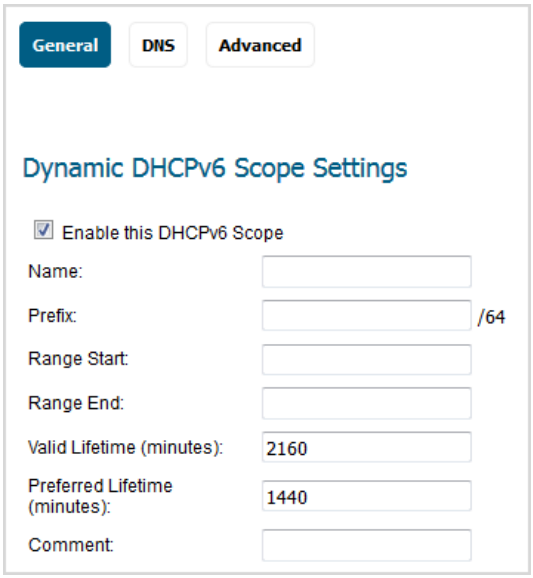
To configure DHCP server for dynamic IP address ranges:

1 Navigate to **MANAGE | System Setup | Network > DHCP Server**.

2 Under **DHCPv4/6 Server Lease Scopes** table, click **ADD DYNAMIC**. For:

- IPv6, the **Add DHCPv6 Dynamic Scope** dialog displays. Go to Add DHCPv6 Dynamic Scope on page 529.
- IPv4, the **Dynamic Ranges Configuration** dialog displays. Go to Dynamic Range Configuration on page 530.

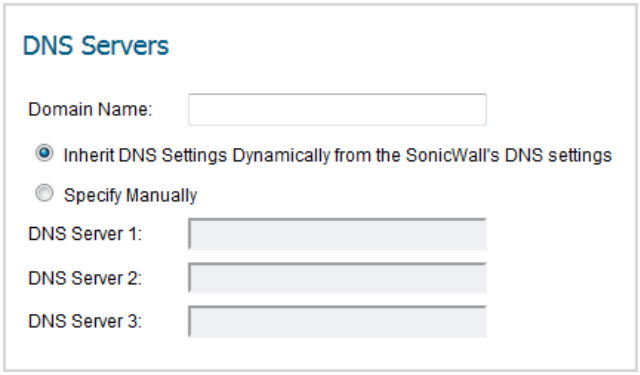
Add DHCPv6 Dynamic Scope



To add a dynamic scope:

1. To enable this scope, ensure **Enable this DHCP Scope** is selected. This option is selected by default.
2. Enter a name for the scope in the **Name** field.
3. Enter the prefix the scope uses to distribute IPv6 addresses in the **Prefix** field.
4. Enter the range start and range end in the **Range Start** and **Range End** fields, respectively. Both addresses must fall within the scope of the prefix.
5. Enter the valid lifetime of an IPv6 address leased by the scope, in minutes, in the **Valid Lifetime** field. The minimum is 0, the maximum is 9999, and the default is **2160**.
6. Enter the preferred lifetime of an IPv6 address leased by the scope, in minutes, in the **Preferred Lifetime** field. The minimum is 0, the maximum is 9999, and the default is **1440**.
7. Optionally, enter a comment in the **Comment** field.
8. Click **DNS**.

DNS



To add a DNS server:

1 Enter a domain name in the **Domain Name** field.

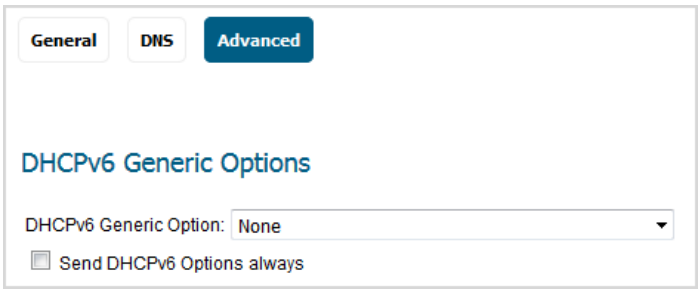
2 Choose whether to:

- **Inherit DNS Settings Dynamically from the SonicWall's DNS settings**; go to Step 4.
- **Specify Manually**. The **DNS Server 1/2/3** fields become available.

3 Enter the IP address(es) of the DNS server(s) in the respective **DNS Server 1/2/3** field(s).

4 Click **Advanced**.

Advanced



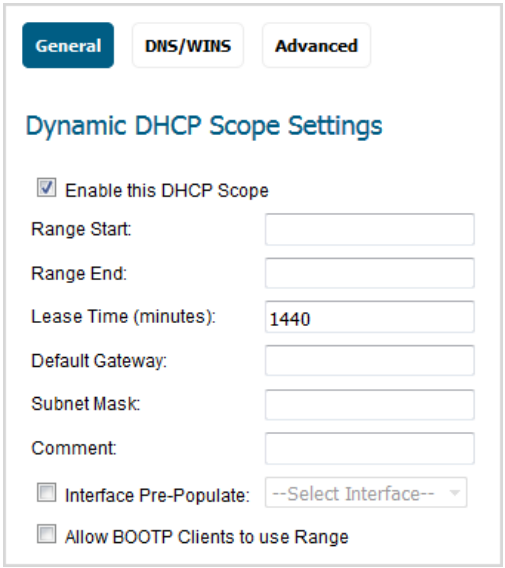
To configure generic DHCP options:

1 Select a DHCP Option Object or Group from **DHCPv6 Generic Option**. The default is **None**. To configure a new DHCPv6 Option or Group, see Configuring DHCP Option Objects on page 523 and/or Configuring DHCP Option Groups on page 525.

2 To send all configured DHCPv6 options for this scope regardless of the Option Request Option contained in the message from the DHCPv6 client, select **Send DHCPv6 Options always**. This option is not selected by default.

3 Click **OK**.

Dynamic Range Configuration



To configure a dynamic range:

1. To enable this scope, ensure **Enable this DHCP Scope** is selected. This option is selected by default.
2. To populate the **Range Start**, **Range End**, **Default Gateway**, and **Subnet Mask** fields:
   a. With default values for a certain interface:
      1) Select **Interface Pre-Populate** near the bottom of the dialog. The selections become available. This option is not selected by default.
      2) Select the interface. The populated IP addresses are in the same private subnet as the selected interface.

      IMPORTANT: To select an interface from **Interface Pre-Populate**, the interface must first be fully configured and it must be either:
      - Of the zone type LAN, WLAN, or DMZ.
      - A VLAN sub-interface.

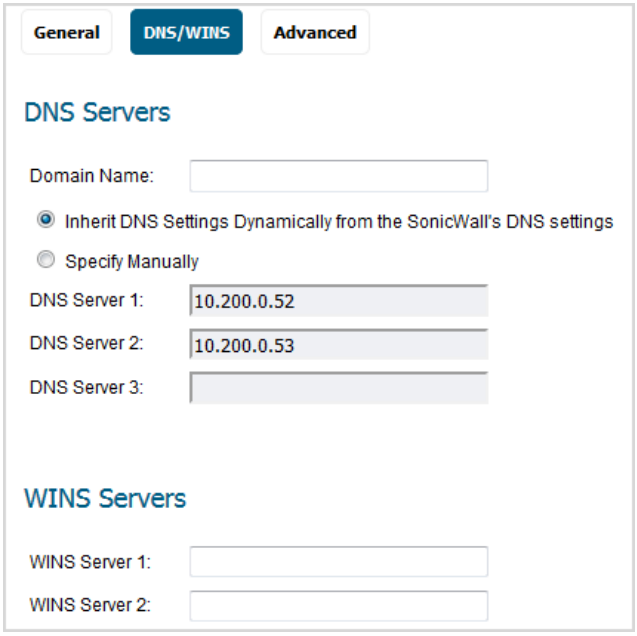
      3) Go to Step 3.

   b. Manually:
      1) Type in your own IP address range.
      2) Enter the number of minutes an IP address is leased by the scope before it is issued another IP address in the **Lease Time (minutes)** field. The minimum is 0, the maximum is 9999, and **1440** minutes (24 hours) is the default.
      3) Enter the IP address of the gateway into the **Default Gateway** field.
      4) Enter the gateway subnet mask into the **Subnet Mask** field.
3. Optionally, enter a comment in the **Comment** field.
4. Select **Allow BOOTP Clients to use Range** if you have BOOTP Clients on your network. This option is not selected by default.

   BOOTP stands for bootstrap protocol, which is a TCP/IP protocol and service that allows diskless workstations to obtain their IP address, other TCP/IP configuration information, and their boot image file from a BOOTP server.

5 Click **DNS/WINS** to continue configuring the DHCP Server feature.

DNS/WINS



To configure DNS/WINS servers:

1 If you have a domain name for the DNS server, enter it in the **Domain Name** field.

2 Choose whether to:

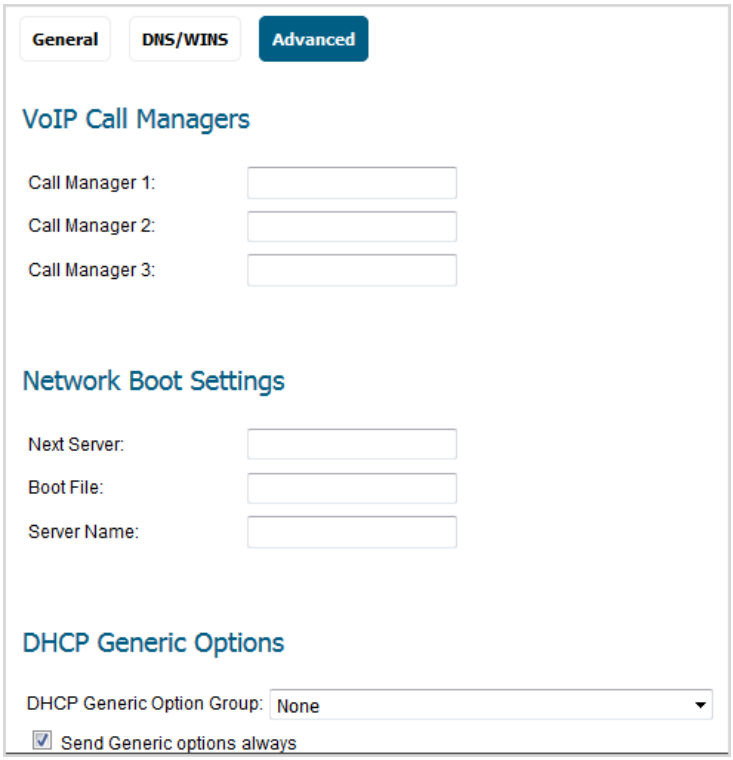
- **Inherit DNS Settings Dynamically from the SonicWall's DNS settings**; go to Step 4.
- **Specify Manually**. The **DNS Server 1/2/3** fields become available.

3 Enter the IP address(es) of the DNS server(s) in the respective **DNS Server 1/2/3** field(s).

4 If you have WINS running on your network, type the WINS server IP address in the **WINS Server 1** field. You can add an additional WINS server.

5 Click **Advanced**. The **Advanced** options allow you to configure DHCP server to send Cisco Call Manager information to VoIP clients on the network.

Advanced



To configure advanced settings:

1 Under **VoIP Call Managers**, enter the IP address or FQDN of your VoIP Call Manager in the **Call Manager 1** field. You can add two additional VoIP Call Manager addresses.

2 Under **Network Boot Settings**, in the **Next Server** field, enter the IP address of the PXE boot server (TFTP server) that a PXE client uses during the next stage of the boot process.

> **IMPORTANT:** The fields under **Network Boot Settings** are used in a Pre-boot Execution Environment (PXE), in which the client boots up using files obtained over a network interface. The PXE client obtains the IP address and name of the PXE boot server, and the boot file name, from the DHCP server.
>
> When using these options, select **PXE** under **DHCP Generic Options**.

3 In the **Boot File** field, enter the name of the boot file that the PXE client can get over TFTP from the PXE boot server.

4 In the **Server Name** field, enter the DNS host name of the PXE boot server (TFTP server).

5 For information on configuring DHCP Generic Options see Configuring DHCP Generic Options for DHCP Lease Scopes on page 536.

6 Click **OK**.

7 Click **Accept** for the settings to take effect on the firewall.

For more information on VoIP support features on the SonicWall Security Appliance, see About VoIP on page 765.

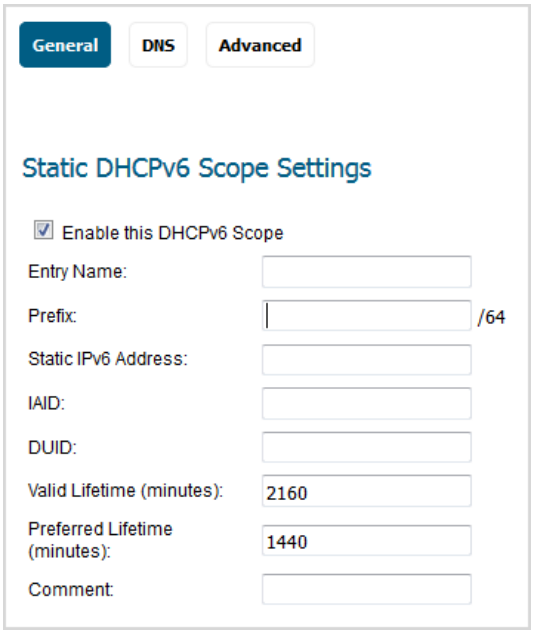
Configuring Static DHCP Entries

Static entries are IP addresses assigned to servers requiring permanent IP settings. Because SonicOS allows multiple DHCP scopes per interface, there is no requirement that the subnet range is attached to the interface when configuring DHCP scopes.

To configure static entries:

1. Navigate to **MANAGE | System Setup | Network > DHCP Server**.
2. Under **DHCPv4/6 Server Lease Scopes** table, click **ADD STATIC**. For:
   - IPv6, the **Add DHCPv6 Static Scope** dialog displays. Go to Add DHCPv6 Static Scope on page 534.
   - IPv4, the **Static Entry Configuration** dialog displays. Go to Static Entry Configuration on page 535.

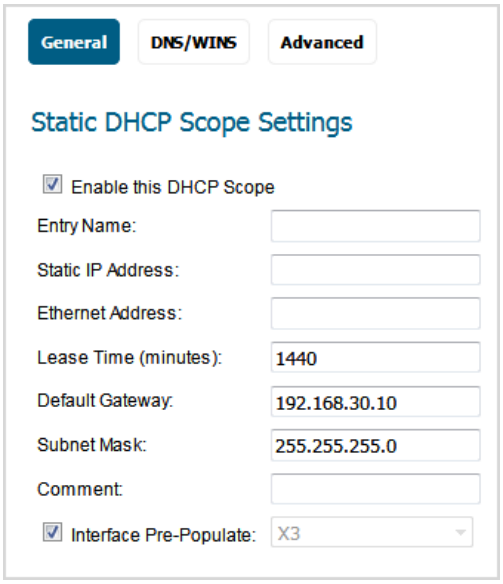
Add DHCPv6 Static Scope



1. To enable this scope, ensure **Enable this DHCP Scope** is selected. This option is selected by default.
2. Enter a name for the static DHCPv6 entry in the **Entry Name** field.
3. Enter the prefix the scope uses to distribute IPv6 addresses in the **Prefix** field.
4. Enter the IPv6 address in the **Static IPv6 Address** field. The address must fall within the scope of the prefix.
5. Enter an IAID (Interface Association Identifier), in decimal format, in the **IAID** field. The maximum length is 10 numbers, and the maximum length is 4294967295.
6. Enter a DUID (Device Unique Identifier) in the **DUID** field. The maximum length is 128 characters.
7. Enter the valid lifetime of an IPv6 address leased by the scope, in minutes, in the **Valid Lifetime** field. The minimum is 0, the maximum is 9999, and the default is **2160**.
8. Enter the preferred lifetime of an IPv6 address leased by the scope, in minutes, in the **Preferred Lifetime** field. The minimum is 0, the maximum is 9999, and the default is **1440**.
9. Optionally, enter a comment in the **Comment** field.

10 For how to configure DNS and Advanced settings, see DNS on page 529 and Advanced on page 530, respectively.

Static Entry Configuration



1 To enable this scope, ensure **Enable this DHCP Scope** is selected. This option is selected by default.

2 Enter a name for the static entry in the **Entry Name** field.

3 Enter the device IP address in the **Static IP Address** field.

4 Enter the device Ethernet (MAC) address in the **Ethernet Address** field.

5 To populate the **Lease Time**, **Default Gateway**, and **Subnet Mask** fields with default values for a certain interface, select **Interface Pre-Populate** near the bottom of the dialog. The drop-down menu becomes available. This option is not selected by default.

   a Select the interface from the drop-down menu. The populated IP addresses are in the same private subnet as the selected interface.

IMPORTANT: To select an interface from the Interface menu, it must first be fully configured and it must be of the zone type, LAN, WLAN, or DMZ, or be a VLAN sub-interface.

6 Enter the number of minutes an IP address is leased by the scope before it is issued another IP address in the **Lease Time (minutes)** field. The minimum is 0, the maximum is 9999, and **1440** minutes (24 hours) is the default.

7 Use the populated gateway address or enter the IP address of the gateway into the **Default Gateway** field.

8 Use the populated subnet mask or enter the gateway subnet mask into the **Subnet Mask** field.

9 Optionally, enter a comment in the **Comment** field.

10 For how to configure DNS/WINS and Advanced settings, see DNS/WINS on page 532 and Advanced on page 533, respectively.

11 Click **OK** to add the settings to the firewall.

12 Click **Accept** for the settings to take effect on the firewall.

For more information on VoIP support features on the SonicWall Security Appliance, see About VoIP on page 765.

Configuring DHCP Generic Options for DHCP Lease Scopes

This section provides configuration tasks for DHCP generic options for lease scopes.

NOTE: Before generic options for a DHCP lease scope can be configured, a static or dynamic DHCP server lease scope must be created.

The RFC-Defined DHCP Option Numbers on page 536 provides a list of DHCP options by RFC-assigned option number.

To configure DHCP generic options for DHCP server lease scopes:

1 If:

- Modifying an existing DHCP lease scope:
  1) Locate the lease scope under DHCP Server Lease Scopes on **Network > DHCP Server**.
  2) Click the **Configure** icon.
  3) Click **Advanced** on the displayed dialog.
- Creating a new DHCP lease scope:
  1) Click the **Advanced** tab after configuring the options under the **General** and **DNS/WINS** tabs (see Configuring DHCP Server for Dynamic Ranges on page 528 or Configuring Static DHCP Entries on page 534).

2 Select a DHCP option or option group in the **DHCP Generic Option Group** drop-down menu.

When the **Network Boot Settings** fields are configured for use with PXE, select **PXE** here.

3 To always use DHCP options for this DHCP server lease scope, check **Send Generic options always**.

4 Click **OK**.

RFC-Defined DHCP Option Numbers

Option Number	IPv6 √	Name	Description
2		Time Offset	Time offset in seconds from UTC
3		Router	N/4 router addresses
4		Time Servers	N/4 time server addresses
5		Name Servers	N/4 IEN-116 server addresses
6		DNS Servers	N/4 DNS server addresses
7		Log Servers	N/4 logging server addresses
8		Cookie Servers	N/4 quote server addresses
9		LPR Servers	N/4 printer server addresses
10		Impress Servers	N/4 impress server addresses
11		RLP Servers	N/4 RLP server addresses
12	√	Host Name	Hostname string, such as **(Server Unicast)**
13		Boot File Size	Size of boot file in 512-byte chunks
14		Merit Dump File	Client to dump and name of file to dump to

Option Number	IPv6 √	Name	Description
15		Domain Name	DNS domain name of the client
16		Swap Server	Swap server addresses
17		Root Path	Path name for root disk
18		Extension File	Patch name for more BOOTP info
19		IP Layer Forwarding	Enable or disable IP forwarding
20		Src route enabler	Enable or disable source routing
21	√	Policy Filter (IPv4) SIP Servers Domain Name List (IPv6)	Routing policy filters (IPv4) Enables listing of SIP Servers domain names (IPv6)
22	√	Maximum DG Reassembly Size (IPv4) SIP Servers IPv6 Address List (IPv6)	Maximum datagram reassembly size (IPv4) Enables listing of SIP Servers IPv6 Addresses (IPv6)
23	√	Default IP TTL (IPv4) DNS Recursive Name Server (IPv6)	Default IP time-to-live (IPv4) Enables listing of DNS Recursive Name servers (IPv6)
24	√	Path MTU Aging Timeout (IPv4) Domain Search List (IPv6)	Path MTU aging timeout (IPv4) Enables listing of domain names for searching (IPv6)
25		MTU Plateau	Path MTU plateau table
26		Interface MTU Size	Interface MTU size
27	√	All Subnets Are Local (IPv4) Network Information Service (NIS) Servers (IPv6)	All subnets are local (IPv4) Enables listing of Network Information Service (NIS) servers (IPv6)
28	√	Broadcast Address (IPv4) Network Information Service V2 (NIS+) Servers (IPv6)	Broadcast address (IPv4) Enables listing of Network Information Service V2 (NIS+) servers (IPv6)
29	√	Perform Mask Discovery (IPv4) Network Information Service (NIS) Domain Name (IPv6)	Perform mask discovery (IPv4) Enables listing of Network Information Service (NIS) domain names (IPv6)
30	√	Provide Mask to Others (IPv4) Network Information Service V2 (NIS+) Domain Name (IPv6)	Provide mask to others (IPv4) Enables listing of Network Information Service V2 (NIS+) domain names (IPv6)
31	√	Perform Router Discovery (IPv4) Simple Network Time Protocol (SNTP) Servers (IPv6)	Perform router discovery (IPv4) Enables listing of Simple Network Time Protocol (SNTP) servers (IPv6)
32	√	Router Solicitation Address (IPv4) Information Refresh Time (IPv6)	Router solicitation address (IPv4) Information refresh time (IPv6)
33		Static Routing Table	Static routing table
34		Trailer Encapsulation	Trailer encapsulation
35		ARP Cache Timeout	ARP cache timeout
36		Ethernet Encapsulation	Ethernet encapsulation
37		Default TCP Time to Live	Default TCP time to live
38		TCP Keepalive Interval	TCP keepalive interval

Option Number	IPv6 √	Name	Description
39		TCP Keepalive Garbage	TCP keepalive garbage
40		NIS Domain Name	NIS domain name
41		NIS Server Addresses	NIS server addresses
42		NTP Servers Addresses	NTP servers addresses
43		Vendor Specific Information	Vendor specific information
44		NetBIOS Name Server	NetBIOS name server
45		NetBIOS Datagram Distribution	NetBIOS datagram distribution
46		NetBIOS Node Type	NetBIOS node type
47		NetBIOS Scope	NetBIOS scope
48		X Window Font Server	X window font server
49		X Window Display Manager	X window display manager
50		Requested IP address	Requested IP address
51		IP Address Lease Time	IP address lease time
52		Option Overload	Overload “sname” or “file”
53		DHCP Message Type	DHCP message type
54		DHCP Server Identification	DHCP server identification
55		Parameter Request List	Parameter request list
56		Message	DHCP error message
57		DHCP Maximum Message Size	DHCP maximum message size
58		Renew Time Value	DHCP renewal (T1) time
59		Rebinding Time Value	DHCP rebinding (T2) time
60		Client Identifier	Client identifier
61		Client Identifier	Client identifier
62		Netware/IP Domain Name	Netware/IP domain name
63		Netware/IP sub Options	Netware/IP sub options
64		NIS+ V3 Client Domain Name	NIS+ V3 client domain name
65		NIS+ V3 Server Address	NIS+ V3 server address
66		TFTP Server Name	TFTP server name
67		Boot File Name	Boot file name
68		Home Agent Addresses	Home agent addresses
69		Simple Mail Server Addresses	Simple mail server addresses
70		Post Office Server Addresses	Post office server addresses
71		Network News Server Addresses	Network news server addresses
72		WWW Server Addresses	WWW server addresses
73		Finger Server Addresses	Finger server addresses
74		Chat Server Addresses	Chat server addresses
75		StreetTalk Server Addresses	StreetTalk server addresses
76		StreetTalk Directory Assistance Addresses	StreetTalk directory assistance addresses
77		User Class Information	User class information
78		SLP Directory Agent	Directory agent information

Option Number	IPv6 √	Name	Description
79		SLP Service Scope	Service location agent scope
80		Rapid Commit	Rapid commit
81		FQDN, Fully Qualified Domain Name	Fully qualified domain name
82		Relay Agent Information	Relay agent information
83		Internet Storage Name Service	Internet storage name service
84		Undefined	N/A
85		Novell Directory Servers	Novell Directory Services servers
86		Novell Directory Server Tree Name	Novell Directory Services server tree name
87		Novell Directory Server Context	Novell Directory Services server context
88		BCMCS Controller Domain Name List	CMCS controller domain name list
89		BCMCS Controller IPv4 Address List	BCMCS controller IPv4 address list
90		Authentication	Authentication
91- 92		Undefined	N/A
93		Client System	Client system architecture
94		Client Network Device Interface	Client network device interface
95		LDAP Use	Lightweight Directory Access Protocol
96		Undefined	N/A
97		UUID/GUID-based Client Identifier	UUID/GUID-based client identifier
98		Open Group’s User Authentication	Open group’s user authentication
99 - 108		Undefined	N/A
109		Autonomous System Number	Autonomous system number
110 - 111		Undefined	N/A
112		NetInfo Parent Server Address	NetInfo parent server address
113		NetInfo Parent Server Tag	NetInfo parent server tag
114		URL:	URL
115		Undefined	N/A
116		Auto Configure	DHCP auto-configuration
117		Name Service Search	Name service search
118		Subnet Collection	Subnet selection
119		DNS Domain Search List	DNS domain search list
120		SIP Servers DHCP Option	SIP servers DHCP option
121		Classless Static Route Option	Classless static route option
122		CCC, CableLabs Client Configuration	CableLabs client configuration
123		GeoConf	GeoConf
124		Vendor-Identifying Vendor Class	Vendor-identifying vendor class
125		Vendor Identifying Vendor Specific	Vendor-identifying vendor specific
126 - 127		Undefined	N/A
128		TFTP Server IP Address	TFTP server IP address for IP phone software load
129		Call Server IP Address	Call server IP address

Option Number	IPv6 √	Name	Description
130		Discrimination String	Discrimination string to identify vendor
131		Remote Statistics Server IP Address	Remote statistics server IP address
132		802.1Q VLAN ID	IEEE 802.1Q VLAN ID
133		802.1Q L2 Priority	IEEE 802.1Q layer 2 priority
134		Diffserv Code Point	Diffserv code point for VoIP signalling and media streams
135		HTTP Proxy For Phone Applications	HTTP proxy for phone-specific applications
136 - 149		Undefined	N/A
150		TFTP Server Address, Etherboot, GRUB Config	TFTP server address, Etherboot, GRUB configuration
151 - 174		Undefined	N/A
175		Ether Boot	Ether Boot
176		IP Telephone	IP telephone
177		Ether Boot PacketCable and CableHome	Ether Boot PacketCable and CableHome
178 - 207		Undefined	N/A
208		`pxelinux.magic` (string) = `241.0.116.126`	`pxelinux.magic` (string) = `241.0.116.126`
209		`pxelinux.configfile` (text)	`pxelinux.configfile` (text)
210		`pxelinux.pathprefix` (text)	`pxelinux.pathprefix` (text)
211		`pxelinux.reboottime`	`pxelinux.reboottime`
212 - 219		Undefined	N/A
220		Subnet Allocation	Subnet allocation
221		Virtual Subnet Allocation	Virtual subnet selection
222 - 223		Undefined	N/A
224 - 254		Private Use	Private use

DHCP and IPv6

For complete information on the SonicOS implementation of IPv6, see IPv6 on page 932.

27

Using IP Helper

Topics:

- About IP Helper on page 541
  - VPN Tunnel Interface Support for IP Helper on page 542
  - DHCPv6 Relay on page 543
- **Network > IP Helper** on page 545
  - Relay Protocols on page 546
  - Policies on page 547
  - DHCP/DHCPv6 Relay Leases on page 547
- Configuring IP Helper on page 548
  - Enabling IP Helper on page 548
  - Viewing Traffic Statistics on page 548
  - Managing Relay Protocols on page 548
  - Managing IP Helper Policies on page 550
  - Filtering What DHCP Relay Leases are Displayed on page 552

About IP Helper

IMPORTANT: IP Helper is not supported for WAN interfaces or for interfaces that are configured for NAT.

Many User Datagram Protocols (UDP) rely on broadcast/multicast to find its respective server, usually requiring their servers to be present on the same broadcast subnet. To support cases where servers lie on different subnets than clients, a mechanism is needed to forward these UDP broadcasts/multicasts to those subnets. This mechanism is referred to as UDP broadcast forwarding. IP Helper helps broadcast/multicast packets to cross a SonicWall Security Appliance's interface and be forwarded to other interfaces based on policy. IP Helper allows the Security Appliance to forward DHCP requests originating from its interfaces to a centralized DHCP server.

IP Helper supports user-defined protocols and extended policies. IP Helper provides better control on existing NetBIOS/DHCP relay applications. Some of the built-in applications that have been extended are:

Extended Built-in Relay Applications

Protocol	UDP Port Number
DHCP	67/68
DHCPv6	546, 547
Net-Bios NS	137
Net-Bios Datagram	138
DNS	53

Extended Built-in Relay Applications (Continued)

Protocol	UDP Port Number
Time Service	37
Wake on LAN (WOL)	
mDNS	5353 Multicast address: 224.0.0.251

VPN Tunnel Interface Support for IP Helper

The VPN Tunnel Interface can support IP Helper. DHCP Replay in IP Helper with Tunnel Interface support shows a simple example of DHCP replay in IP Helper:

- PC is the device needed to get an IPv4 address from the DHCP protocol.
- GatewayA is the gateway-enabled IP helper.
- GatewayB is the gateway with a DHCP server.

DHCP Replay in IP Helper with Tunnel Interface support



To configure IP Helper with a VPN Tunnel Interface:

NOTE: The numbers in DHCP Replay in IP Helper with Tunnel Interface support correspond to the numbered tasks.

1 In PC:
   a Connect to the LAN (X0) subnet of GatewayA.
   b Set to obtain an IP address through DHCP mode.
2 Set up a VPN tunnel between GatewayA and GatewayB.
   - Add a VPN Tunnel Interface.
3 In GatewayB:
   a Add a route entry from the Tunnel Interface's IP address to GatewayA's X0 interface.
   b Add the outbound interface of the Tunnel Interface.
   c Add an IP address range as the DHCP scope for PC.
4 In GatewayA:
   a Enable IP Helper.
   b Add an IP Helper DHCP relay protocol from X0 to GatewayB's Tunnel Interface address. The protocol is DHCP.

DHCPv6 Relay

Topics:

- About DHCPv6 Relay on page 543
- Configuring DHCPv6 Relay on page 544

About DHCPv6 Relay

SonicOS supports DHCPv6 Relay. A DHCP relay agent is a node that acts as an intermediary to deliver DHCP messages between clients and server, and is on the same link as the client. A DHCPv6 relay agent is used to relay messages between the client and the server when they are not on the same IPv6 link. The DHCPv6 relay agent operation is transparent to the client.

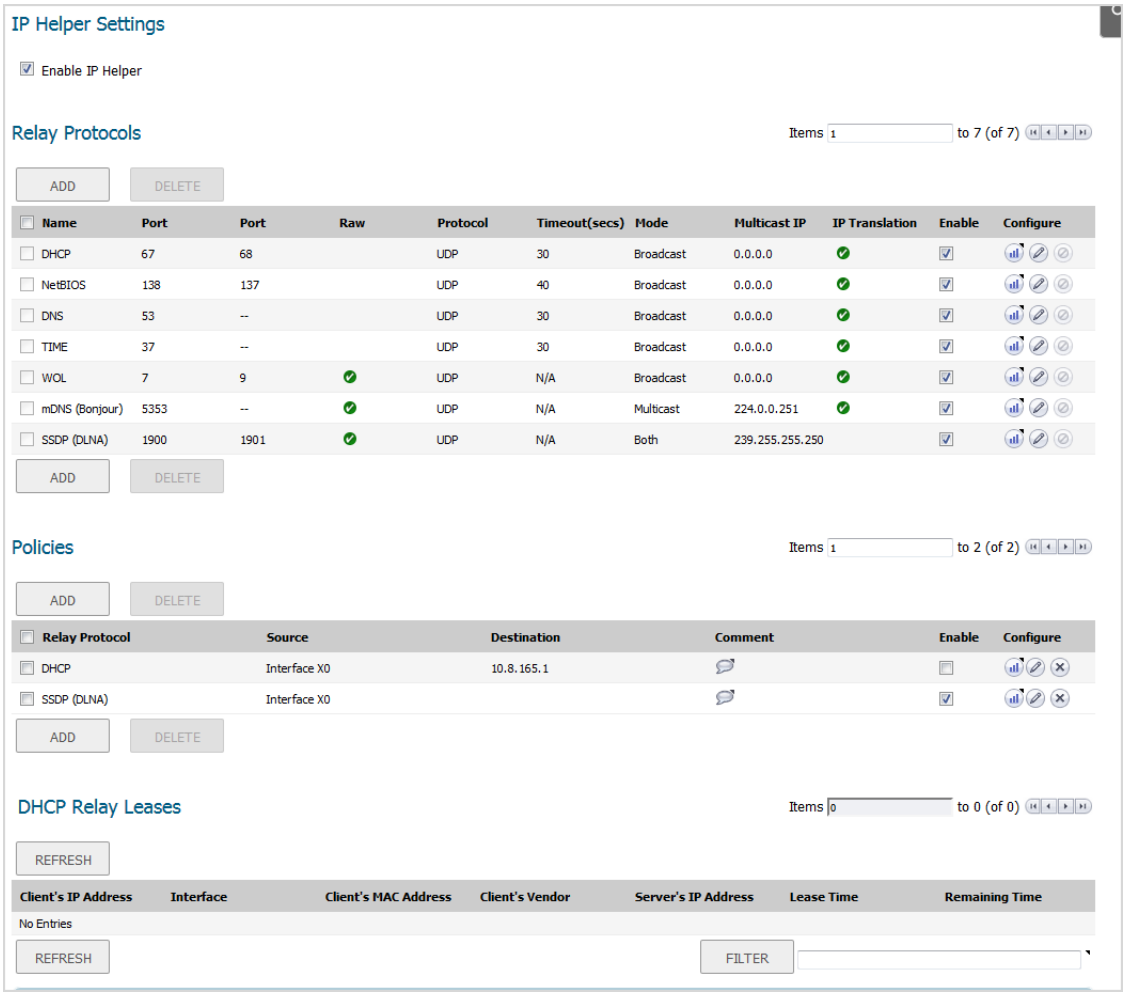
In SonicOS, supported destination addresses can be global addresses or link-local addresses, but not multicast addresses.

DHCPv6 relay can be enabled on both physical and virtual interfaces. DHCPv6 is a built-in application in IP Helper protocols.

Configuring DHCPv6 Relay

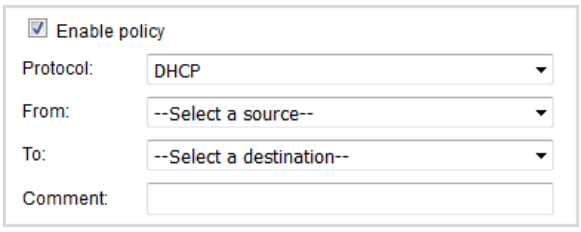
To configure DHCPv6 Relay:

1 Navigate to the **MANAGE | System Setup > Network > IP Helper** page.



2 Scroll to the **Policies** section.

3 Click **ADD**. The **Add IP Helper Policy** dialog displays.



4 Select **DHCPv6** from **Protocol**.

5 Select the desired interface from **From**.

6 In the **To** field, type in the destination IPv6 address. This may be a list of destination addresses, which may include unicast addresses, or other addresses you select. The address cannot be a multicast address.

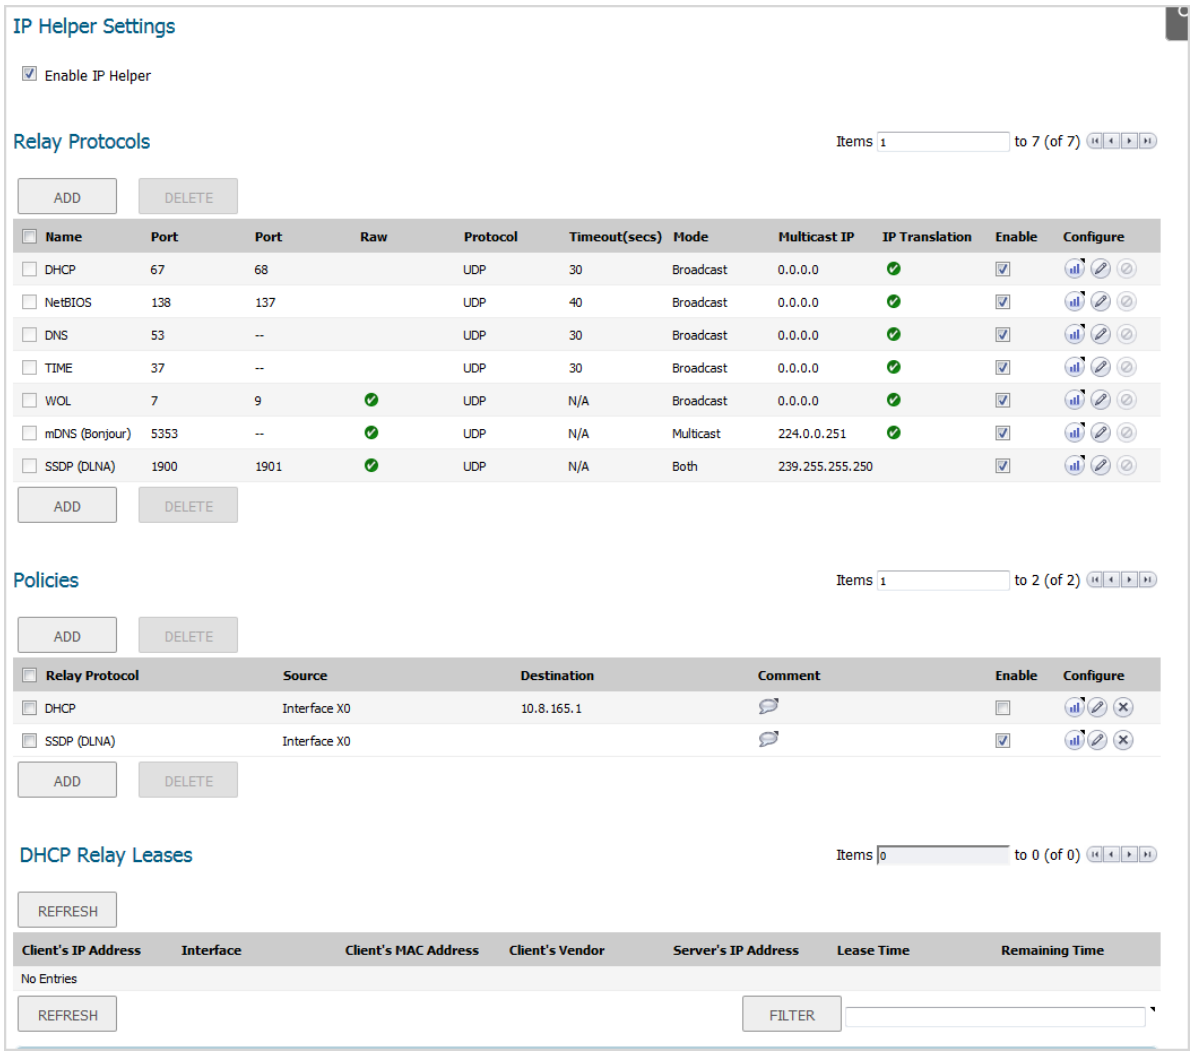
7 If the destination in the **To** field is a:

- Global address, there is no need to select an egress interface. Go to Step 8.
- Link-local address, select an egress interface from **Egress Interface**.

8 Click **OK**.

A new DHCP lease appears in the **DHCPv6 Relay Leases** section of the page when the client gets a new IP address from the server.

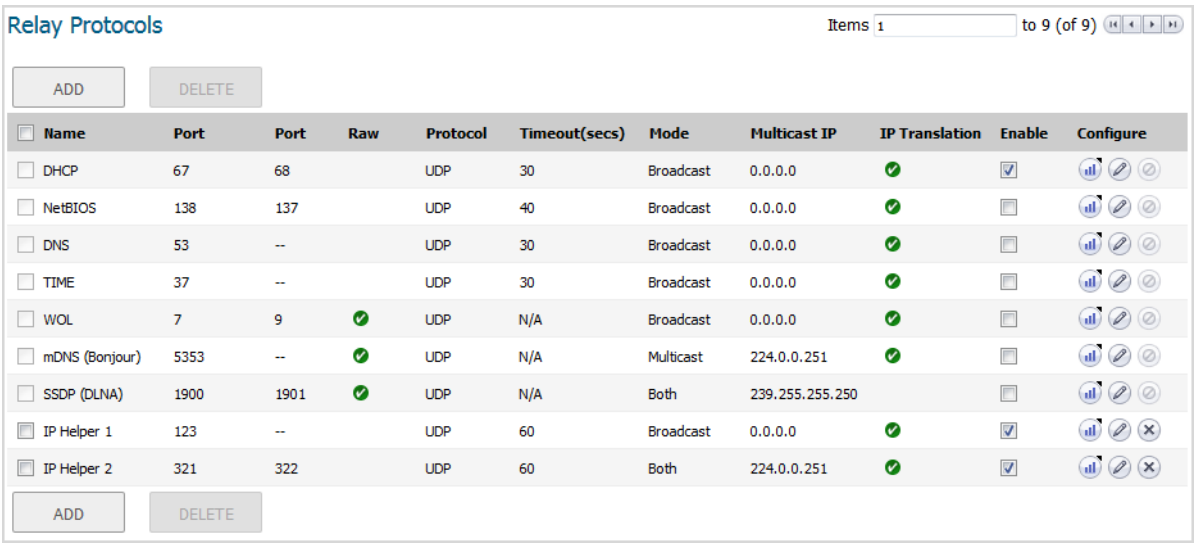
Network > IP Helper



Topics:

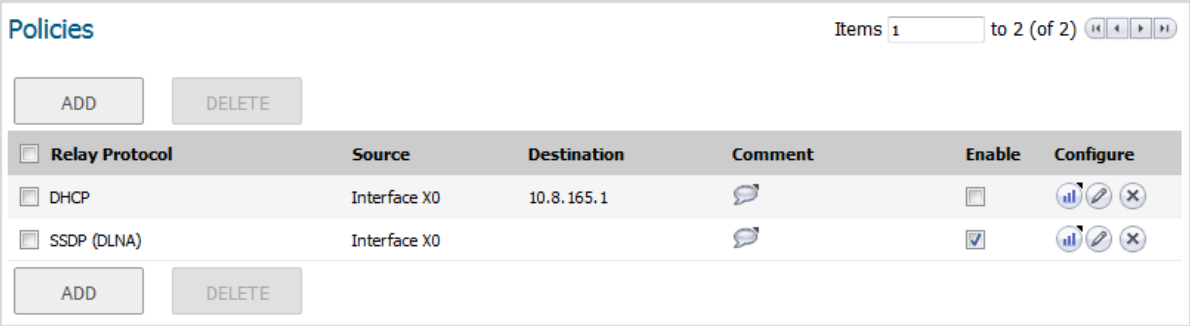
- Relay Protocols on page 546
- Policies on page 547
- DHCP/DHCPv6 Relay Leases on page 547

Relay Protocols



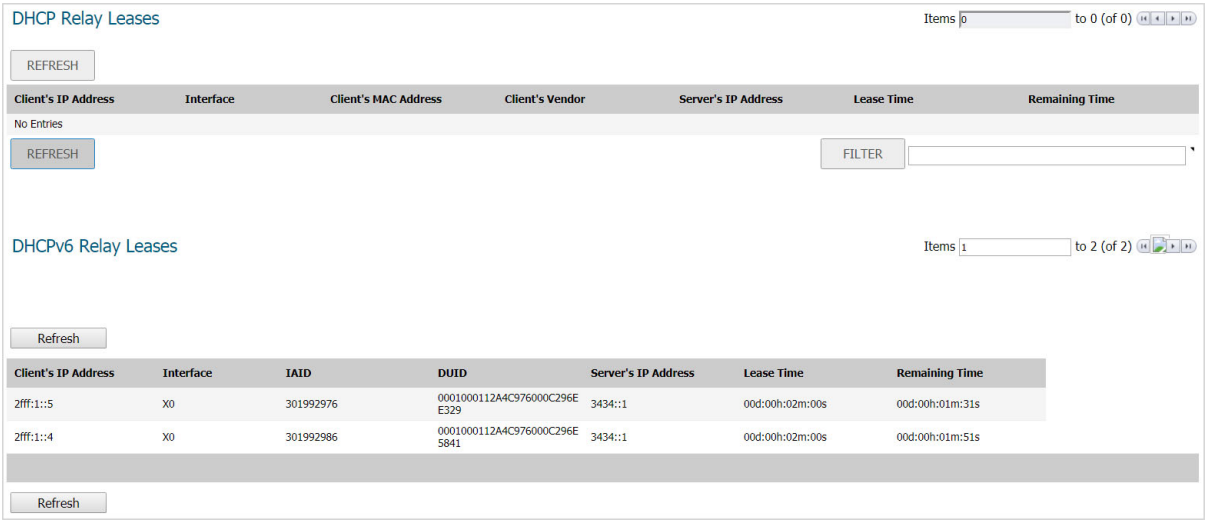
Name	IP Helper application name.
Port	First UDP port number for the IP Helper application.
Port	Optional second UDP port number for the IP Helper application.
Raw	Indicates whether raw mode was selected when the IP Helper application was configured. Timeout is ignored if this option is enabled.
Protocol	UDP.
Timeout (secs)	Timeout for the IP Helper cache. **N/A** indicates Raw mode is selected and the timeout is ignored.
Mode	Indicates the mode the protocol supports: • **Broadcast** • **Multicast** • **Both**
Multicast IP	Multicast IP the protocol uses.
IP Translation	Indicates whether the source IP address is translated when packets are forwarded by an IP Helper policy.
Enable	Indicates whether the IP Helper policy is enabled.
Configure	Contains the **Statistics**, **Edit**, and **Delete** icons for the entries. **NOTE:** Only user-generated Relay protocols can be deleted.

Policies



Relay Protocol	Protocol for the policy.
Source	Interface or zone for the policy.
Destination	Network destination.
Comment	Comment entered when the policy was configured.
Enable	Indicates whether the IP Helper policy is enabled.
Configure	Contains the Statistics, Edit, and Delete icons for each entry.

DHCP/DHCPv6 Relay Leases



Client's IP Address	IP address of the client device.
Interface	Receiving interface on the Security Appliance.
DHCP Relay Leases:	
Client's MAC Address	MAC address of the client device.
Client's Vendor	Manufacturer of the client device.
DHCPv6 Relay Leases:	
IAID	Interface ID; an Interface Association Identifier that is a binding between the interface and one or several IP addresses.
DUID	Device (host) ID; a DHCP Unique Identifier for a DHCP participant.

Server's IP Address	IP address of the DHCP server.
Lease Time	Time of the relay lease.
Remaining Time	Time remaining on the relay lease.

To refresh the DHCP Relay Leases table:

1 Click **REFRESH**.

Configuring IP Helper

Topics:

- Enabling IP Helper on page 548
- Managing Relay Protocols on page 548
- Managing IP Helper Policies on page 550

Enabling IP Helper

To activate IP Helper features:

1 Navigate to **Network > IP Helper.**

2 Select **Enable IP Helper** in **IP Helper Settings**.

Managing Relay Protocols

Topics:

- Viewing Traffic Statistics on page 548
- Adding User-Defined Relay Protocols on page 549
- Deleting Custom Protocols on page 550

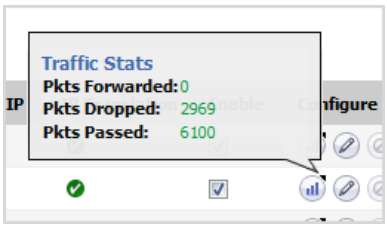
Viewing Traffic Statistics

You can view traffic statistics for both the Relay Protocols table and the Policies table.

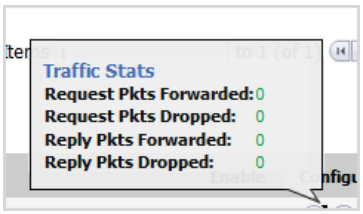
To view traffic statistics:

1 Hover the cursor over a protocol or policy's **Statistics** icon. A popup displays the traffic status for that entry.

Relay Protocols table



Policies table

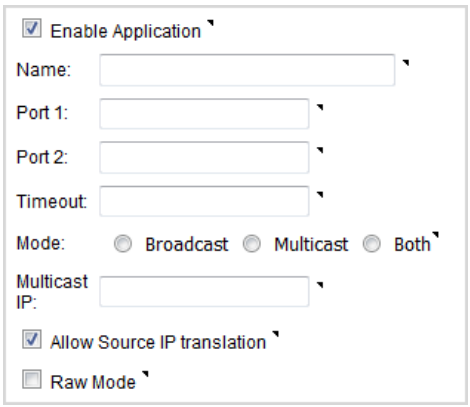


Adding User-Defined Relay Protocols

To add a relay protocol:

1 Navigate to **Network > IP Helper**.

2 Click **Add** in the **Relay Protocols** section. The **Add IP Helper Application** dialog displays.



3 Enable the IP Helper application by selecting **Enable Application**.

NOTE: If this option is disabled, all IP Helper cache is deleted.

4 Enter a unique, case-sensitive name for the IP Helper application in the **Name** field.

5 In the **Port 1** field, specify a unique UDP port number for the application.

6 Optionally, in the **Port 2** field, specify a second unique UDP port number for the application.

7 Optionally, specify the IP Helper cache timeout, in seconds, in an increment of 10 from 10 to 60, in the **Timeout** field. If a timeout is not specified, a default value of **30** seconds is selected.

TIP: This field is ignored if **Raw Mode** is selected.

8 Choose a **Mode**:

- **Broadcast**
- **Multicast**
- **Both**

9 If you selected **Multicast** or **Both** for **Mode**, specify a valid multicast IP that this protocol is used in the **Multicast IP** field.

10 To allow the source IP address to be translated when a packet is forwarded by an IP Helper policy, select **Allow Source IP Translation**. This option is selected by default.

11 To prevent a cache from being created when a packet is forwarded by an IP Helper policy, select **Raw Mode**. Unidirectional forwarding is supported. This option is not selected by default.

NOTE: Any time set in the **Timeout** field is ignored.

12 Click **OK**.

Deleting Custom Protocols

To delete a custom protocol:

1 Navigate to **Network > IP Helper**.

2 Select the **Delete** icon for that protocol.

To delete one or more custom relay protocols:

1 Navigate to **Network > IP Helper**.

2 Select the left-most checkbox(es) (by the protocol name) of the desired protocol(s). **DELETE** becomes available.

3 Click **DELETE**.

To delete all custom relay protocols:

1 Navigate to **Network > IP Helper**.

2 Select the checkbox in the **Relay Protocols** table header. **DELETE** becomes available.

3 Click **DELETE**.

Managing IP Helper Policies

IP Helper policies allow you to forward DHCP and NetBIOS broadcasts from one interface to another interface.

IMPORTANT: IP Helper is not supported for WAN interfaces or for interfaces that are configured for NAT.

Topics:

- Adding an IP Helper Policy on page 550
- Editing an IP Helper Policy on page 551
- Deleting IP Helper Policies on page 551
- Displaying IP Helper Cache from TSR on page 552

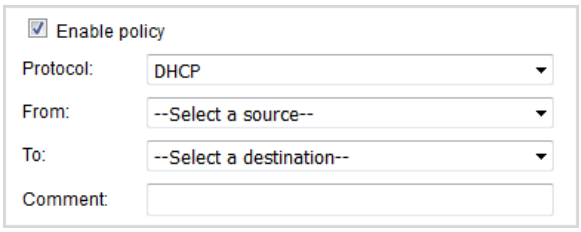
Adding an IP Helper Policy

You can add up to 256 policies.

To add an IP Helper policy:

1 Navigate to **Network > IP Helper**.

2 Click **Add** for the **IP Helper Policies** table. The **Add IP Helper Policy** dialog displays.



3 The policy is enabled by default. To configure the policy without enabling it, clear the **Enabled** checkbox.

4 Select a protocol from the **Protocol** menu. The default is **DHCP**.

5 Select a source interface or zone from **From**.

6 From **To**, select either:

- A destination Address Group or Address Object.
- **Create a new network** to create a new Address Object. The **Add Address Object** dialog displays. For further information about creating an Address Object, see the *SonicOS Policies Guide*.

7 Enter an optional comment in the **Comment** field.

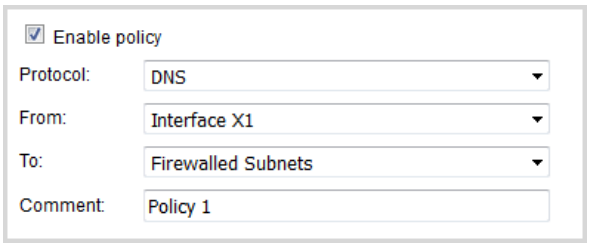
8 Click **OK**.

Editing an IP Helper Policy

To edit an IP Helper policy:

1 Navigate to **Network > IP Helper**.

2 Click the **Edit** icon in the **Configure** column of the entry in the **IP Helper Policies** table. The **Edit IP Helper Policy** dialog displays.



3 The settings are the same as the **Add IP Helper Policy** dialog. For information about the dialog, see Adding an IP Helper Policy on page 550.

Deleting IP Helper Policies

To delete a custom policy:

1 Navigate to **Network > IP Helper**.

2 Select the **Delete** icon in the **Policies** table for that policy.

To delete one or more custom policies:

1 Navigate to **Network > IP Helper**.

2 Select the left-most checkbox(es) (by the relay protocol) of the desired policies. **DELETE** becomes available.

3 Click **DELETE**.

To delete all custom policies:

1 Navigate to **Network > IP Helper**.

2 Select the checkbox in the **Policies** table header. **DELETE** becomes available.

3 Click **DELETE**.

Filtering What DHCP Relay Leases are Displayed

You can display only a specific device(s) in the **Anti-Spoof Cache** and **Spoof Detected List** tables by using the **Filter** function.

To filter the table display:

1 Navigate to **Network > MAC-IP Anti-spoof**.

2 In the **Filter** field below the table to be filtered, specify either the device's IP address, interface, MAC address, host name, or name. The field must be filled using the appropriate syntax for operators shown in Filter operator syntax options.

Filter operator syntax options

Operator	Syntax Options
Value with a type	• Ip=1.1.1.1 or ip=1.1.1.0/24 • Mac=00:01:02:03:04:05 • Iface=x1
String	• X1 • 00:01 • Tst-mc • 1.1.
AND	• Ip=1.1.1.1;iface=x1 • Ip=1.1.1.0/24;iface=x1;just-string
OR	• Ip=1.1.1.1,2.2.2.2,3.3.3.0/24 • Iface=x1,x2,x3
Negative	• !ip=1.1.1.1;!just-string • !iface=x1,x2
Mixed	• Ip=1.1.1.1,2.2.2.2;mac=00:01:02:03:04:05; just-string;!iface=x1,x2

Displaying IP Helper Cache from TSR

The TSR shows all the IP Helper caches, current policies, and protocols:

```
#IP_HELPER_START
IP Helper
-----IP Helper Global Run-time Data-------
```

```
 IP Helper is OFF
 IP Helper - DHCP Relay is OFF
 IP Helper - Netbios Relay is OFF
  Total Number Of Fwded Packets           :0
  Total Number Of Dropped Packets         :0
  Total Number Of Passed Packets          :0
  Total Number Of Unknown Packets         :0
  Total Number Of record create failure   :0
  Total Number Of element create failure  :0User-defined
-----IP Helper Applications -------
 Name: DHCP
  Port: 67, 68, Max Record: 4000, Status: OFF
  CanBeDel: NO, ChangeIp: 1, Raw: NO
  Max Element: 8000, Timeout: 3, index: 1, proto: 1,
  Record Count: 0, Element Count: 0,
  Fwded: 0, Dropped: 0, Passed: 0
 Name: NetBIOS
  Port: 138, 137, Max Record: 4000, Status: OFF
  CanBeDel: NO, ChangeIp: 1, Raw: NO
  Max Element: 8000, Timeout: 4, index: 2, proto: 1,
  Record Count: 0, Element Count: 0,
  Fwded: 0, Dropped: 0, Passed: 0
 Name: DNS
  Port: 53, 0, Max Record: 8000, Status: OFF
  CanBeDel: NO, ChangeIp: 1, Raw: NO
  Max Element: 16000, Timeout: 3, index: 3, proto: 1,
  Record Count: 0, Element Count: 0,
  Fwded: 0, Dropped: 0, Passed: 0
 Name: TIME
  Port: 37, 0, Max Record: 8000, Status: OFF
  CanBeDel: NO, ChangeIp: 1, Raw: NO
  Max Element: 16000, Timeout: 3, index: 4, proto: 1,
  Record Count: 0, Element Count: 0,
  Fwded: 0, Dropped: 0, Passed: 0
 Name: WOL
  Port: 7, 9, Max Record: 8000, Status: OFF
  CanBeDel: NO, ChangeIp: 1, Raw: YES
  Max Element: 16000, Timeout: 3, index: 5, proto: 1,
  Record Count: 0, Element Count: 0,
  Fwded: 0, Dropped: 0, Passed: 0
 Name: mDNS
  Port: 5353, 0, Max Record: 8000, Status: OFF
  CanBeDel: NO, ChangeIp: 1, Raw: YES
  Max Element: 16000, Timeout: 3, index: 6, proto: 1,
  Record Count: 0, Element Count: 0,
  Fwded: 0, Dropped: 0, Passed: 0
----------GEN APP Relay Policy-----------
---------------------------------Record Table-----------------------------
Record(hash)[ClientIP, ClientIf, ClientMac, Proto, Vpn, transId, Age(pkts)]
 Elmnt(hash)[serverIp, serverIf, srcIp, dhcpMac, transId, Vpn, proto(fm,to)]
-------------------------------------------------------------------------
----------DHCP Relay Policy-----------
--------NETBIOS Relay Policy-----------#IP_HELPER_END
```

28

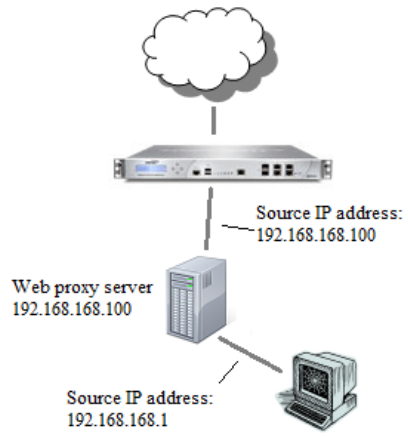
Setting Up Web Proxy Forwarding

Topics:

- **Network > Web Proxy** on page 554
  - Configuring Automatic Proxy Forwarding (Web Only) on page 555
  - Configuring User Proxy Servers on page 556

Network > Web Proxy

When users access the web through a proxy server located on the internal network (between the user and the SonicWall Security Appliance), the HTTP/HTTPS connections seen by the Security Appliance originate from the proxy server, not from the user.



A web proxy server intercepts HTTP requests and determines if it has stored copies of the requested Web pages. If it does not, the proxy completes the request to the server on the Internet, returning the requested information to the user and also saving it locally for future requests. Setting up a web proxy server on a network can be cumbersome, because each computer on the network must be configured to direct web requests to the server.

If you have a proxy server on your network, instead of configuring each computer's web browser to point to the proxy server, you can move the server to the WAN or DMZ zone and enable Web Proxy Forwarding using the settings on the **Network > Web Proxy** page. The Security Appliance automatically forwards all web proxy requests to the proxy server without requiring all the computers on the network to be configured.

Topics:

- Configuring Automatic Proxy Forwarding (Web Only) on page 555
- Configuring User Proxy Servers on page 556

Configuring Automatic Proxy Forwarding (Web Only)

NOTE: To enable Web Proxy, enable CFS on the related zones where clients are from (this is not necessary when using the WXA's Web Cache on TZ series appliances).

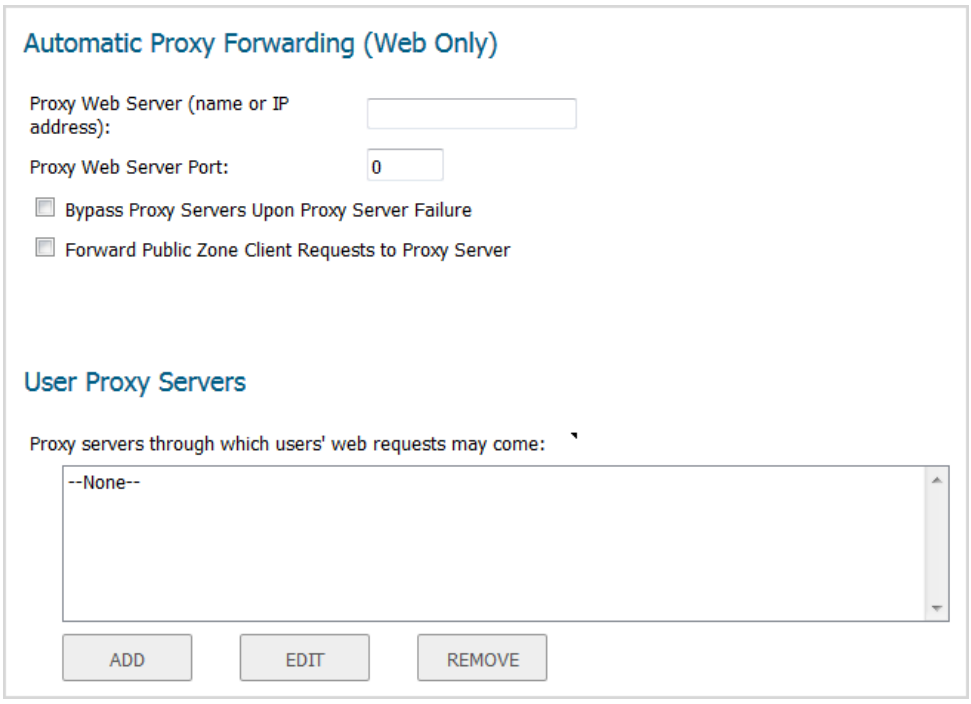
To configure Automatic Proxy Forwarding (Web Only):

1 Connect the Web proxy server to a hub.

2 Connect the hub to the firewall WAN or DMZ port.

NOTE: The proxy server must be located on a WAN or DMZ zone; it can not be located on the LAN.

3 Go to **Network > Web Proxy**.



4 To have all web proxy requests forwarded to the proxy server automatically, in the **Automatic Proxy Forwarding (Web Only)** section, enter the name or IP address of the proxy server in the **Proxy Web Server (name or IP address)** field. The minimum length is 0 and the maximum is 39 characters.

5 Enter the proxy IP port in the **Proxy Web Server Port** field. The default is 0.

6 To have clients access the Internet directly if the web proxy server becomes unavailable, select **Bypass Proxy Servers Upon Proxy Server Failure**. This option is disabled by default.

NOTE: The **Bypass Proxy Servers Upon Proxy Server Failure** checkbox allows clients behind the firewall to bypass the Web proxy server if it becomes unavailable. Instead, the client's browser accesses the Internet directly as if a Web proxy server is not specified.

7 To force clients on public zones to use the proxy server as well. select **Forward Public Zone Client Requests to Proxy Server**. This option is disabled by default.

8 Click **Accept**.

After the Security Appliance has been updated, a message confirming the update is displayed at the bottom of the browser window.

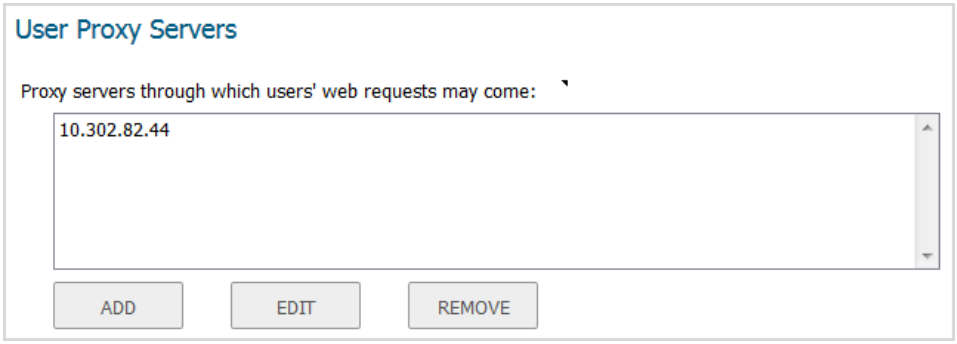
Configuring User Proxy Servers

You can configure up to 32 user proxy servers by entering the host name or IP address.

To configure a user proxy sever:

1 Navigate to **Network > Web Proxy**.

2 Go to the **User Proxy Servers** section.



3 Click **Add**. The **Add Proxy Servers** pop-up dialog displays.



NOTE: If users' web requests go through a proxy server before reaching the SonicWall Security Appliance, then the web requests seen by the Security Appliance come from the proxy server, not directly from the user. Thus, the Security Appliance cannot identify the user from the source IP address. The proxy server to identify the source of each web request, however, normally includes this information in the HTTP headers.

If any internal proxy servers are configured here, then the Security Appliance uses the information from the servers to identify the users.

This works for both identification of users accessing the web through proxy servers on the internal network and for remote HTTP management of the Security Appliance through a WAN-side external proxy server.

4 Enter the name or IP address of the proxy server.

5 Click **OK**.

6 Repeat Step 3 through Step 5 to add more proxy servers.

7 Click **Accept**.

8 After you have configured the interface, you can connect it to the host. See Configuring Interfaces on page 259.

Editing User Proxy Servers

To edit the name or IP address of a proxy server:

1 Navigate to **Network > Web Proxy**.

2 Go to the **User Proxy Servers** section.

3 In the **Users Proxy Servers** table, select the proxy server you want to edit.

4 Click **Edit**. The **Edit Proxy Server** pop-up dialog displays.



5 Change the name or IP address of the proxy server.

6 Click **OK**.

Removing User Proxy Servers

To remove a proxy server:

1 Navigate to **Network > Web Proxy**.

2 Go to the **User Proxy Servers** section.

3 In the **Users Proxy Servers** table, select the proxy server you want to remove.

4 Click **Remove**.

5 Click **Accept**.

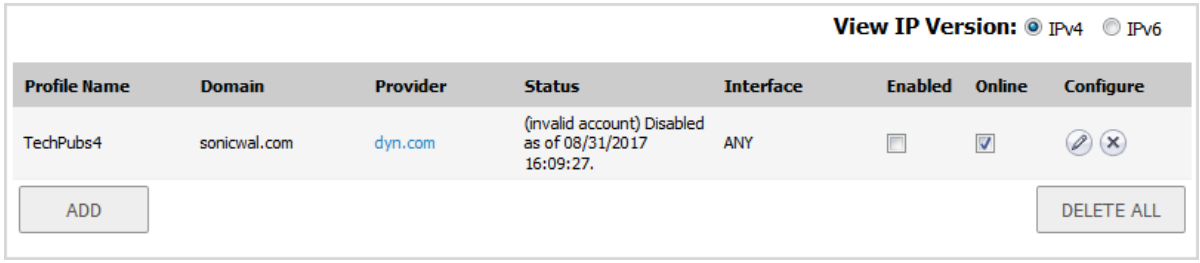
29

Configuring Dynamic DNS

Topics:

- **Network > Dynamic DNS** on page 558
  - About Dynamic DNS on page 558
  - Supported DDNS Providers on page 559
  - Dynamic DNS Profiles Table on page 560
  - Configuring a Dynamic DNS Profile on page 561
  - Editing a DDNS Profile on page 564
  - Deleting DDNS Profiles on page 564

Network > Dynamic DNS



Topics:

- About Dynamic DNS on page 558
- Supported DDNS Providers on page 559
- Dynamic DNS Profiles Table on page 560
- Configuring a Dynamic DNS Profile on page 561
- Editing a DDNS Profile on page 564
- Deleting DDNS Profiles on page 564

About Dynamic DNS

Dynamic DNS (DDNS) is a service provided by various companies and organizations that allows for dynamic changing IP addresses to automatically update DNS records without manual intervention. This service allows for network access using domain names rather than IP addresses, even when the target's IP addresses change. For example, is a user has a DSL connection with a dynamically assigned IP address from the ISP, the user can use

DDNS to register the IP address, and any subsequent address changes, with a DDNS service provider so that external hosts can reach it using an unchanging domain name.

Dynamic DNS implementations change from one service provider to another. There is no strict standard for the method of communication, for the types of records that can be registered, or for the types of services that can be offered. Some providers offer premium versions of their services, as well, for a fee. As such, supporting a particular DDNS provider requires explicit interoperability with that provider's specific implementation.

Most providers strongly prefer that DDNS records only be updated when IP address changes occur. Frequent updates, particularly when the registered IP address is unchanged, may be considered abuse by providers, and could result in your DDNS account getting locked out. Please refer to the use policies posted on the provider's pages and abide by the guidelines. SonicWall does not provide technical support for DDNS providers; the providers themselves must be contacted.

Supported DDNS Providers

Not all services and features from all providers are supported, and the list of supported providers is subject to change. SonicOS currently supports the services from providers listed in Dynamic DNS providers:

Dynamic DNS providers

dns.org	SonicOS requires a username, password, Mail Exchanger, and Backup MX to configure DDNS from `Dyndns.org`.
changeip.com	A single, traditional Dynamic DNS service requiring only username, password, and domain name for SonicOS configuration.
no-ip.com	Dynamic DNS service requiring only username, password, and domain name for SonicOS configuration. Also supports hostname grouping.
Yi.org	Dynamic DNS service requiring only username, password, and domain name for SonicOS configuration. Requires that an RR record be created on the `yi.org` administrative page for dynamic updates to occur properly.

Additional Services offered by Dynamic DNS Providers

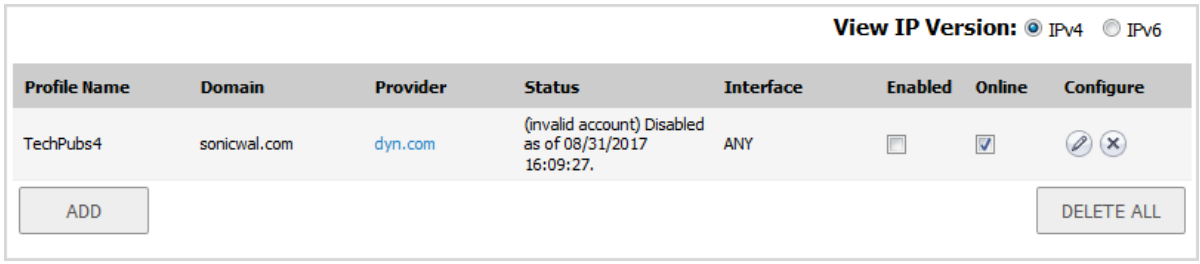
Some common additional services offered by Dynamic DNS providers include:

Wildcards	Allows for wildcard references to sub-domains. For example, if you register `yourdomain.dyndns.org`, your site would be reachable at `*.yourdomain.dyndyn.org`, for example, `server.yourdomain.dyndyn.org`, `www.yourdomain.dyndyn.org`, `ftp.yourdomain.dyndyn.org`.
Mail Exchangers	Creates MX record entries for your domain so that SMTP servers can locate it through DNS and send mail. **NOTE:** Inbound SMTP is frequently blocked by ISPs; check with your provider before attempting to host a mail server.
Backup MX (offered by `dns.org, yi.org`)	Allows for the specification of an alternative IP address for the MX record in the event that the primary IP address is inactive.
Groups	Allows for the grouping of hosts so that an update can be performed once at the group level, rather than multiple times for each member.
Off-Line IP Address	Allows for the specification of an alternative address for your registered host names if primary registered IP is offline.

For information on setting up DDNS Profiles, see Configuring a Dynamic DNS Profile on page 561.

Dynamic DNS Profiles Table

The **Dynamic DNS Profiles** table provides information about configured DDNS profiles.



View IP Version	Allows you to toggle the table between IPv4 and IPv6 DDNS profiles.
Profile Name	Name assigned to the DDNS entry during its creation. This can be any value, and is used only for identification.
Domain	Fully qualified domain name (FQDN) of the DDNS entry.
Provider	DDNS provider with whom the entry is registered.
Status	Last reported/current status of the DDNS entry:

Online	DDNS entry is administratively online. The current IP setting for this entry is shown with a timestamp.
Taken Offline Locally	DDNS entry is administratively offline. If the entry is **Enabled**, the action configured in the **Offline Settings** section of the **Advanced** page of **Add DDNS Profile** is taken.
Abuse	DDNS provider has considered the type or frequency of updates to be abusive. Check with the DDNS provider's guidelines to determine what is considered abuse.
No IP change	Abuse possible. A forced update without an IP address change is considered by some DDNS providers to be abusive. Automatic updates only occur when address or state changes occur. Manual or forced updates should only be made when absolutely necessary, such as when registered information is incorrect.
Disabled	Account has been disabled because of a configuration error or a policy violation. Check the profile's settings, and verify the DDNS account status with the provider.
Invalid Account	Account information provided is not valid. Check the profile's settings, and verify the DDNS account status with the provider.
Network Error	Unable to communicate with the DDNS provider due to a suspected network error. Verify that the provider is reachable and online. Try the action again later.
Provider Error	DDNS provider is unable to perform the requested action at this time. Check the profile's settings, and verify the DDNS account status with the provider. Try the action again later.
Not Donator Account	Certain functions provided from certain provider, such as offline address settings, are only available to paying or donating subscribers. Check with the provider for more details on which services may require payment or donation.

Enabled	When selected, this profile is administratively enabled, and the Security Appliance takes the **Online Settings** action configured on the **Advanced** page of **Add DDNS Profile**. This setting can also be controlled using the **Enable this DDNS Profile** option of the entry's **Add DDNS Profile**. Deselecting this option disable the profiles, and no communications with the DDNS provider occurs for this profile until the profile is again enabled.
Online	When selected, this profile is administratively online. The setting can also be controlled using the **Use Online Settings** option on the entry's **Add DDNS Profile**. Deselecting this option while the profile is enabled takes the profile offline, and the Security Appliance takes the **Offline Settings** action that is configured on the **Advanced** page.
Configure	Includes the **Edit** icon for configuring the DDNS profile settings and the **Delete** icon for deleting the DDNS profile entry.

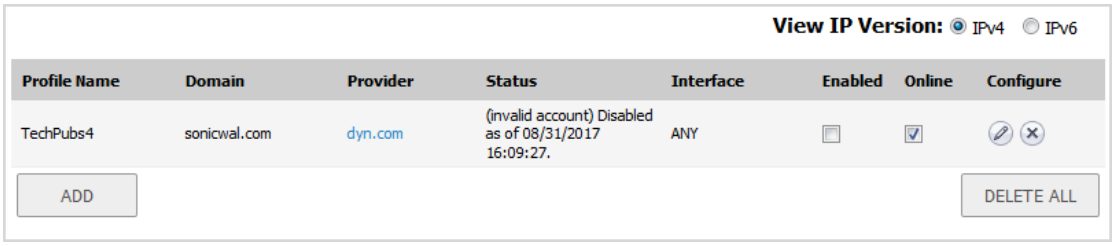
Configuring a Dynamic DNS Profile

For general information on setting up DDNS Profiles, see About Dynamic DNS on page 558.

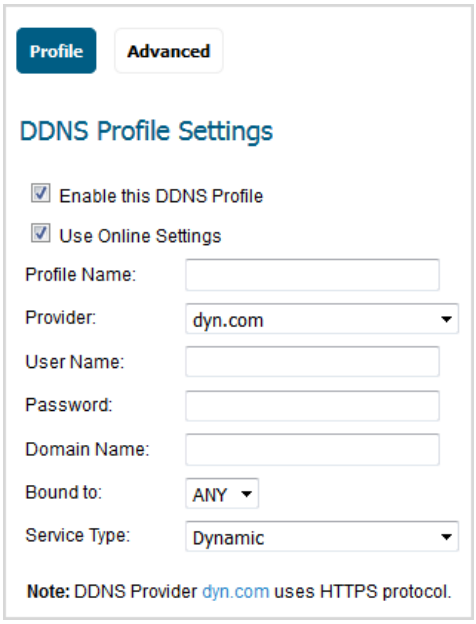
Using any Dynamic DNS service begins with settings up an account with the DDNS service provider (or providers) of your choice. It is possible to use multiple providers simultaneously. Refer to the various providers listed in Dynamic DNS providers. The registration process normally involves a confirmation email from the provider, with a final acknowledgment performed by visiting a unique URL embedded in the confirmation email. After logging in to the selected provider's page, you should visit the administrative link (typically add or manage), and create your host entries. This must be performed prior to attempting to use the dynamic DNS client on SonicOS. The **Network > Dynamic DNS** page provides the settings for configuring the SonicWall Security Appliance to use your DDNS service.

To configure Dynamic DNS on the SonicWall Security Appliance:

1 Navigate to **Network > Dynamic DNS**.



2 Click **Add**. The **Add DDNS Profile** dialog displays.



3 If **Enable this DDNS Profile** is checked, the profile is administratively enabled, and the Security Appliance takes the actions defined in the **Online Settings** section on the **Advanced** page. This option is selected by default.

4 If **Use Online Settings** is checked, the profile is administratively online. This option is selected by default.

5 Enter a name to assign to the DDNS entry in the **Profile Name** field. This can be any value used to identify the entry in the **Dynamic DNS Settings** table. The minimum length is 1 character, and the maximum length is 63 characters.

6 From **Provider**, select the dynamic DNS provider; these providers are described in Dynamic DNS providers. The default is **dyn.com**.

IMPORTANT: You must have created a dynamic service record with the DNS provider you select.

TIP: Not all options are available for all DNS providers. Also, the **Note** at the bottom of the page displays whether the DNS provider uses HTTP or HTTPS protocol along with a link to the provider's website.

7 Enter the username for your DNS-provider account in the **User Name** field. The minimum length is 1 character, and the maximum length is 63 characters.

8 Enter your DNS password in the **Password** field. The minimum length is 1 character, and the maximum length is 31 characters.

9 Enter the fully qualified domain name (FQDN) of the host name you registered with the DNS provider in the **Domain Name** field. Make sure you provide the same host name and domain as you configured. The minimum length is 1 character, and the maximum length is 63 characters.

10 Optionally, to assign this DDNS profile to a specific WAN interface, select that WAN interface from **Bound to**. If you are configuring multiple-WAN load balancing, this option allows you to advertise a predictable IP address to the DDNS service. By default, this is set to **ANY**, which means the profile is free to use any of the WAN interfaces on the Security Appliance.

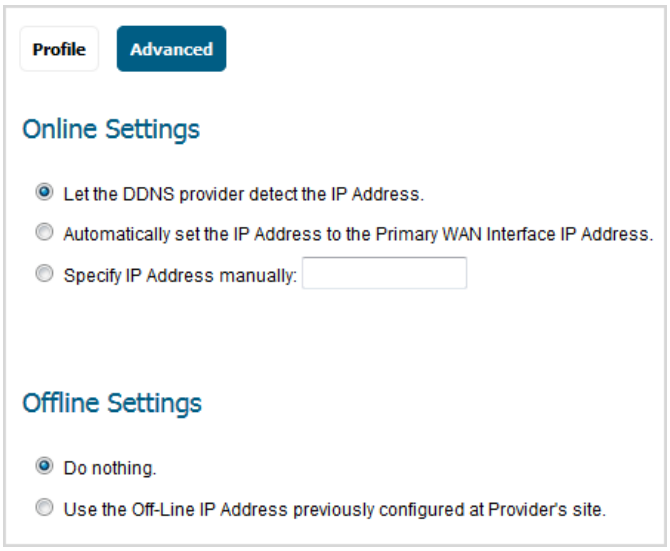
11 If you selected **dyn.com** for **Provider**, go to Step 13.

12 When using `dyn.org`, select the service type that corresponds to your type of service from **Service Type**:

Dynamic	Free Dynamic DNS service. This is the default.
Custom	Managed primary DNS solution that provides a unified primary/secondary DNS service and a Web-based interface. Supports both dynamic and static IP addresses.
Static	Free DNS service for static IP addresses.

13 Click **Advanced**.

TIP: You can typically leave the default settings on this page.



14 The **On-line Settings** section provides control over what address is registered with the dynamic DNS provider. Choose:

Let the DDNS provider detect the IP Address	The Security Appliance allows the DNS provider to specify the IP address **NOTE:** IPv4 only. This option is selected by default.
Automatically set IP Address to the Primary WAN Interface IP Address	Causes the Security Appliance to assert its WAN IP address as the registered IP address, overriding auto-detection by the dynamic DNS server. Useful if detection is not working correctly. This option is selected by default. **NOTE:** In IPv6: This option is selected by default.
Specify IP Address manually	Allows for the IP address to be registered to be manually specified and asserted.

15 The **Off-line Settings** section controls what IP address is registered with the dynamic DNS service provider if the dynamic DNS entry is taken off-line locally (disabled) on the Security Appliance. Choose:

Do nothing	Allows the previously registered address to remain current with the dynamic DNS provider. This option is selected by default.
Use the Off-Line IP address previously configured at Providers site	If your provider supports manual configuration of Off-Line Settings, you can select this option to use those settings when this profile is taken administratively offline.

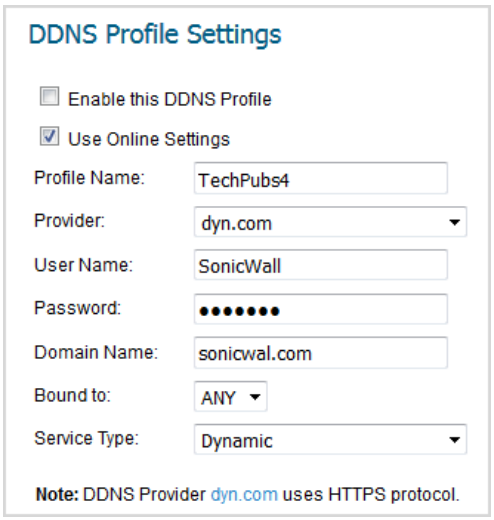
16 Click **OK**.

Editing a DDNS Profile

To edit a DDNS profile:

1 Navigate to **Network > Dynamic DNS**.

2 In the Dynamic DNS Profiles table, click the **Edit** icon of the profile. The Edit DDNS Profile dialog displays.



3 Make changes; for a description of the options, follow the instructions for Configuring a Dynamic DNS Profile on page 561.

4 Click **OK**.

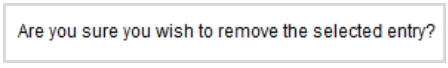
Deleting DDNS Profiles

You can delete one or all DDNS profiles.

To delete a DDNS profile.

1 Navigate to **Network > Dynamic DNS**.

2 Click the **Delete** icon of the profile to be deleted. A confirmation message displays:

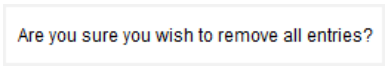


3 Click **OK**.

To delete all DDNS entries:

1 Navigate to **Network > Dynamic DNS**.

2 Click **Delete All**. A confirmation message displays:



3 Click **OK**.

30

Configuring AWS Credentials

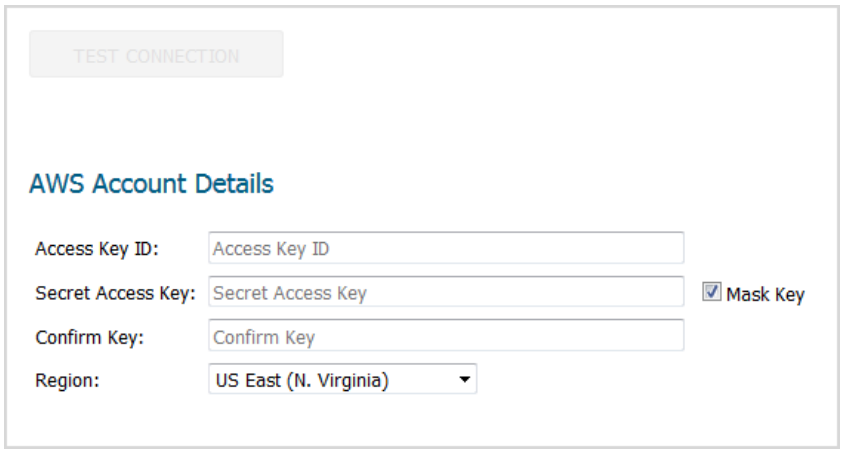
IMPORTANT: To use the SonicOS-AWS integration feature, you must:

- Be registered with Amazon Web Services (AWS).
- Have an AWS Identity and Access Management (IAM) User's Access Key ID and Access Key.
- Be familiar with IAM Best Practices.

Topics:

- **Network > AWS Configuration** on page 565
  - About AWS on page 565
  - Configuring AWS on page 566
  - Troubleshooting the Connection on page 567

Network > AWS Configuration



Topics:

- About AWS on page 565
- Configuring AWS on page 566
- Troubleshooting the Connection on page 567

About AWS

SonicOS integration with Amazon Web Service (AWS) enables you to:

- Store your logs on the AWS CloudWatch Logs service monitor and troubleshoot your systems and applications.
- Use AWS-hosted analysis tools such as ElasticSearch and Kibana.

To integrate SonicOS with AWS and allow the Security Appliance to communicate with the various application programming interfaces (APIs) of AWS, you need to:

1 Provide AWS security credentials; see Configuring AWS on page 566.

2 Create AWS Objects, such as Address Objects and Address Groups, that correspond to AWS EC2 Instances For further information about creating AWS Objects, see *SonicOS 6.5 Policies*.

3 Create VPN connections from the Security Appliance to the AWS Virtual Private Clouds (VPCs). For further information about creating VPN connections, see *SonicOS 6.5 Connectivity*.

4 Create a Log Stream and enable logging. For more information about logging to Amazon CloudWatch Logs, see *SonicOS 6.5 Logs and Reporting*.

Configuring AWS

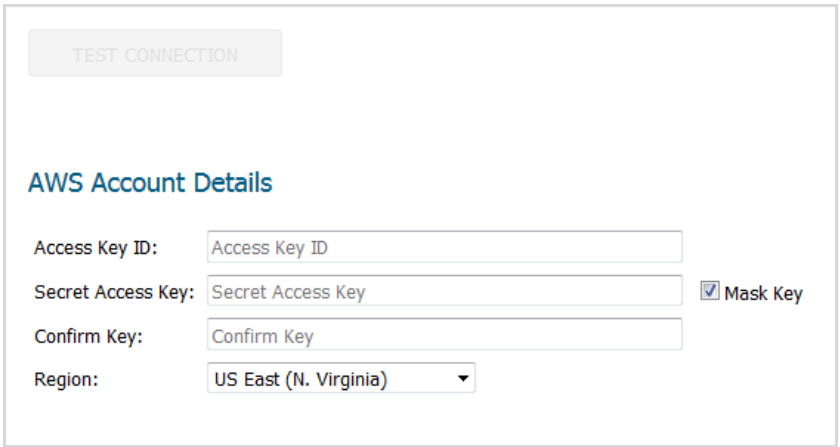
NOTE: To configure SonicOS to allow TLS v1.0 for AWS, contact SonicWall Support.

To configure AWS:

1 Ensure you have:

- Registered with Amazon Web Services (AWS).
- An AWS Identity and Access Management (IAM) User's Access Key ID and Secret Access Key.
- Familiarity with IAM Best Practices.

2 Navigate to **MANAGE | System Setup > Network > AWS Configuration**.



3 Enter the AWS Access Key ID in the **Access Key ID** field. The AWS Access Key ID is used by the Security Appliance to access AWS APIs. This option is not selected by default.

4 To mask the key for security, ensure **Mask Key** is selected. This option is selected by default.

5 Enter the AWS Secret Access Key in the **Secret Access Key** field. The Secret Access Key is used by the Security Appliance to access AWS APIs. If **Mask Key** is selected, the field is a series of bullets.

6 Reenter the AWS Secret Access Key in the **Confirm Key** field.

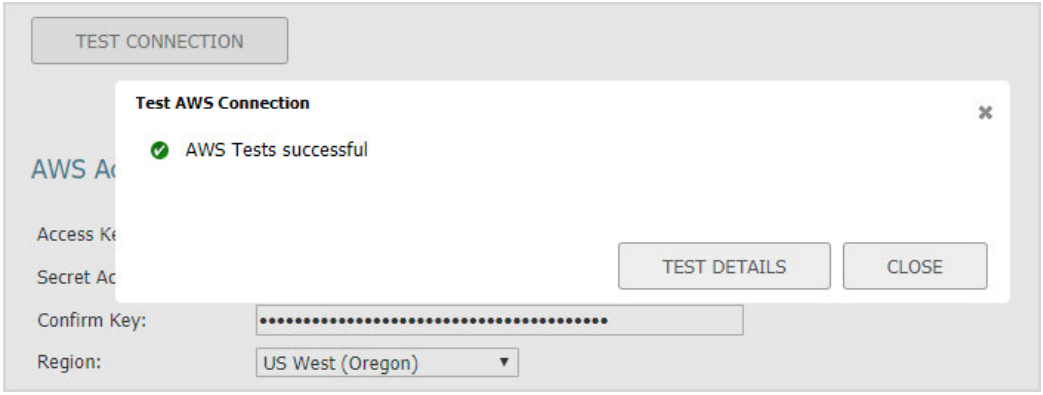
7 From **Region**, select the default region used to initialize the **MANAGE | Policies > Objects > AWS Objects** and **MANAGE | Connectivity > VPN > AWS VPN** pages. The default is **US East (N. Virginia)**.

IMPORTANT: If the default region is the region used when sending Security Appliance event logs to AWS CloudWatch Logs, it is affected by changes on the **MANAGE | Logs & Reporting > Log Settings > AWS Logs** page.

8 Click **ACCEPT**. **TEST CONNECTION** becomes available.

CAUTION: It is important to test the connection and configuration at this time as any error at this point will result in issues later.

9 To test validity of the credentials and that Security Appliance can successfully communicate with AWS, click **TEST CONNECTION**. Several tests are run to test the credentials and the connection to AWS. The results display.



TIP: If there were problems with the test, see Troubleshooting the Connection on page 567

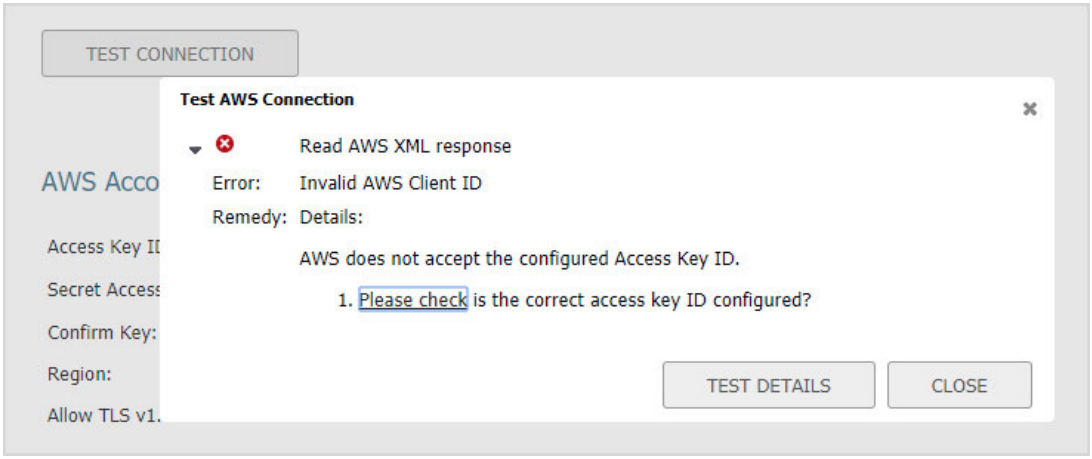
10 Click **Close**.

Troubleshooting the Connection

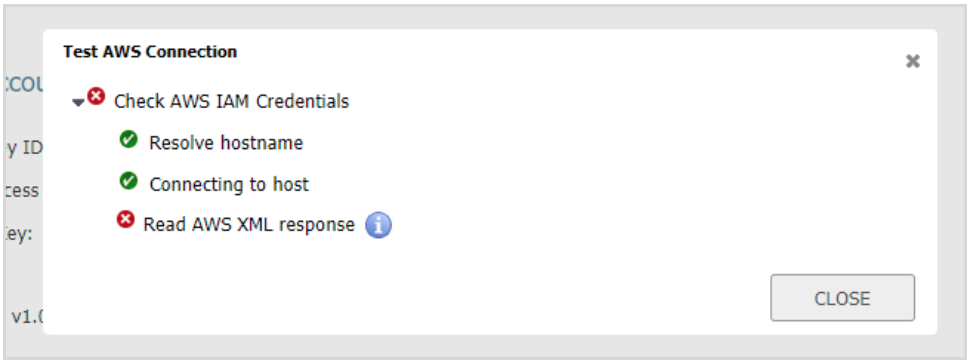
This example shows an invalid Access Key ID.

To troubleshoot the connection:

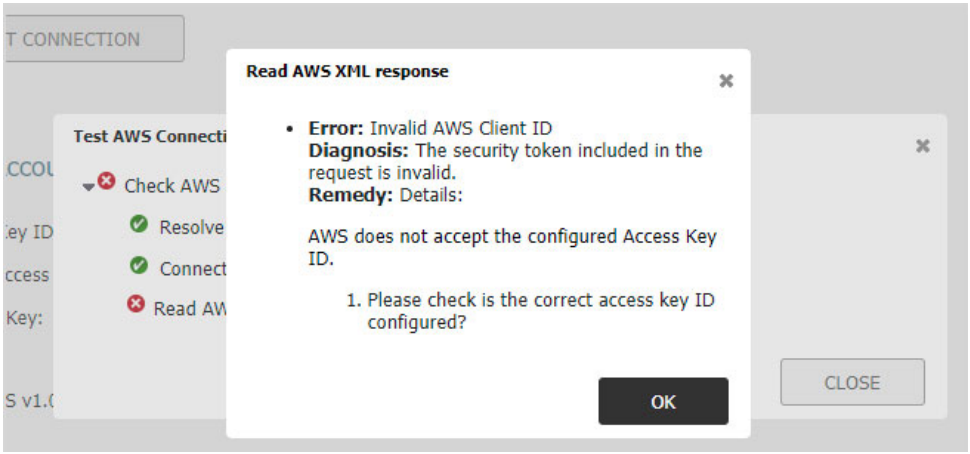
1 Click **TEST CONNECTION**. The results display.



2 Click **TEST DETAILS**. More information displays.



3 To display more information, click the **Information** icon. Another popup displays.



4 Note the diagnosis.

5 Click **OK**.

6 Click **CLOSE**.

7 Correct the problem described in **Diagnosis**.

8 Click **TEST CONNECTION**.

9 Repeat Step 1 through Step 8 until you solve the problem(s).

10 Click **CLOSE**.

Part 5

System Setup | SD-WAN

- About SD-WAN
- Configuring SD-WAN Groups
- Configuring Performance Probes
- Configuring Performance Class Objects
- Configuring Path Selection Profiles
- Configuring SD-WAN Route Policies
- Monitoring SD-WAN
- Viewing SD-WAN Route Policy Connections

31

About SD-WAN

SD-WAN (Software-Defined Wide Area Network) provides software-based control over wide area network (WAN) connections. SonicOS SD-WAN offers these features:

- SD-WAN Interface Groups
  - WAN and VPN
  - Scalable from one to N interfaces
- Dynamic path selection based on:
  - Latency, jitter, and/or packet loss
  - User-defined thresholds for quality assessment
- Application-aware routing
- Path Performance Probes for metrics
- Connection-based traffic distribution
- Automatic connection failover over VPN
- Provisioning and management (GMS and Capture Security Center)

SD-WAN is best used for specific traffic types and/or applications requiring dynamically chosen optimal destination interfaces depending on how the network paths are behaving. To operate well, each application has a certain requirement from the network path. For example the network quality for VoIP to operate well requires the optimal latency be 100 ms or less while a latency of 150 ms or higher results in choppy calls. SD-WAN helps in such scenarios by first dynamically measuring the various network performance metrics, such as latency, jitter and packet loss, on multiple network paths. SD-WAN then compares these metrics with the performance threshold for a particular traffic flow and determines the optimal network that meets the flow's network quality accordingly.

Topics:

- Configuring SD-WAN Groups on page 571
- Configuring Performance Probes on page 575
- Configuring Performance Class Objects on page 581
- Configuring Path Selection Profiles on page 586
- Configuring SD-WAN Route Policies on page 591
- Monitoring SD-WAN on page 597
- Viewing SD-WAN Route Policy Connections on page 599

32

Configuring SD-WAN Groups

Topics:

- SD-WAN Groups on page 571
  - Creating an SD-WAN Group on page 572
  - Editing an SD-WAN Group on page 572
  - Deleting SD-WAN Groups on page 573

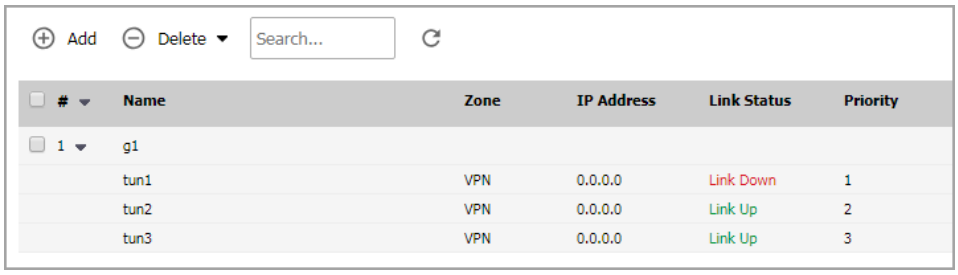
SD-WAN Groups

SD-WAN supports physical and virtual WAN interface types as well as VPN Numbered and Unnumbered Tunnel Interface instances, all choices provided during SD-WAN Group Congifuration.

Numbered Tunnel Interface groups are more evolved versions of the route-based implementations, and they are also far more intuitive to work with. However, they do scale poorly as they borrow from other actual entries in the interface table. Unnumbered Tunnel Interface configurations can meet the necessary scalability requirements because they are not tied to an interface table entry.

SD-WAN Groups are logical groups of interfaces that can be used for load-balancing as well as dynamic path selection based on the performance criterion through each interface path. You can create your own custom groups. For a description of SonicOS SD-WAN and its features, see About SD-WAN on page 570.

The **SD-WAN Groups** page displays the custom pool of interfaces used for optimized and resilient traffic flow.



Name	Name of the SD-WAN group.
Zone	The zone of the interface member: • WAN • VPN
IP Address	IP address of the interface.
Link Status	Indicates whether the link is: • **Link Up** (green) • **Link Down** (red)
Priority	Priority of the interface in the group.
Configure	Contains the **Edit** and **Delete** icons for the group.

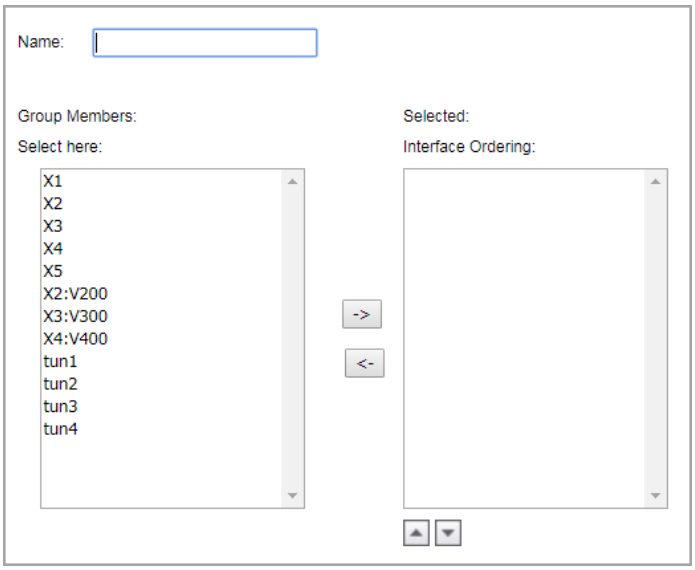
Creating an SD-WAN Group

You can create multiple SD-WAN Groups to meet your requirements.

To add an SD-WAN group:

1 Navigate to **MANAGE | System Setup > SD-WAN > SD-WAN Groups**. Note the Unnumbered Tunnel Interfaces included as an interface configuration option in an SD-WAN group.

2 Click the **Add** icon. The **Add SD-WAN Group** dialog displays.



3 Enter a descriptive name in the **Name** field.

4 Select one or more interfaces from the **Group Members Select here** list. Member interfaces can be only WAN, Numbered or Unnumbered Tunnel Interfaces.

IMPORTANT: An interface cannot be a member of more than one SD-WAN group.

5 Click the **Right Arrow** to move the selected interfaces to the **Selected Interface Ordering** column.

6 To change the priority of the selected group members:

a Select the interface.

b Click the **Up Arrow** or **Down Arrow**.

7 Repeat Step 6 for each interface to prioritize.

8 Click **OK**.

Editing an SD-WAN Group

To edit an SD-WAN group:

1 Navigate to **MANAGE | System Setup > SD-WAN > SD-WAN Groups**.

2 Click the **Edit** icon of the group to edit. The **Edit SD-WAN group** dialog displays, which is the same as the **Add SD-WAN group** dialog.

3 Make changes as described in Creating an SD-WAN Group on page 572.

4 Click **OK**.

Deleting SD-WAN Groups

You can delete one, multiple, or all SD-WAN groups.

Topics:

- Deleting an SD-WAN Group on page 573
- Deleting Multiple SD-WAN Groups on page 573
- Deleting All SD-WAN Groups on page 574

Deleting an SD-WAN Group

To delete an SD-WAN group:

1 Navigate to **MANAGE | System Setup > SD-WAN > SD-WAN Groups**.

2 Click the **Delete** icon of the group to delete. A confirmation message displays.



3 Click **OK**.

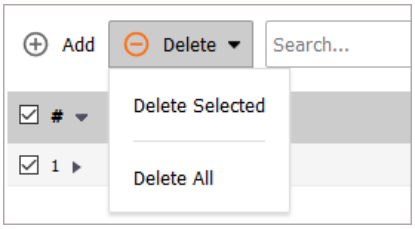
Deleting Multiple SD-WAN Groups

To delete multiple SD-WAN groups:

1 Navigate to **MANAGE | System Setup > SD-WAN > SD-WAN Groups**.

2 Select the groups to delete.

3 From **Delete** above the **SD-WAN Groups** table, select **Delete Selected**.



A confirmation message displays.



4 Click **OK**.

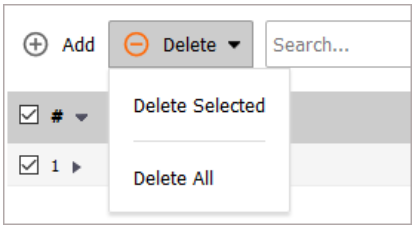
Deleting All SD-WAN Groups

To delete all SD-WAN groups:

1 Navigate to **MANAGE | System Setup > SD-WAN > SD-WAN Groups**.

2 Select the checkbox in the header of the **SD-WAN Groups** table. All the groups are selected.

3 From **Delete** above the **SD-WAN Groups** table, select **Delete All**.



A confirmation message displays.



4 Click **OK**.

33

Configuring Performance Probes

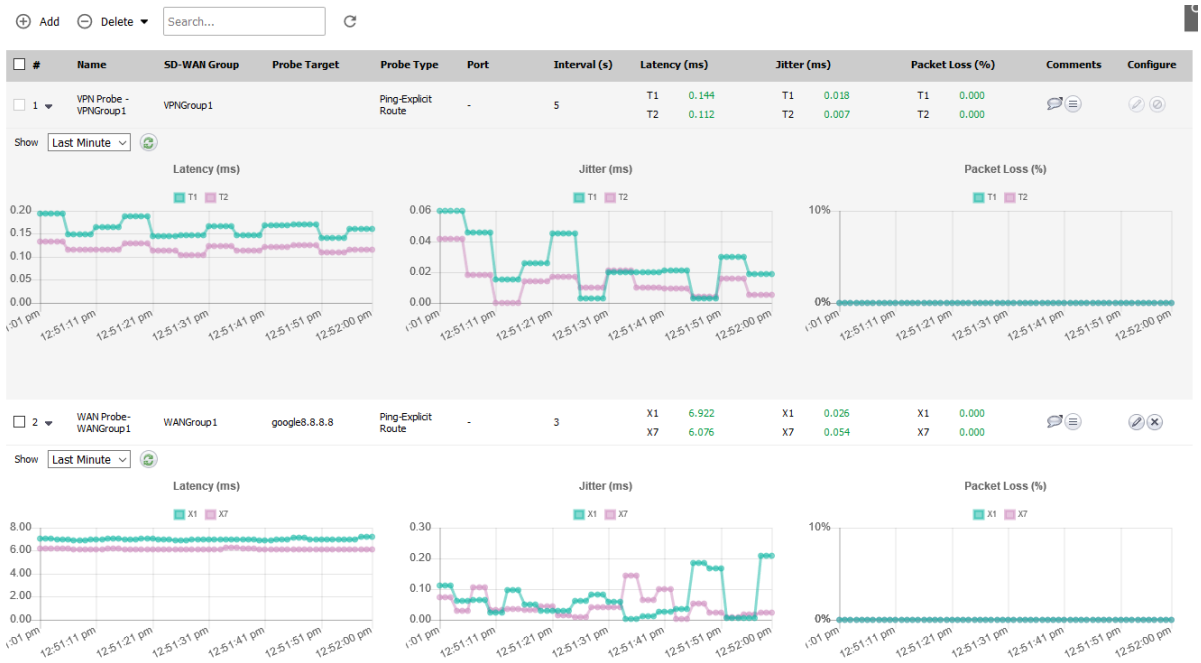
Topics:

- About Performance Probes on page 575
- Configuring Performance Probes on page 578

About Performance Probes

Network path performance metrics are determined using SD-WAN performance probes, which are similar to Network Monitor Probes. SonicOS supports ICMP and TCP probe types. A SD-WAN performance probe can be used by multiple Path Selection profiles (for further information, see Configuring Path Selection Profiles on page 586; for a description of SonicOS SD-WAN and its features, see About SD-WAN on page 570.).

The **MANAGE | System Setup > SD-WAN > Performance Probes** page shows the dynamic performance data (latency/jitter/packet loss) and probe status for each path (interface) in the SD-WAN group, in both tabular and graphic displays. The display can show data for the last minute (default), last day, last week, or last month.

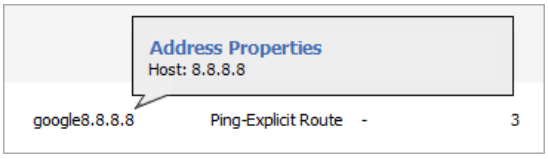


#	Number of the probe. The **Collapse/Expand** icon toggles the display of the graphs.
Name	Name of the SD-WAN performance probe.

SD-WAN Group Name of the SD-WAN group associated with the SD-WAN performance probe. Mousing over the entry displays properties about the SD-WAN group:



Probe Target Target address object of the SD-WAN performance probe. Mousing over the entry displays the host address:



Probe Type Type of performance probe:

- **Ping – Explicit Route**
- **TCP – Explicit Route**

Note: When - **TCP – Explicit Route** is selected along with the **RST Response Counts as Miss** field, the **Port** field also becomes available.

Port Port for the SD-WAN performance probe. The minimum/maximum values are 1 to 65535.

NOTE: Ports are displayed only for **TCP - Explicit Route** probe types. A hyphen (**–**) displays for **Ping - Explicit Route** probe types.

Interval (s) Time between SD-WAN performance probes, in seconds.

Latency (ms) Round trip delay for the probes sent through a particular path/interface to reach the probe target and acknowledge back, in milliseconds. This is also displayed as a graph below the probe's entry in the **Performance Probe** table.

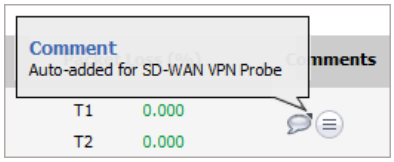
Jitter (ms) Variation in the latency measurements for the probes through a particular path/interface, in milliseconds. This is also displayed as a graph below the probe's entry in the **Performance Probe** table.

Packet Loss (%) Percentage of probes that are missed of the probes sent through a particular path/interface. This is also displayed as a graph below the probe's entry in the **Performance Probe** table.

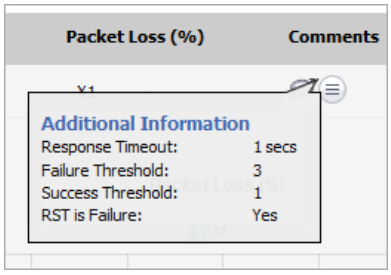
Comments

Displays the:

- **Comment** icon; mousing over the icon displays the comment entered when the performance probe was configured.



- **Notes** icon; mousing over the displays statistics about the performance probe.



  - **Response Timeout**: Maximum delay for a response.
  - **Failure Threshold**: Number of missed intervals before the probe state is set to DOWN.
  - **Success Threshold**: Number of successful intervals until the probe state is set to UP.
  - **RST is Failure**: For probe types of **TCP – Explicit Route**, whether RST responses count as misses.

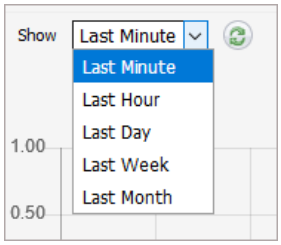
Configure

Displays the **Edit** and **Delete** icons for the SD-WAN performance probe.

NOTE: These icons are dimmed for auto-added probes, such as those for SD-WAN groups containing a numbered tunnel interface. These probes are deleted automatically when its SD-WAN group is deleted.

Show

Controls the display interval of graph data:



The default is **Last Minute**.

Graphs

Graphical depictions of the columnar data:

- **Latency (ms)**
- **Jitter (ms)**
- **Packet Loss (%)**

TIP: You can toggle the display of the graphs for each performance probe by clicking the **Collapse/Expand** icon by the probe number.

TIP: These graphs display a snapshot of historical data for the time frame specified in the **Show** drop-down menu. The graphs in the **MONITOR | Appliance Health > SD-WAN Monitor** and **MANAGE | System Setup > SD-WAN > SD-WAN Monitor** pages display moving real-time data. For information about the **SD-WAN Monitor** graphs, see *SonicOS 6.5 Monitoring* or Monitoring SD-WAN on page 597.

When configuring SD-WAN Performance Probes, default row(s) are created for each of the interfaces used by the SD-WAN Group in the Network Monitor Policies screen.

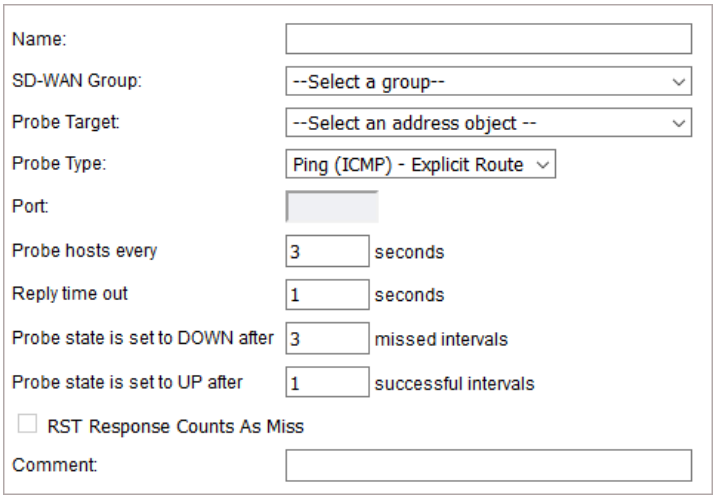
Configuring Performance Probes

IMPORTANT: A Performance Probe is created automatically for a SD-WAN Group containing a VPN numbered tunnel interface. You do not need to create an additional performance probe.

To add a performance probe for non-VPN SD-WAN Groups:

1 Navigate to **MANAGE | System Setup > SD-WAN > Performance Probes**.

2 Click the **Add** icon. The **Add SD-WAN Performance Probe** dialog displays.



3 Enter a meaningful name in the **Name** field.

4 Select a SD-WAN group from **SD-WAN Group**.

5 Select an address object from **Probe Target**.

6 From **Probe Type**, select:

- **Ping (ICMP) - Explicit Route** (default); go to Step 8.
- **TCP - Explicit Route**; the **Port** field becomes available.

7 Enter the port number of the explicit route in the **Port** field.

8 Enter the interval between probes in the **Probe hosts every … seconds** field. The minimum is 1 second, the maximum is 3600 seconds, and the default is **3** seconds.

TIP: The probe interval must be greater than the reply timeout.

9 Enter the maximum delay for a response in the **Reply time out … seconds** field. The minimum is 1 second, the maximum is 60 seconds, and the default is **1** second.

10 Enter the maximum number of missed intervals before the performance probe is set to the DOWN state in the **Probe state is set to DOWN after … missed intervals** field. The minimum number is 1, the maximum is 100, and the default is **3**.

11 Enter the maximum number of successful intervals before the performance probe is set to the UP state in the **Probe state is set to UP after … successful intervals** field. The minimum number is 1, the maximum is 100, and the default is **1**.

12 If you selected **TCP - Explicit Route** for **Probe Type**, the **RST Response Counts As Miss** option becomes available. Select the option to count RST responses as missed intervals. This option is not selected by default.

13 Optionally, enter a comment in the **Comment** field.

14 Click **ADD**.

15 Repeat Step 3 through Step 14 to add more probes.

16 Click **CLOSE**.

Editing an SD-WAN Performance Probe

To edit an SD-WAN performance probe:

1 Navigate to **MANAGE | System Setup > SD-WAN > Performance Probes**.

2 Click the **Edit** icon of the performance probe to edit. The **Edit SD-WAN Performance Probe** dialog displays, which is the same as the **Add SD-WAN Performance Probe** dialog.

3 Make changes as described in Configuring Performance Probes on page 578.

4 Click **OK**.

Deleting SD-WAN Performance Probes

You can delete one, multiple, or all SD-WAN performance probes.

Topics:

- Deleting a Performance Probe on page 579
- Deleting Multiple Performance Probes on page 579
- Deleting All Performance Probes on page 580

Deleting a Performance Probe

To delete a SD-WAN performance probe:

1 Navigate to **MANAGE | System Setup > SD-WAN > Performance Probes**.

2 Click the **Delete** icon of the performance probe to delete. A confirmation message displays.



3 Click **OK**.

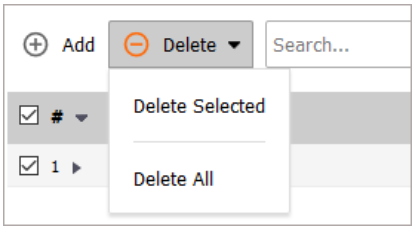
Deleting Multiple Performance Probes

To delete multiple SD-WAN performance probes:

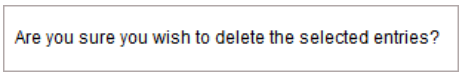
1 Navigate to **MANAGE | System Setup > SD-WAN > Performance Probes**.

2 Select the performance probes to delete.

3 From **Delete** above the **SD-WAN Performance Probe** table, select **Delete Selected**.



A confirmation message displays.



4 Click **OK**.

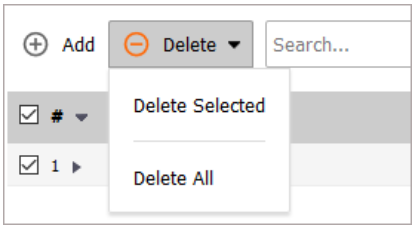
Deleting All Performance Probes

To delete all SD-WAN performance probes:

1 Navigate to **MANAGE | System Setup > SD-WAN > Performance Probes**.

2 Select the checkbox in the header of the **SD-WAN Performance Probe** table. All the performance probes are selected.

3 From **Delete** above the **SD-WAN Performance Probe** table, select **Delete All**.



A confirmation message displays.



4 Click **OK**.

34

Configuring Performance Class Objects

Topics:

- About Performance Class Objects on page 581
- Configuring Performance Class Objects on page 583
- Editing a Performance Class Object on page 584
- Deleting SD-WAN Performance Class Objects on page 584

About Performance Class Objects

A Performance Class specifies the performance criterion for selecting the optimal path. It could be the:

- Best latency/jitter/packet loss among the existing paths.
- Performance Class Object that defines the metric thresholds for latency, jitter and packet loss.

You use SD-WAN Performance Class Objects to configure the desired performance characteristics for the application/traffic categories. These objects are used in the Path Selection Profile to automate the selection of paths based on these metrics. (For a description of SonicOS SD-WAN and its features, see About SD-WAN on page 570.)

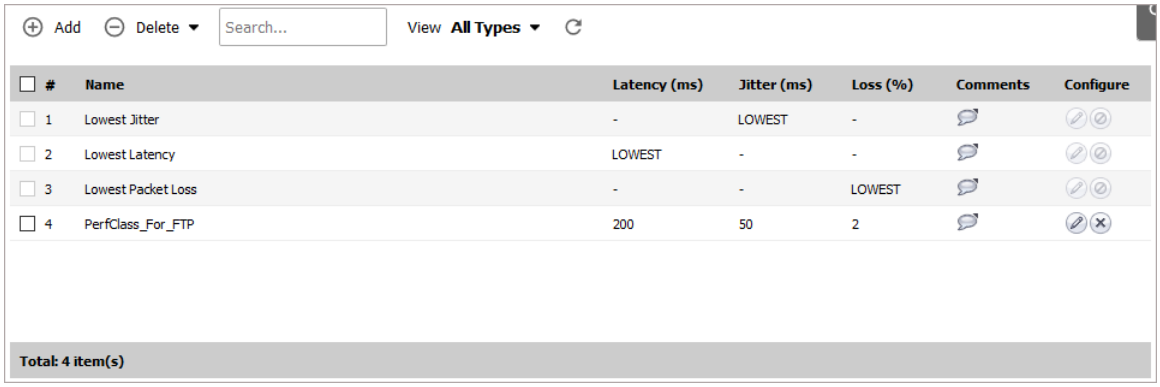
These are the default Performance Class Objects:

- **Lowest Jitter**
- **Lowest Latency**
- **Lowest Packet Loss**

NOTE: These default Performance Class Objects cannot be edited or deleted.

You configure custom performance thresholds that best meet the needs of your application/traffic categories with custom Performance Class Objects. Starting in SonicOS 6.5.4.5, you can include or exclude the Latency, Jitter, or Packet Loss attributes in your custom object, although you cannot exclude all three attributes in the same object. When excluded, the value of that attribute is not used as a criterion or threshold when determining whether a particular path is qualified or not. For example, if you want to evaluate a particular path only on the Latency attribute but you don't care about the other attributes, you can include Latency and exclude Jitter and Packet Loss in your custom object.

NOTE: When SonicOS is upgraded to 6.5.4.5 from a previous version, all existing custom performance class objects are upgraded with the Include option enabled for Latency, Jitter, and Packet Loss.



Name Name of the Performance Class Object

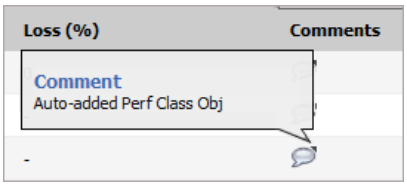
Latency (ms) Threshold time for the round trip delay for the probes sent through a particular path/interface to reach the probe target and acknowledge back, in milliseconds, in milliseconds. For the **Lowest Latency** Performance Class Object, the time is always **LOWEST**; for the other default Performance Class Objects, a hyphen (**–**) displays.

Jitter (ms) Threshold variation in the latency measurements for the probes through a particular path/interface, in milliseconds. For the **Lowest Jitter** Performance Class Object, the time is always **LOWEST**; for the other default Performance Class Objects, a hyphen (**–**) displays.

Loss (%) Threshold percentage of probes that are missed of the probes sent through a particular path/interface. For the **Lowest Packet Loss** Performance Class Object, the percentage is always **LOWEST**; for the other default Performance Class Objects, a hyphen (**–**) displays.

Comments Displays the **Comment** icon; mousing over the icon displays the comment entered when the Performance Class Object was configured.

NOTE: When mousing over the **Comment** icon for a default Performance Class Object, the comment is the same:



Configure Displays the **Edit** and **Delete** icons for the SD-WAN Performance Class Object.

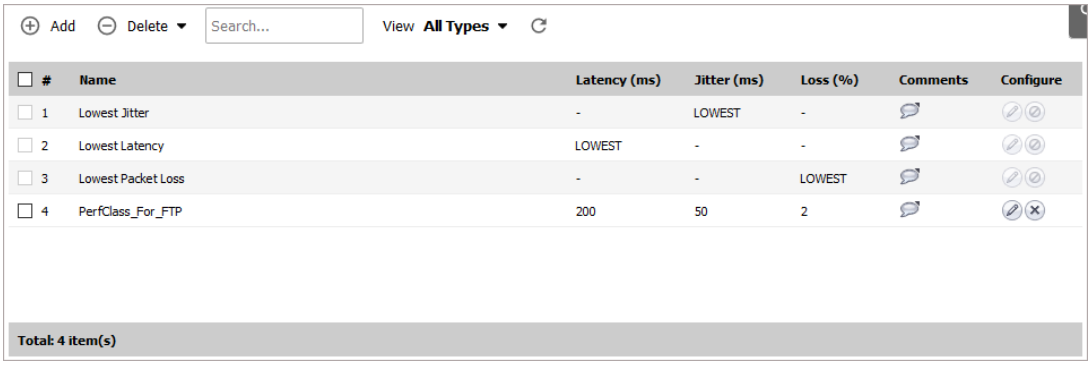
NOTE: You cannot edit or delete the default Performance Class Objects, and these icons are dimmed.

Total Displays the total number of Performance Class Objects in the **Performance Class Object** table.

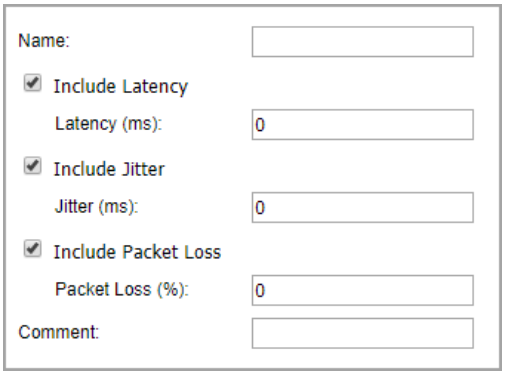
Configuring Performance Class Objects

To add a Performance Class Object:

1 Navigate to **MANAGE | System Setup | SD-WAN > Performance Class Objects**.



2 Click the **Add** icon. The **Add Performance Class Object** dialog displays.



3 Enter a meaningful name in the **Name** field.

4 Select the **Include Latency** checkbox to include the performance class latency attribute for this object, or clear the checkbox to exclude the latency attribute. This option is selected by default.

5 If **Include Latency** is selected, enter the acceptable latency, in milliseconds, in the **Latency (ms)** field. The minimum is 0 milliseconds, the maximum is 1000 milliseconds, and the default is **0** milliseconds.

6 Select the **Include Jitter** checkbox to include the performance class jitter attribute for this object, or clear the checkbox to exclude the jitter attribute. This option is selected by default.

7 If **Include Jitter** is selected, enter the acceptable jitter, in milliseconds, in the **Jitter (ms)** field. The minimum is 0 milliseconds, the maximum is 1000 milliseconds, and the default is **0** milliseconds.

8 Select the **Include Packet Loss** checkbox to include the performance class packet loss attribute for this object, or clear the checkbox to exclude the packet loss attribute. This option is selected by default.

9 If **Include Packet Loss** is selected, enter the acceptable percentage of packet loss in the **Packet Loss (%)** field. The minimum is 0, the maximum is 100, and the default is **0**.

> **NOTE:** You cannot exclude all three attributes (Latency, Jitter, Packet Loss) in the same object.

10 Optionally, enter a comment in the **Comment** field.

11 Click **OK**.

12 To configure more Performance Class Objects, repeat Step 3 through Step 11 for each Performance Class Object.

13 When finished, click **CANCEL**.

Editing a Performance Class Object

To edit a SD-WAN performance probe:

1 Navigate to **MANAGE | System Setup | SD-WAN > Performance Class Objects**.

2 Click the **Edit** icon of the performance probe to edit. The **Edit SD-WAN Performance Class Object** dialog displays, which has the same options as the **Add SD-WAN Performance Class Object** dialog.

3 Make changes as described in Configuring Performance Class Objects on page 583.

4 Click **OK**.

Deleting SD-WAN Performance Class Objects

You can delete one, multiple, or all SD-WAN Performance Class Objects, but you cannot delete the default Performance Class Objects.

Topics:

- Deleting a Performance Class Object on page 584
- Deleting Multiple Performance Class Objects on page 584
- Deleting All Performance Class Objects on page 585

Deleting a Performance Class Object

To delete a SD-WAN Performance Class Object:

1 Navigate to **MANAGE | System Setup | SD-WAN > Performance Class Objects**.

2 Click the **Delete** icon of the Performance Class Object to delete. A confirmation message displays.



3 Click **OK**.

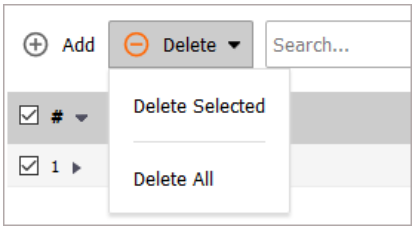
Deleting Multiple Performance Class Objects

To delete multiple SD-WAN Performance Class Objects:

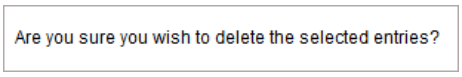
1 Navigate to **MANAGE | System Setup | SD-WAN > Performance Class Objects**.

2 Select the Performance Class Objects to delete.

3 From **Delete** above the **SD-WAN Performance Class Object** table, select **Delete Selected**.



A confirmation message displays.



4 Click **OK**.

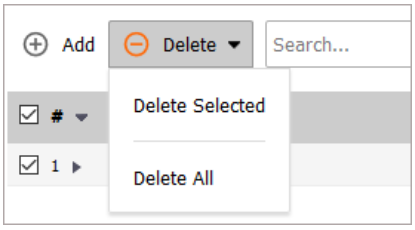
Deleting All Performance Class Objects

To delete all SD-WAN Performance Class Objects:

1 Navigate to **MANAGE | System Setup | SD-WAN > Performance Class Objects**.

2 Select the checkbox in the header of the **SD-WAN Performance Class Object** table. All the Performance Class Objects are selected.

3 From **Delete** above the **SD-WAN Performance Class Object** table, select **Delete All**.



A confirmation message displays.



4 Click **OK**.

35

Configuring Path Selection Profiles

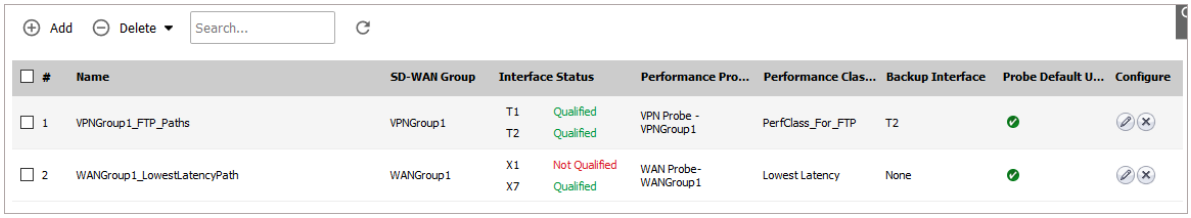
Topics:

- About Path Selection Profiles on page 586

About Path Selection Profiles

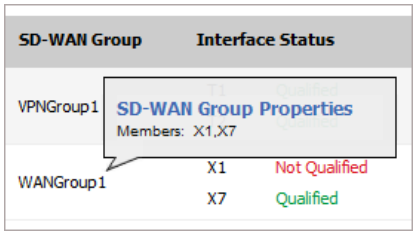
Path Selection Profiles (PSPs) are the settings that help to determine the network path or interface that satisfies a specific network performance criteria, from a pool of available network paths/interfaces. A qualified path or interface meets the performance class criteria.

The dynamic path selection mechanism is implemented using the PSP settings when associated with Policy-based Routes (PBR). When more than one network path meets the criterion (as per the performance class in the PSP), then traffic is load balanced among the qualified network paths/interfaces. When associated with a policy-based routing policy, a path selection profile helps select the optimal path among the SD-WAN interfaces for the application/service.



Name Name of the Path Selection Profile.

SD-WAN Group SD-WAN interface group to which the profile applies. Mousing over the SD-WAN group name displays the members of the group:

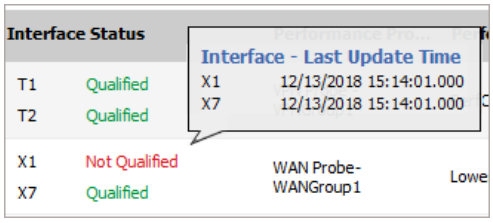


Interface Status

Status of the members of the SD-WAN interface group:

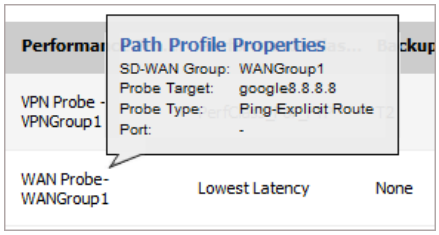
- Interface member
- Member Status:
  - **Qualified** (green)
  - **Not Qualified** (red)

Mousing over the status displays a popup with last-update time:



Performance Probe

Performance Probe used by the Path Selection Profile. Mousing over the name of the performance probe displays a popup with the properties of the performance probe:



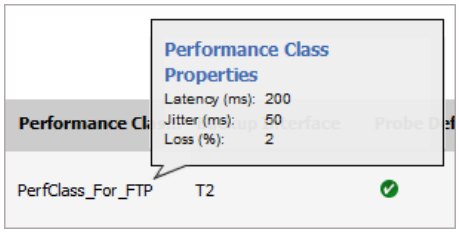
- **SD-WAN Group**: Name of the SD-WAN interface group associated with the group.
- **Probe Target**: Address object of the probe.
- **Probe Type**: Type of probe:
  - **Ping (ICMP) - Explicit Route**
  - **TCP - Explicit Route**
- **Port**: Port associated with the Probe Type; displayed only for a probe type of **TCP - Explicit Route**.

Performance Class Object

Performance Class Object used by the Path Selection Profile:

- **Lowest Latency**
- **Lowest Jitter**
- **Lowest Packet Loss**
- Custom Performance Class Object

Mousing over the entry displays a popup displaying the Performance Class Object's properties:



When Latency, Jitter, or Packet Loss attribute(s) are excluded in a Custom Performance Class Object, those attributes are not used by the Path Selection Profile to evaluate the optimal path.

Backup Interface

Indicates the interface chosen when none of the SD-WAN group interfaces meet the performance criteria. If a backup interface was not chosen, **None** displays.

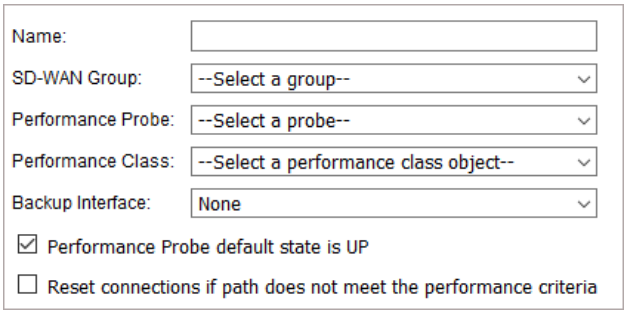
Probe Default UP	Indicates whether the default state of the performance probe is: • UP (green checkmark icon) • DOWN (red X icon)
Configure	Displays the **Edit** and **Delete** icons for the Path Selection Profile.

Configuring Path Selection Profiles

To add a Path Selection Profile:

1 Navigate to **MANAGE | System Setup | SD-WAN > Path Selection Profiles**.

2 Click the **Add** icon. The **Add Path Selection Profile** dialog displays.



3 Add a meaningful name in the **Name** field.

4 From **SD-WAN Group**, select the SD-WAN interface group to which the profile applies.

5 From **Performance Probe**, select the probe to use in the profile.

6 From **Performance Class**, select the Performance Class Object for the dynamic selection of the optimal network path:

- **Lowest Latency**
- **Lowest Jitter**
- **Lowest Packet Loss**
- Custom Performance Class Object

7 From **Backup Interface**, select the interface to use when all the SD-WAN Group interfaces fail to meet the performance criteria specified **Performance Class**; in or all the interfaces are down:

- **None** (default)
- Individual interface
- VPN Tunnel Interface (if any)

8 To specify whether the default state of the performance probe should be treated as UP, select **Performance Probe default state is UP**. If this option is not selected, the performance probe is treated as DOWN. This option is selected by default.

9 For path selection profiles with Non-VPN SD-WAN groups, if existing connections on the path should be reset when the path does not meet the performance criteria any more, select **Reset conditions if path does not meet the performance criteria**. This option is not selected by default.

10 Click **ADD**.

11 To add more Path Selection Profiles, repeat Step 3 through Step 10 for each additional profile.

12 Click **CLOSE**.

Editing a Path Selection Profile

To edit a SD-WAN Path Selection Profile:

1 Navigate to **MANAGE | System Setup | SD-WAN > Path Selection Profiles**.

2 Click the **Edit** icon of the Path Selection Profile to edit. The **Edit SD-WAN Path Selection Profile** dialog displays, which is the same as the **Add SD-WAN Path Selection Profile** dialog.

3 Make changes as described in Configuring Path Selection Profiles on page 588.

4 Click **OK**.

Deleting SD-WAN Path Selection Profiles

You can delete one, multiple, or all SD-WAN Path Selection Profiles

Topics:

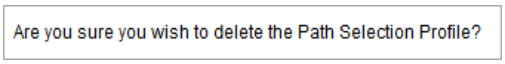
- Deleting a Path Selection Profile on page 589
- Deleting Multiple Path Selection Profiles on page 589
- Deleting All Path Selection Profiles on page 590

Deleting a Path Selection Profile

To delete a SD-WAN Path Selection Profile:

1 Navigate to **MANAGE | System Setup | SD-WAN > Path Selection Profiles**.

2 Click the **Delete** icon of the Performance Class Object to delete. A confirmation message displays.



3 Click **OK**.

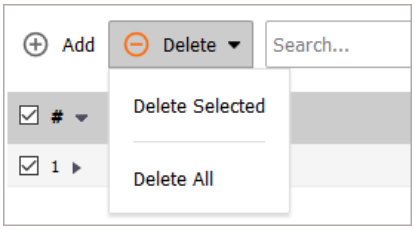
Deleting Multiple Path Selection Profiles

To delete multiple SD-WAN Path Selection Profiles:

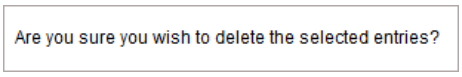
1 Navigate to **MANAGE | System Setup | SD-WAN > Path Selection Profiles**.

2 Select the Path Selection Profiles to delete.

3 From **Delete** above the **SD-WAN Path Selection Profile** table, select **Delete Selected**.



A confirmation message displays.



4 Click **OK**.

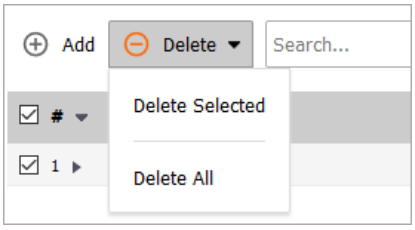
Deleting All Path Selection Profiles

To delete all SD-WAN Path Selection Profiles:

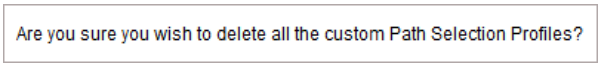
1 Navigate to **MANAGE | System Setup | SD-WAN > Path Selection Profiles**.

2 Select the checkbox in the header of the **SD-WAN Path Selection Profile** table. All the Path Selection Profiles are selected.

3 From **Delete** above the **SD-WAN Path Selection Profile** table, select **Delete All**.



A confirmation message displays.



4 Click **OK**.

36

Configuring SD-WAN Route Policies

TIP: For non-SD-WAN network routing and policies, see Configuring Route Advertisements and Route Policies on page 450.

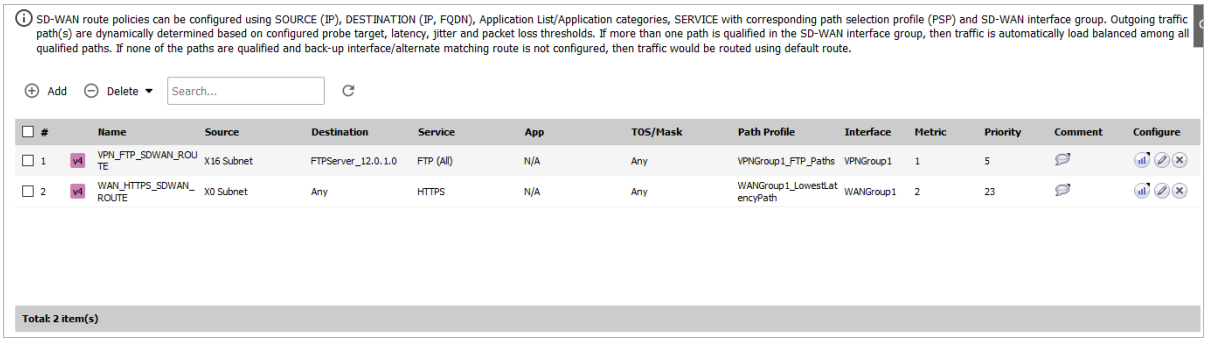
Topics:

- About SD-WAN Route Policies on page 591
- Configuring SD-WAN Route Policies on page 592
- Editing SD-WAN Route Policies on page 594
- Deleting SD-WAN Route Policies on page 594

About SD-WAN Route Policies

Dynamic Path selection for specific traffic flows uses Policy-based Routes. A SD-WAN Policy-based Route is used to configure the route policy for the specific source/destination service/app combination, with a corresponding Path Selection Profile that determines the outgoing path dynamically based on the Path Selection Profile. If there is more than one path qualified by the Path Selection Profile, the traffic is automatically load balanced among the qualified paths. If none of the paths are qualified by the path selection profile and the backup interface in the profile is not configured or is down, the route is disabled. For information about SonicOS SD-WAN, see About SD-WAN on page 570.

TIP: SD-WAN routing can be configured from the **MANAGE | System Setup > Network > Routing** page or **MANAGE | System Setup > SD-WAN > SD-WAN Route Policies** page. The **SD-WAN > SD-WAN Route Policies** page, however, only shows the SD-WAN Routes and only allows configuration of SD-WAN-type routes.



Name	Name of the SD-WAN Route Policy. The IP version is shown by an icon. Included in this column is an icon showing whether the Route Policy is for IPv4 and/or IPv6.
Source	Source address object for the SD-WAN route.

Destination	Destination address object for the SD-WAN route.
Service	Service object for the for the SD-WAN route. If **App** was selected instead of **Service** for the type of route policy, **N/A** appears.
App	App object for the for the SD-WAN route. If **Service** was selected instead of **App** for the type of route policy, **N/A** appears.
TOS/Mask	Hexadecimal TOS and TOS Mask. If these options were not configured, **Any** appears.
Path Profile	Path Selection Profile for the SD-WAN route.
Interface	SD-WAN interface group associated with the SD-WAN route.
Metric	Metric used for the SD-WAN route.
Priority	Priority of the Route Policy.
Comment	Information icon that, when moused over, displays the comment entered when the SD-WAN route policy was configured. It the **Comment** field was left blank, the popup displays the name of the SD-WAN route policy.
Configure	Displays the **Statistics**, **Edit**, and **Delete** icons for the route policy. Mousing over the **Statistics** icon displays the number of active connections for the route.

Configuring SD-WAN Route Policies

To add a SD-WAN route policy:

1 Navigate to **MANAGE | System Setup > SD-WAN > SD-WAN Route Policies**.

2 Click the **Add** icon. The **Add SD-WAN Route Policy** dialog displays.

IMPORTANT: When configuring an SD-WAN route from the **MANAGE | System Setup > Network > Routing > Route Policies** page, choose **SD-WAN Route** on the **Add Route Policy** dialog; the options change to match those of the **Add SD-WAN Route Policy** dialog.



NOTE: The **Interface** and **Disable route when the interface is disconnected** options are dimmed because these options cannot be edited in SD-WAN policies. The **Interface** option is populated with the SD-WAN group name in the associated Path Selection Profile (PSP) and cannot be changed. The interface for the SD-WAN route is selected from the SD-WAN group that is part of the PSP associated with the SD-WAN route and, therefore, cannot be configured.

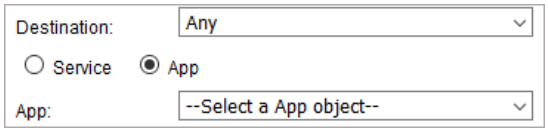
3 Enter a meaningful name in the **Name** field.

4 From **Source**, select the source address object for the static route or select **Create new address object** to dynamically create a new address object. The default is **Any**.

5 From **Destination**, select the destination address object or select **Create new address object** to dynamically create a new address object. The default is **Any**.

6 Choose the type of route policy:

- **Service** (default)
- **App** – Service changes to **App**:



IMPORTANT: Application Control Licensing is required for application-based routing.

Go to Step 9.

7 From **Service**, select a service object. For a generic static route that allows all traffic types, simply select **Any** (the default).

8 Go to Step 10.

9 From **App**, select an App Object.

10 From **Path Profile**, select a Path Selection Profile.

11 Enter the **Metric** (weighted cost) for the route. The minimum is 1, and the maximum is 254.

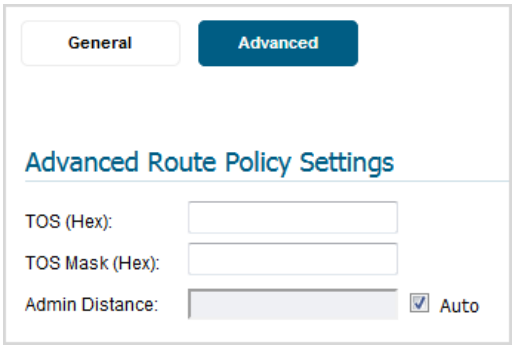
For more information on metrics, see About Metrics and Administrative Distance on page 451 and Policy-based Routing on page 453.

TIP: Lower metrics are considered better and take precedence over higher metrics (costs).

12 Optionally, enter a **Comment** for the route. This field allows you to enter a descriptive comment for the new static route policy.

13 If WXA is licensed, select the WXA group from **WXA Group**. The default is **None**.

14 Click **Advanced**.



15 Enter a TOS value in the **TOS (Hex)** field. The maximum value is FF. If the **TOS** and **TOS Mask** fields are not configured, a value of 0 is used. For further information about TOS and TOS Mask values, see Policy-based TOS Routing on page 453.

16 Enter the same value in the **TOS Mask (Hex)** field.

17 To manually specify an administration distance:

a Deselect **Auto**. The **Admin Distance** field becomes available. This option is selected by default. For information about administration distance, see About Metrics and Administrative Distance on page 451.

b Enter the administration distance in the **Admin Distance** field.

18 Click **OK**.

Editing SD-WAN Route Policies

To edit SD-WAN Route Policies:

1 Navigate to **MANAGE | System Setup > SD-WAN > SD-WAN Route Policies**.

2 Click the **Edit** icon of the Path Selection Profile to edit. The **Edit SD-WAN SD-WAN Route Policy** dialog displays, which is the same as the **Add SD-WAN Route Policy** dialog.

3 Make changes as described in Configuring SD-WAN Route Policies on page 592.

4 Click **OK**.

Deleting SD-WAN Route Policies

You can delete one, multiple, or all SD-WAN Route Policies.

Topics:

- Deleting SD-WAN Route Policies on page 595
- Deleting Multiple SD-WAN Route Policies on page 595
- Deleting All SD-WAN Route Policies on page 595

Deleting SD-WAN Route Policies

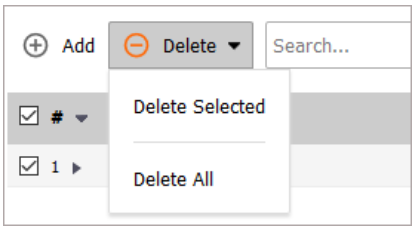
To delete a SD-WAN route policy:

1. Navigate to **MANAGE | System Setup > SD-WAN > SD-WAN Route Policies**.
2. Click the **Delete** icon of the SD-WAN route policy to delete.

Deleting Multiple SD-WAN Route Policies

To delete multiple SD-WAN SD-WAN Route Policies

1. Navigate to **MANAGE | System Setup > SD-WAN > SD-WAN Route Policies**.
2. Select the Path Selection Profiles to delete.
3. From **Delete** above the **SD-WAN Path Selection Profile** table, select **Delete Selected**.



A confirmation message displays.

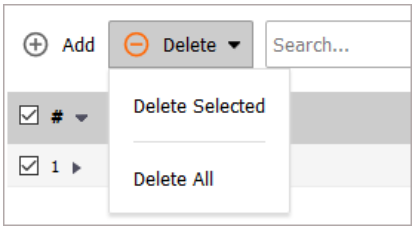


4. Click **OK**.

Deleting All SD-WAN Route Policies

To delete all SD-WAN Route Policies:

1. Navigate to **MANAGE | System Setup > SD-WAN > SD-WAN Route Policies**.
2. Select the checkbox in the header of the **SD-WAN SD-WAN Route Policies** table. All the Path Selection Profile are selected.
3. From **Delete** above the **SD-WAN Path Selection Profile** table, select **Delete All**.



A confirmation message displays.



4 Click **OK**.

37

Monitoring SD-WAN

Topics:

- **SD-WAN > SD-WAN Monitor** on page 597

SD-WAN > SD-WAN Monitor

TIP: You can also view these graphs on the **MONITOR | Appliance Health > SD-WAN Monitor** page; for further information, see *SonicOS 6.5 Monitoring*.



NOTE: A chart may be empty or blank if there are no recent data entries received within the viewing range.

To monitor SD-WAN performance:

1. Navigate to **MONITOR | Appliance Health > SD-WAN Monitor.**
2. From **SD-WAN Probes**, select the performance probe you would like to use to monitor.

3 Indicate the Refresh rate, in seconds, in the **Refresh every:** field.

4 Select a **View Range**:

- **60 seconds**
- **2 Minutes**
- **5 minutes**
- **10 minutes** (default).

5 Choose an interface to track or select **All Interfaces** from the drop-down menu on the right side.

6 For scaling ratios, you can enter values such as:

- **Auto** - Auto Y-Scaling
- **<num>[<unit>]** - *num* is a numeric integer. The unit is optional, but can also be empty, K for Kilo, M for Mega, G for Giga, or % for percentages.

  NOTE: An invalid value defaults to Auto Y-scaling.

7 The two small icons on the right allow you to toggle between line and block displays.

38

Viewing SD-WAN Route Policy Connections

Topics:

- SD-WAN > SD-WAN Connection Logs on page 599

SD-WAN > SD-WAN Connection Logs

You can view the connections that have been associated with SD-WAN Route Policies on the **MANAGE | System Setup > SD-WAN > SD-WAN Connection Logs** page. For information about SonicOS SD-WAN and its functions, see About SD-WAN on page 570.

TIP: You can also view this information on the **INVESTIGATE | Logs > SD-WAN Connection Logs** page.



Src MAC	MAC address of the appliance that is the source of the connection.
Src Vendor	Name of the vendor of the appliance that is the source of the connection.
Src IP	IP address of the appliance that is the source of the connection.
Src Port	Port on the appliance that is the source of the connection.
Dst MAC	MAC address of the appliance that is the destination of the connection.
Dst Vendor	Name of the vendor of the appliance that is the destination of the connection.
Dst IP	IP address of the appliance that is the destination of the connection.
Dst Port	Port on the appliance that is the destination of the connection.
Protocol	Protocol used for the connection.
Src Iface	Interface on the appliance that is the source of the connection.
Dst Iface	Interface on the appliance that is the destination of the connection.
Flow Type	Type of data flow control, such as **FTP Control**.
IPS Category	Internet Provider Security (IPS) category. If this information is not available or relevant, the column displays **N/A**.
ABR App ID	App-Based Routing Application ID.
ABR Category ID	App-Based Routing Category ID.
Expiry (sec)	Number of seconds until the connection expires.

TX Bytes	Number of bytes transmitted on the connection.
Rx Bytes	Number of bytes received on the connection.
TX Pkts	Number of packets transmitted on the connection.
Rx Pkts	Number of packets received on the connection.
Flush	Displays the **Flush** icon. Clicking the icon flushes the connection.
Total	Total number of entries in the **SD-WAN > SD-WAN Connection Logs** page.

Part 6

System Setup | Switching

- About Switching
- Configuring VLAN Trunking
- Managing Layer 2 Discovery and LLDP/LLTD
- Configuring Link Aggregation
- Configuring Port Mirroring
- Configuring Switch Shield
- Configuring the SonicWall Switch

39

About Switching

About Switching

NOTE: Switching is available on all NSa 2650 and above and SuperMassive appliances.

NOTE: This section describes advanced switching in SonicOS, which is different from managing a Dell X-Series switch or N-Series switch from a TZ appliance. For more information about managing X-Series switches, see SonicOS Support of X-Series/N-Series Switches on page 348.

SonicOS supports L2 switch functionality; this support is summarized in SonicOS Switching.

Sonicwall appliances can manage two different switch product lines: the X-Series/N-Series extended switches and the SonicWall SWS-series switches.

Configure the X-Series/N-Series extended switches through the **MANAGE > Networks > Portshield Groups > External Switch Configuration** menu. For complete details refer to the *SonicWall SoniccOS 6.5 X-Series / N-Series Solution Deployment Guide* at https://www.sonicwall.com/support/technical-documentation/?language=English&category=Firewalls&resources=Deployment%20Guide.

Configure and manage the SonicWall SWS-series switches through the **MANAGE | System Setup > Switch Controller** pages. For complete details refer to the *SonicWall Switch Getting Started Guide* at https://www.sonicwall.com/support/technical-documentation/?language=English. These features are described in detail in the chapter entitled Configuring from the Firewall on page 640.

Topics:

- SonicOS Switching on page 602
- Benefits of Switching on page 603
- How Switching Works on page 604
- Glossary on page 604

SonicOS Switching

SonicOS provides Layer 2 (data link layer) switching functionality that supports these switching features:

NOTE: Switching is available on all NSA 2650 and above and SuperMassive appliances.

- **VLAN Trunking** – Provides the ability to trunk different VLANs between multiple switches.
- **Layer 2 Network Discovery** – Uses IEEE 802.1AB (LLDP) and Microsoft LLTD protocols and switch forwarding table to discover devices visible from a port.

- **Link Aggregation** – Provides the ability to aggregate ports for increased performance and redundancy.

  NOTE: Link Aggregation is supported on NSA 3600 and higher Security Appliances. On the NSA 2600, Link Aggregation for Network Interfaces is a separate feature from Link Aggregation for Switching. The NSA 2600 does support Link Aggregation for Network Interfaces (see Configuring Link Aggregation on page 627), but the NSA 2600 does not support Switching and, therefore, does not support Link Aggregation for Switching.

- **Port Mirroring** – Allows you to assign a mirror port to mirror ingress, egress or bidirectional packets coming from a group of ports.
- **Jumbo Frames** – Supporting jumbo frames allows the SonicOS to process Ethernet frames with payloads ranging from 1500-9000 bytes.

  NOTE: Jumbo frames are supported on NSA 3600 and higher appliances.

Benefits of Switching

SonicOS provides a combined security and switching solution. Layer 2 switching features enhance the deployment and interoperability of SonicWall devices within existing Layer-2 networks.

NOTE: Advanced switching is supported on NSA 3600 and higher appliances.

The advanced switching features on a network Security Appliance provide these benefits:

- **Increased port density** – With one appliance providing up to 26 interfaces, including up to 24 switch ports, you can decrease the number of devices on your internal network.
- **Increased security across multiple switch ports** – The PortShield architecture provides the flexibility to configure all LAN switch ports into separate security zones such as LANs, WLANs and DMZs, providing protection not only from the WAN and DMZ, but also between devices inside the LAN. Effectively, each security zone has its own wire-speed “mini-switch” that benefits from the protection of a dedicated deep packet inspection firewall.
- **VLAN Trunking** – Simplifies VLAN management and configuration by reducing the need to configure VLAN information on every switch; provides the ability to trunk different VLANs between multiple switches.
- **Layer 2 Network Discovery** – Provides Layer 2 network information for all devices attached to the appliance; uses IEEE 802.1AB (LLDP) and Microsoft LLTD protocols and switch forwarding table to discover devices visible from a port.
- **Link Aggregation** – Aggregated ports provide increased performance through load balancing when connected to a switch that supports aggregation, and provide redundancy when connected to a switch or server that supports aggregation.
- **Port Mirroring** – Allows you to easily monitor and inspect network traffic on one or more ports and to assign a mirror port to mirror ingress, egress or bidirectional packets coming from a group of ports.
- **Jumbo Frames** – Allows increased throughput and reduces the number of Ethernet frames to be processed by allowing SonicOS to process Ethernet frames with payloads ranging from 1500-9000 bytes. Throughput increase may not be seen in some cases. However, there are some improvement in throughput if the packets traversing are really jumbo size.

  NOTE: Jumbo frames are supported on NSA 3600 and higher appliances.

How Switching Works

Some switching features operate on PortShield Groups and require preliminary configuration on the **Network > PortShield Groups** page. Some operate on existing **Network > Interfaces** configurations. For more information about configuring these related features in SonicOS, see:

- Configuring Interfaces on page 259
- Configuring PortShield Interfaces on page 347

For details about the operation of each switching feature, see:

- Configuring VLAN Trunking on page 605
- Managing Layer 2 Discovery and LLDP/LLTD on page 612
- Configuring Link Aggregation on page 627
- Configuring Port Mirroring on page 633

Glossary

BPDU	Bridge Protocol Data Unit – Used in RSTP, BPDUs are special data frames used to exchange information about bridge IDs and root path costs. BPDUs are exchanged every few seconds to allow switches to keep track of network topology and start or stop port forwarding.
CoS	Class Of Service – Cos (IEEE 802.1p) defines eight different classes of service that are indicated in a 3-bit user_priority field in an IEEE 802.1Q header added to an Ethernet frame when using tagged frames on an 802.1 network.
DSCP	Differentiated Services Code Point – Also known as DiffServ, DSCP is a networking architecture that defines a simple, coarse-grained, class-based mechanism for classifying and managing network traffic and providing Quality of Service (QoS) guarantees on IP networks. RFC 2475, published in 1998 by the IETF, defines DSCP. DSCP operates by marking an 8-bit field in the IP packet header.
IETF	Internet Engineering Task Force – The IETF is an open standards organization that develops and promotes Internet standards.
L2	OSI Layer 2 (Ethernet) – Layer 2 of the seven layer OSI model is the Data Link Layer, on which the Ethernet protocol runs. Layer 2 is used to transfer data among network entities.
LACP	Link Aggregation Control Protocol – LACP is an IEEE specification that provides a way to combine multiple physical ports together to form a single logical channel. LACP allows load balancing by the connected devices.
LLDP	Link Layer Discovery Protocol (IEEE 802.1AB) – LLDP is a Layer 2 protocol used by network devices to communicate their identity, capabilities, and interconnections. This information is stored in a MIB database on each host, which can be queried with SNMP to determine the network topology. The information includes system name, port name, VLAN name, IP address, system capabilities (switching, routing), MAC address, link aggregation, and more.
LLTD	Link Layer Topology Discovery (Microsoft Standard) – LLTD is a Microsoft proprietary protocol with functionality similar to LLDP. It operates on wired or wireless networks (Ethernet 802.3 or wireless 802.11). LLTD is included on Windows Vista and Windows 7, and can be installed on Windows XP.
PDU	Protocol Data Unit – In the context of the Switching feature, the Layer 2 PDU is the frame. It contains the link layer header followed by the packet.
RSTP	Rapid Spanning Tree Protocol (IEEE 802.1D-2004) – RSTP was defined in 1998 as an improvement to Spanning Tree Protocol. It provides faster spanning tree convergence after a topology change.

40

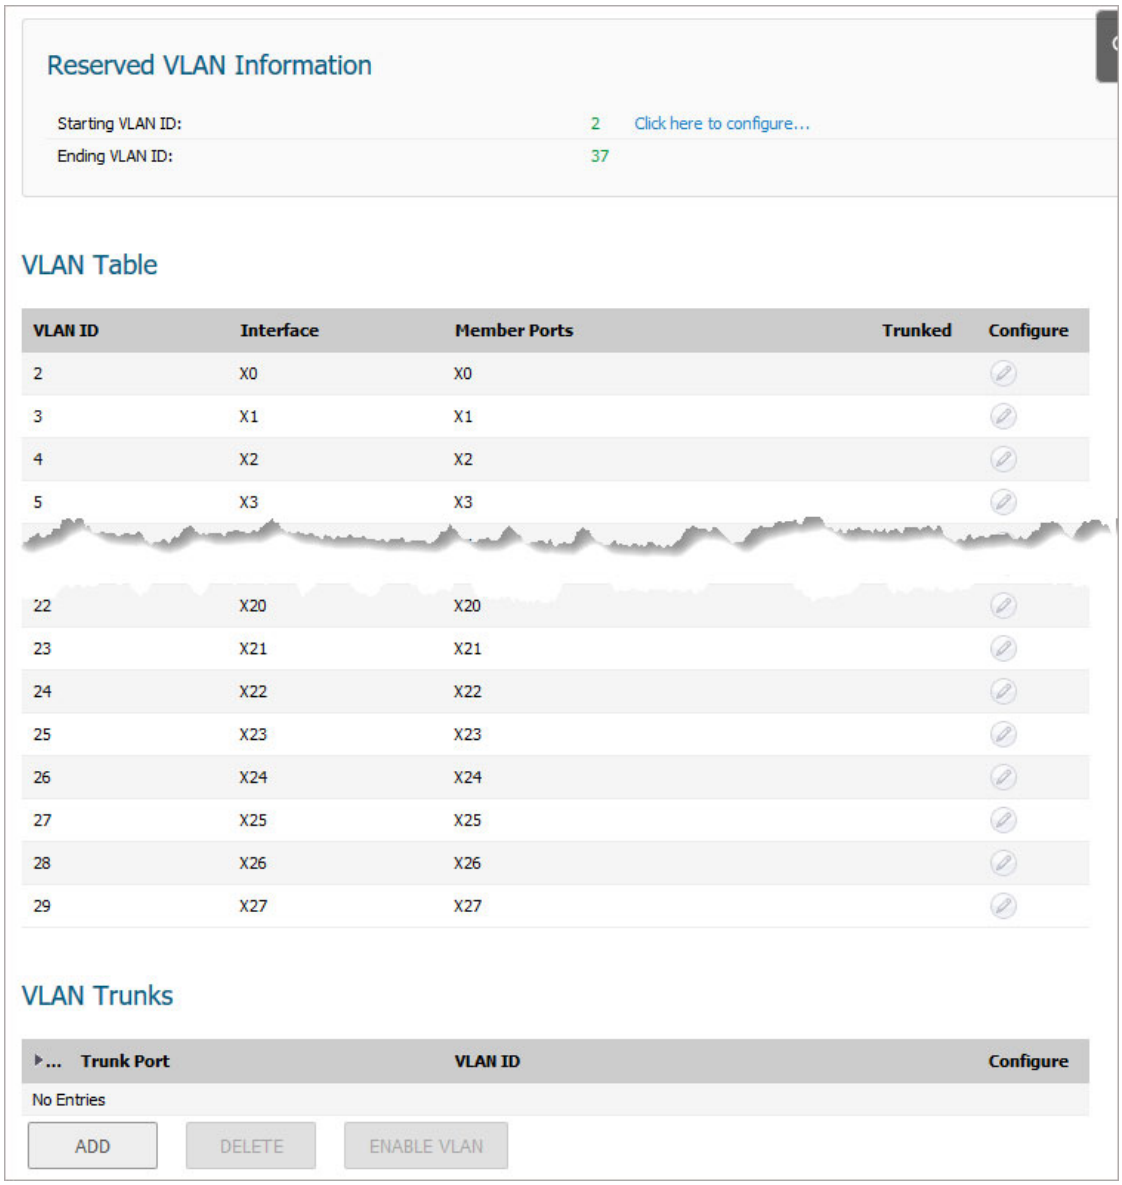
Configuring VLAN Trunking

NOTE: Switching is available on all NSA 2650 and above and SuperMassive appliances.

Topics:

- **Switching > VLAN Trunking** on page 606
  - About Trunking on page 607
  - Viewing VLANs on page 607
  - Editing VLANs on page 609
  - Adding a VLAN Trunk Port on page 609
  - Enabling a VLAN on a Trunk Port on page 610
  - Deleting VLAN Trunk Ports on page 610

Switching > VLAN Trunking



Topics:

- About Trunking on page 607
- Viewing VLANs on page 607
- Editing VLANs on page 609
- Adding a VLAN Trunk Port on page 609
- Deleting VLAN Trunk Ports on page 610
- Enabling a VLAN on a Trunk Port on page 610

About Trunking

Unassigned switch ports on SonicOS can function as VLAN trunk ports. You can enable or disable VLANs on the trunk ports, allowing the existing VLANs on SonicOS to be bridged to respective VLANs on another switch connected through the trunk port. SonicOS support 802.1Q encapsulation on the trunk ports. A maximum of 32 VLANs can be enabled on each trunk port.

The VLAN trunking feature provides these functions:

- Change VLAN ID's of existing PortShield groups
- Add/delete VLAN trunk ports
- Enable/disable customer VLAN IDs on the trunk ports

The allowed VLAN ID range is 1-4094. Some VLAN IDs are reserved for PortShield use, and the reserved range is displayed on **MANAGE | System Setup** | **Switching > VLAN Trunking**.

You can mark certain PortShield groups as "Trunked." If the PortShield group is dismantled, the associated VLAN is automatically disabled on the trunk ports.

VLANs can exist locally in the form of PortShield groups or can be totally remote VLANs. You can change the VLAN ID of PortShield groups on SonicOS. This allows easy integration with existing VLAN numbering.

SonicOS does not allow changing port VLAN membership in an ad-hoc manner. VLAN membership of a port must be configured through PortShield configuration in the SonicOS management interface. For more information about configuring PortShield groups, see Configuring PortShield Interfaces on page 347.

A virtual interface (called the VLAN Trunk Interface) is automatically created for remote VLANs. When the same remote VLAN is enabled on another trunk port, no new interface is created. All packets with the same VLAN tag ingressing on different trunk ports are handled by the same virtual interface. This is a key difference between VLAN sub-interfaces and VLAN trunk interfaces.

The **Name** column on **MANAGE | System Setup** | **Network > Interfaces** displays the VLAN IDs of the VLAN Trunk Interfaces for the VLAN trunks.

You can enable any VLAN, local or remote, on a VLAN trunk to allow bridging to two respective VLANs on another switch. For example, local VLAN 345 can be enabled on the VLAN trunk for port X2, which also has two remote VLANs enabled on it.

VLAN trunking interoperates with Link Aggregation and Port Mirroring features. A VLAN trunk port can be mirrored, but cannot act as a mirror port itself.

Ports configured as VLAN trunks cannot be used for any other function and are reserved for use in Layer 2 only. For example, you cannot configure an IP Address for the trunk ports.

When a Trunk VLAN interface has been configured on a particular trunk port, that trunk port cannot be deleted until the VLAN interface is removed, even though the VLAN is enabled on multiple trunk ports. This is an implementation limitation.

Viewing VLANs

Topics:

- Reserved VLAN Information on page 608
- VLAN Table on page 608
- VLAN Trunks Table on page 609

Reserved VLAN Information

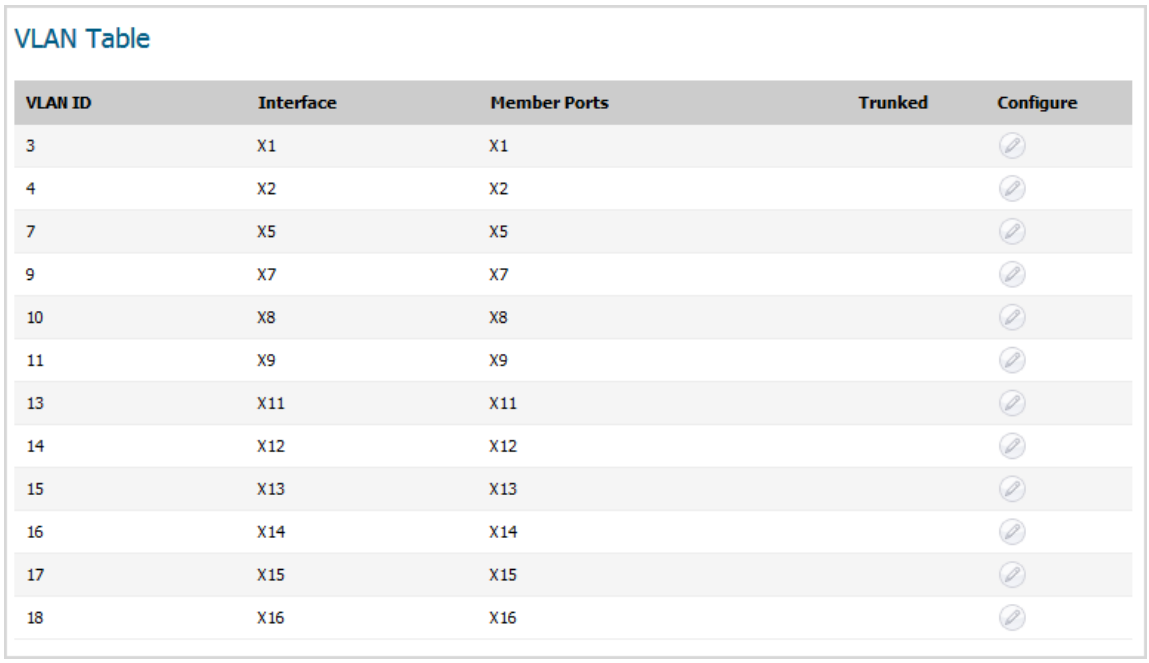


The **Reserved VLAN Information** table lists the range of reserved VLAN IDs:

- **Starting VLAN ID**
- **Ending VLAN ID**

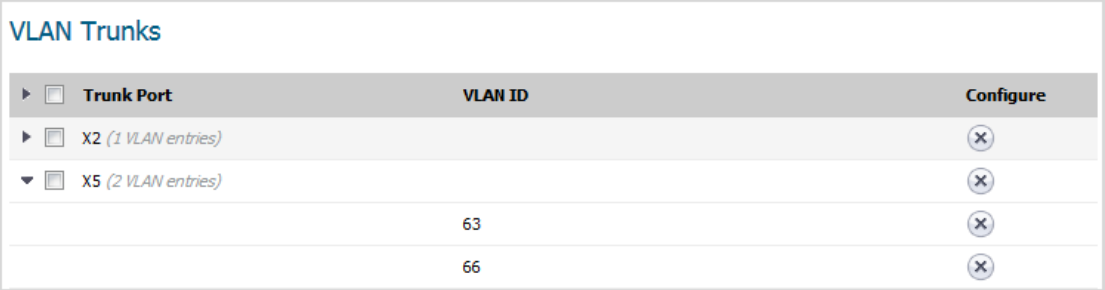
To change the starting VLAN ID, click the `Click here to configure...` link. The **MANAGE | Security Configuration** > **Firewall Settings > Advanced Settings** page displays. For more information, see *SonicOS 6.5 System Setup*.

VLAN Table



VLAN ID	ID of the VLAN.
Interface	Interface assigned to the VLAN.
Member Ports	Ports associated with the interface.
Trunked	Indicates whether this VLAN is trunked.
Configure	Contains **Edit** icons for the VLANs.

VLAN Trunks Table



Trunk Port	Interface for the Trunk port and the number of VLAN entries associated with it
VLAN ID	ID(s) of the VLAN(s)
Configure	Contains **Delete** icons for the VLANs

To display the VLAN ID(s) of the Trunk Port, click the **Expand** icon for the Trunk port. To display the VLAN ID(s) of all the Trunk Ports, click the **Expand** icon in the **VLAN Trunks** table header. To hide the VLAN ID(s), click the appropriate **Collapse** icon.

Editing VLANs

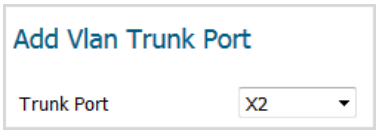
To edit a VLAN:

1. Navigate to **Switching > VLAN Trunking**.
2. Click the **Configure** icon in the **VLAN Table** row for the VLAN ID you want to edit. The **Edit VLAN for PortShield Host** dialog displays.
3. Do one of the following:
   - Type a different VLAN ID into the **VLAN ID** field. You can enter any VLAN ID except the original system-specified VLAN ID or any others in the **Reserved VLAN Information** table.
   - Use the VLAN ID number in the **VLAN ID** field, which matches the one for which you clicked the **Configure** icon.
4. To enable trunking for this VLAN, select the **Trunked** checkbox. To disable trunking for this VLAN, clear the checkbox.
5. Click **OK**.

Adding a VLAN Trunk Port

To add a VLAN trunk port:

1. Navigate to **Switching > VLAN Trunking**.
2. Under **VLAN Trunks**, click **Add**. The **Add VLAN Trunk Port** dialog displays.



3 Select the port to add from the **Trunk Port** drop-down menu.

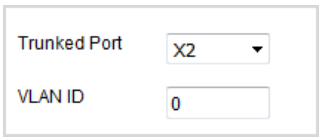
4 Click **OK**.

Enabling a VLAN on a Trunk Port

To enable a custom VLAN ID on a specific trunk port:

1 Navigate to **Switching > VLAN Trunking**.

2 Under the **VLAN Trunks** table, click **Enable VLAN**. The **Enable VLAN** dialog displays.



3 Select a trunked port from the **Trunked Port** drop-down menu. This is the port that you want to use to trunk the VLAN ID indicated in the **VLAN ID** field.

4 In the **VLAN ID** field, type in the VLAN ID to be trunked. This can be a VLAN ID on another switch.

5 Click **OK**.

Deleting VLAN Trunk Ports

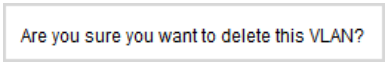
You can delete one VLAN trunk port, multiple ports at a time, or all ports.

To delete a VLAN trunk port:

1 Navigate to **Switching > VLAN Trunking**.

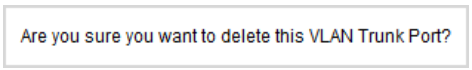
2 Expand the VLAN trunk port to be deleted.

3 Click the **Delete** icon in the **Configure** column for the VLANs to be deleted. A confirmation message displays:



4 Click **OK**.

5 Click the **Delete** icon in the **Configure** column for the port to be deleted. A confirmation message displays:



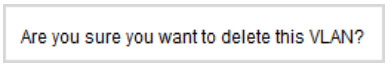
6 Click **OK**.

To delete multiple VLAN trunk ports:

1 Navigate to **Switching > VLAN Trunking**.

2 In the **VLAN Trunks** table, expand the VLAN trunk ports to be deleted.

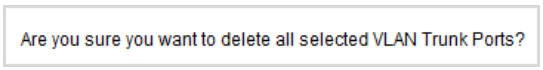
3 Click the **Delete** icon in the **Configure** column for each VLAN to be deleted. A confirmation message displays:



4 Click **OK** for each one.

5 Select the checkboxes for the VLAN trunk ports you want to delete. **Delete** becomes available.

6 Click **Delete**. A confirmation message displays.



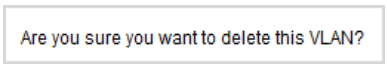
7 Click **OK**.

To delete all VLAN trunk ports:

1 Navigate to **Switching > VLAN Trunking**.

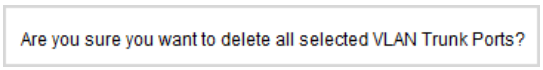
2 In the **VLAN Trunks** table, expand the VLAN trunk ports by clicking the **Expand** icon in the **VLAN Trunks** table heading.

3 Click the **Delete** icon in the **Configure** column for each VLAN to be deleted. A confirmation message displays:



4 Select the checkbox in the **VLAN Trunks** table heading. **Delete** becomes available.

5 Click **Delete**. A confirmation message displays.



6 Click **OK**.

41

Managing Layer 2 Discovery and LLDP/LLTD

NOTE: Switching is available on all NSA 2650 and above and SuperMassive appliances.

LLDP is supported on the NSA 3600 and higher Security Appliances and is also supported when High Availability is enabled.

Topics:

- **Switching > Port Mirroring** on page 612
  - About L2 Discovery and LLDP on page 613
  - Viewing L2 Discovery and LLDP/LLTD Interfaces on page 616
  - Associating an LLDP Profile with a L2 Discovery Interface on page 619
  - Refreshing the Page on page 619
  - Globally Enabling/Disabling LLDP on page 619
  - Discovering Neighbors on page 620
- **Switching > Port Mirroring > LLDP Profile** on page 621
  - Viewing LLDP Profiles on page 621
  - Adding a LLDP Custom Profile on page 623
  - Editing a Custom LLDP Profile on page 625
  - Deleting Custom Profiles on page 625

Switching > Port Mirroring

Topics:

- About L2 Discovery and LLDP on page 613
- Viewing L2 Discovery and LLDP/LLTD Interfaces on page 616
- Associating an LLDP Profile with a L2 Discovery Interface on page 619
- Refreshing the Page on page 619
- Globally Enabling/Disabling LLDP on page 619
- Discovering Neighbors on page 620

About L2 Discovery and LLDP

To discover neighboring devices and their capabilities, the SonicWall Security Appliance uses:

- IEEE 802.1AB (LLDP: Link Layer Discovery Protocol)/Microsoft LLTD (Link Layer Topology Discovery)
- IEEE 802.3-2012 protocols
- A switch-forwarding table

LLDP operates at Layer 2 and exchanges LLDP Protocol Data Units (LLDPDUs) between the neighbors containing a sequence of variable length information elements that include type-length-values (TLV). The information is stored in the SNMP MIBs. These Layer 2 protocols are used by networking devices to advertise their identities and capabilities and to identify their directly connected Layer 2 neighbors/peers on wired Ethernet networks; they do not cross a broadcast domain.

More information about these protocols is available at:

- https://en.wikipedia.org/wiki/Link_Layer_Topology_Discovery
- https://en.wikipedia.org/wiki/Link_Layer_Discovery_Protocol

TIP: SonicOS supports LLDP Transmit and Transmit-Receive Modes.

- https://msdn.microsoft.com/en-us/library/windows/desktop/dn594471(v=vs.85).aspx

LLDP makes troubleshooting easier, especially in cases where peers are not detected by ping or traceroute.

Topics:

- Supported LLDP Modes on page 613
- Type-Length-Values on page 614
- Effect of Interface Link on LLDP Functions on page 615

Supported LLDP Modes

These LLDP modes are supported in SonicOS:

- LLDP-receive (already supported in previous versions of SonicOS 6.5)
- LLDP-transmit
- LLDP-transmit-receive
- LLDP-disabled

You can create custom LLDP profiles for individual interfaces.

These interface types and modes support LLDP:

L2 Interface	If the physical port is configured in L2 Mode
L3 Interface	If the physical port is configured in L3 Mode
Wire-Mode Interface	Supported for secure and inspect mode for wire-mode interfaces, but not for VLAN interfaces
L2 Bridge Interface	Supported for the physical interface, but not for VLAN interfaces
VLAN Sub-Interface	Not supported
LAG/LACP	Supported for learn only on the aggregate port and not a member, but is supported for send on individual interfaces. An aggregate port shows neighbor information for both itself and its members.

Type-Length-Values

Each LLDP frame starts with three mandatory type-length-values TLVs: Chassis ID, Port ID and TTL. The mandatory TLVs are followed by any number of optional TLVs. The LLDP frame ends with a mandatory End-of-frame TLV.

Topics:

- Mandatory TLVs on page 614
- Optional TLVs on page 614

Mandatory TLVs

Mandatory TLVs describes the mandatory LLDP TLVs supported for both transmit and receive.

Mandatory TLVs

TLV Name	TLV Type	Description	SonicOS Usage
Chassis ID TLV	1	Identifies the firewall chassis. Each firewall must have exactly one unique Chassis ID.	SonicOS sends the MAC address of the Security Appliance in the Chassis ID field. The MAC address is same as the Security Appliance serial number.
Port ID TLV	2	Identifies the port from which the LLDPDU is sent. The Security Appliance uses the interface's `ifname` as the Port ID. For example, Port ID can be X1, X2, X3.	The Port ID subtype 5 (interface name) is used to identify the transmitting port.
Time-to-live (TTL) TLV	3	Specifies how long (in seconds) LLDPDU information received from the peer is retained as valid in the local Security Appliance (range is 0-65535). The value is a multiple of the LLDP Hold Time Multiplier. When the TTL value is 0, the information associated with the device is no longer valid and SonicOS removes that entry from the database.	Calculated internally
End of LLDPDU frame TLV	0	Indicates the end of the TLVs in the LLDP Ethernet frame.	

Optional TLVs

Optional TLVs describes the optional LLDP TLVs supported for both transmit and receive.

Optional TLVs

TLV Name	TLV Type	Description	SonicOS Usage
Port Description	4	The port description in alpha-numeric format.	Advertises the values/string added in the comment section of the network interface field.
System Name	5	The Security Appliance name in alpha-numeric format.	Advertises the **Firewall Name** configured on the **MANAGE | System Setup > Appliance > Base Settings** page.
System Description	6	The full name and version identification of the system's hardware type, software operating system, and networking software in alpha-numeric format.	Advertised as Firewall in this field.
System Capabilities	7	This field contains a bit-map of the capabilities that define primary functions of the system. Describes the deployment mode of the interface: • An L3 interface is advertised with router (bit 6) capability and the “other” bit (bit 1). • An L2 interface is advertised with MAC Bridge (bit 3) capability and the “other” bit (bit 1). A virtual wire interface is advertised with Repeater (bit 2) capability and the “other” bit (bit 1).	Advertises the features supported by the Security Appliance and the enabled features.
Management Address	8	IP address used for the management of the device: • IP address of the management (MGT) interface • IPv4 address of the `interfaceLoopback` address • User-defined address entered in the management address field; If no management IP address is provided, the default is the MAC address of the transmitting interface. The interface number of the specified management address is included. Also included is the OID of the hardware interface with the specified management address (if applicable). If more than one management address is specified, they are sent in the order they are specified, starting at the top of the list. One Management Address is supported. This is an optional parameter and can be left disabled.	Advertises the management IP address of an interface if it is configured. **NOTE:** IPv6 is not supported.

Effect of Interface Link on LLDP Functions

LLDP only functions when the interface link is up. When the mode is changed:

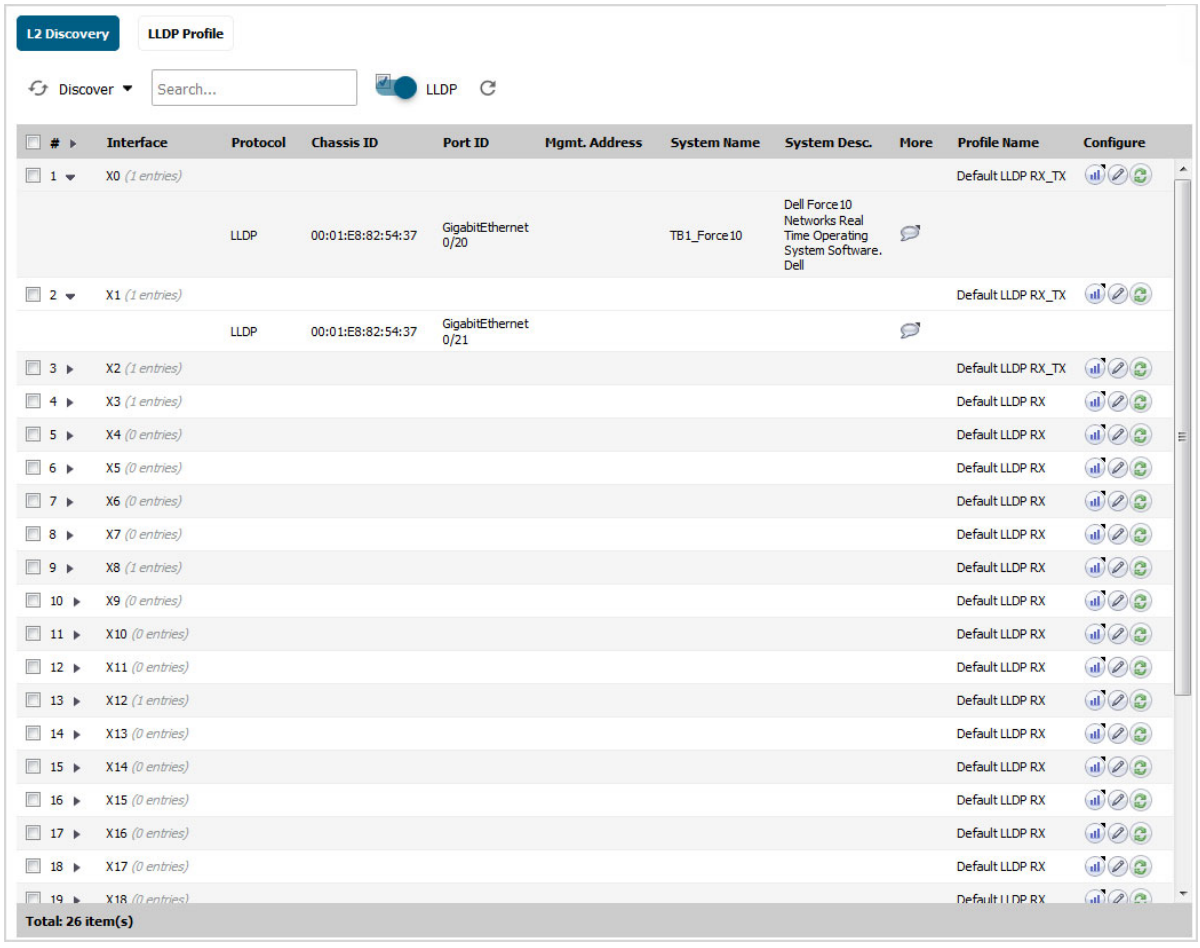
- From Receive to Transmit ,
- From Transmit-Only to Receive-Only,
- To Disabled,

a final LLDP shutdown LLDPDU is sent with these mandatory TLVs:

- Chassis ID TLV
- Port ID TLV
- TTL TLV
- End of LLDPDU TLV

The statistics counters are reset after the link goes down.

Viewing L2 Discovery and LLDP/LLTD Interfaces



Interface	Lists the Security Appliance's interfaces along with either the number of entries.
Profile Name	Name of the default or custom profile name.
Configure	Contains the **Statistics**, **Edit**, and **Refresh** icons for the interfaces. **NOTE:** The **Refresh** icon refreshes only LLTD discovery, not LLDP discovery. To refresh LLDP discovery, click the **Refresh** icon above the **L2 Discovery** table.

NOTE: Only the **Interface** and **Profile Name** columns contain information about interfaces, and the Configure column icons apply only to the interface. The other columns display information about the entries under an interface; for information about these columns, see Displaying Peer Information on page 617.

Topics:

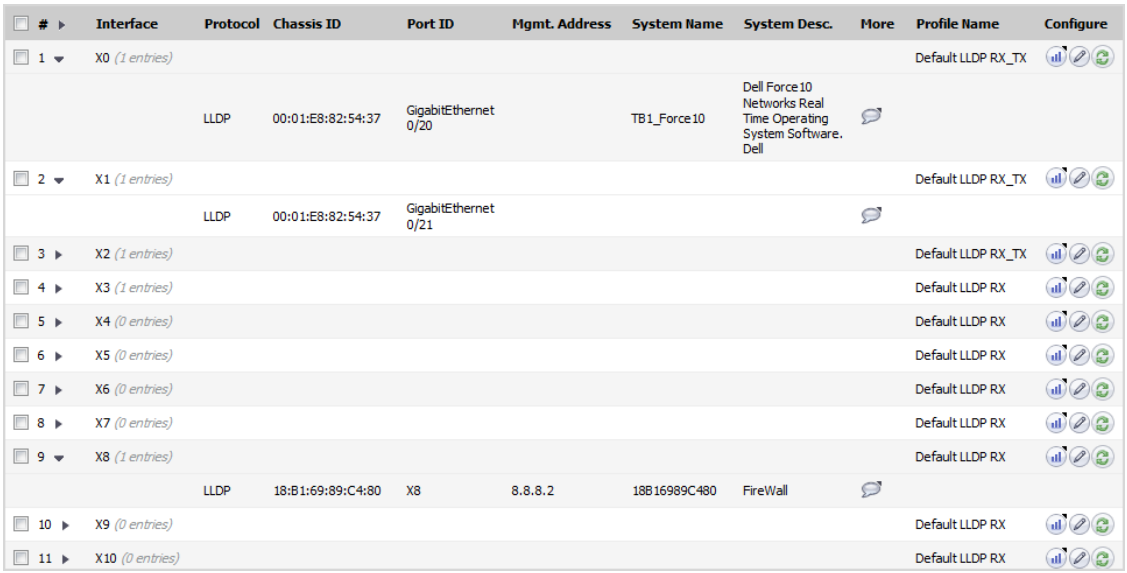
- Displaying Peer Information on page 617
- Displaying Statistics on page 618
- Searching the L2 Discovery Table on page 618

Displaying Peer Information

To display L2 discovery information:

1 Navigate to **MANAGE | System Setup > Switching > Port Mirroring**.

2 In the **L2 Discovery** table, click the **Expand** icon for the desired interface. Information about the nodes (entries) discovered for the interface are displayed.

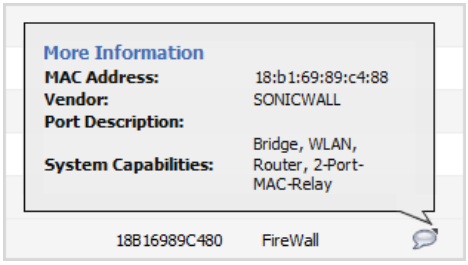


#	Interface	Protocol	Chassis ID	Port ID	Mgmt. Address	System Name	System Desc.	More	Profile Name	Configure
1	X0 *(1 entries)*								Default LLDP RX_TX	
		LLDP	00:01:E8:82:54:37	GigabitEthernet 0/20		TB1_Force10	Dell Force10 Networks Real Time Operating System Software. Dell			
2	X1 *(1 entries)*								Default LLDP RX_TX	
		LLDP	00:01:E8:82:54:37	GigabitEthernet 0/21						
3	X2 *(1 entries)*								Default LLDP RX_TX	
4	X3 *(1 entries)*								Default LLDP RX	
5	X4 *(0 entries)*								Default LLDP RX	
6	X5 *(0 entries)*								Default LLDP RX	
7	X6 *(0 entries)*								Default LLDP RX	
8	X7 *(0 entries)*								Default LLDP RX	
9	X8 *(1 entries)*								Default LLDP RX	
		LLDP	18:B1:69:89:C4:80	X8	8.8.8.2	18B16989C480	FireWall			
10	X9 *(0 entries)*								Default LLDP RX	
11	X10 *(0 entries)*								Default LLDP RX	

Chassis ID	Identifies the Security Appliance's chassis. Each Security Appliance must have exactly one unique Chassis ID that is a string value consisting of mostly the MAC address of the peer.
Port ID	Identifies the port from which the LLDPDU is sent and is a string value of the port name or number. The Security Appliance uses the interface's `ifname` as the Port ID. For example, Port ID can be X1, X2, X3.
Mgmt. Address	Lists the IP or MAC address of the peer used for the management of the device. If multiple management addresses are returned, only the first address is shown.
System Name	Name of the Security Appliance, in alpha-numeric format.
System Desc.	Full name and version identification of the Security Appliance's hardware type, software operating system, and networking software, in alpha-numeric format.
More	Contains an **Information** icon that displays additional peer information.

NOTE: For information about the other columns, see Viewing L2 Discovery and LLDP/LLTD Interfaces on page 616.

3 To display additional peer information for a peer entry, mouse over the **Information** icon in the **More** column for that peer. A popup displays.



MAC Address	MAC address of the peer.
Vendor	Vendor name from the main menu.
Port Description	String value from the **Comments** field for the interface on SonicWall Security Appliances.
System Capabilities	String value representing the list of capabilities supported by the peer device

Displaying Statistics

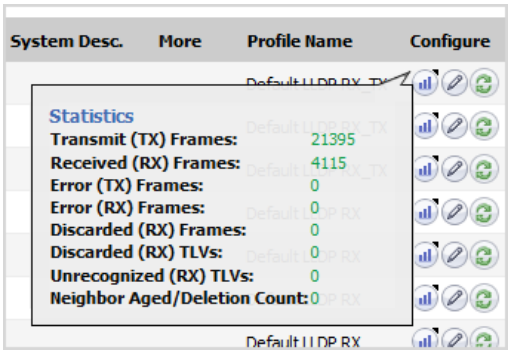
For each interface, you can display the number of:

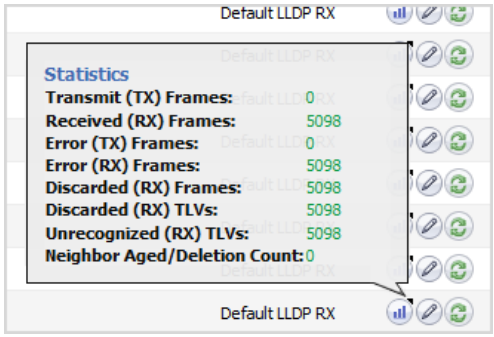
- Transmitted, received, erroneous, and discarded frames.
- Discarded and unrecognized TLVs.
- Aged or deleted neighbors.

To display an interface's statistics:

1 Navigate to **MANAGE | System Setup > Switching > Port Mirroring**.

2 In the **L2 Discovery** table, mouse over the **Statistics** icon for the interface. The **Statistics** popup displays.



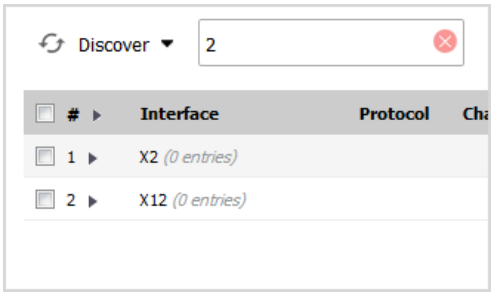


Searching the L2 Discovery Table

To limit the interfaces displayed in the L2 Discovery table:

1 Navigate to **MANAGE | System Setup > Switching > Port Mirroring**.

2 In the **Search** field, enter the search criterion. The display changes.



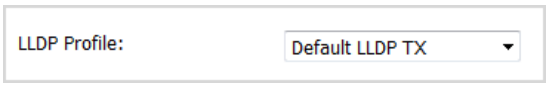
3 To clear the search and redisplay the entire table. click the red **Delete** icon in the **Search** field.

Associating an LLDP Profile with a L2 Discovery Interface

To associate an LLDP profile to a L2 Discovery interface:

1 Navigate to **MANAGE | System Setup > Switching > Port Mirroring**.

2 Click the **Edit** icon in the **Configure** column for the interface. The **Discover on Interface** dialog displays.



3 Select the default or custom profile from **LLDP Profile**:

- **Default LLDP Disabled**
- **Default LLDP RX** (default)
- **Default LLDP TX**
- **Default LLDP RX_TX**
- Custom profile

4 Click **Save**. The name of the profile displays in the **Profile Name** column of the **L2 Discovery** table.

Refreshing the Page

To refresh the data displayed on the page:

1 Click the **Refresh** icon above the **L2 Discovery** table.

Globally Enabling/Disabling LLDP

By default, LLDP is enabled globally. You can toggle the LLDP switch to enable or disable LLDP transmit and receive globally.

To globally enable/disable LLDP:

1 Navigate to **MANAGE | System Setup > Switching > Port Mirroring**.

2 Click **LLDP** above the **L2 Discovery** table. A confirmation message displays.





3 Click **OK**.

Discovering Neighbors

You can discover neighbors for:

- A single interface.
- Multiple interfaces.
- All interfaces.

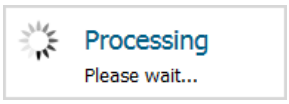
TIP: For LAG with trunk mode, all ports can discover neighbors; LAG with PortShield mode learns neighbors only under the aggregator port.

To discover neighbors for a single interface:

1 Navigate to **MANAGE | System Setup > Switching > Port Mirroring**.

2 Click the **Refresh** icon in the **Configure** column for the interface.

A processing message displays.



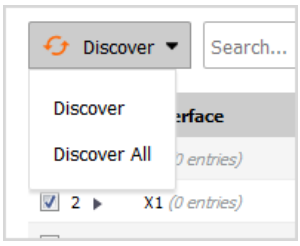
The information for the interface is updated.

To discover neighbors for multiple interfaces:

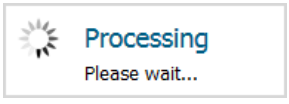
1 Navigate to **MANAGE | System Setup > Switching > Port Mirroring**.

2 Select the interfaces in the **L2 Discovery** table.

3 Select **Discover** from **Discover** above the table. This option is dimmed unless an interface is selected.



A processing message displays.



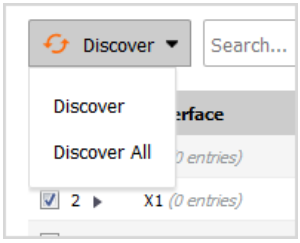
The information for the interfaces is updated.

To discover neighbors for all interfaces:

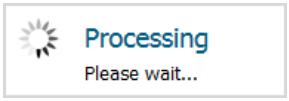
1 Navigate to **MANAGE | System Setup > Switching > Port Mirroring**.

2 Select an interface in the **L2 Discovery** table.

3 Select **Discover All** from **Discover** above the table.



A processing message displays.



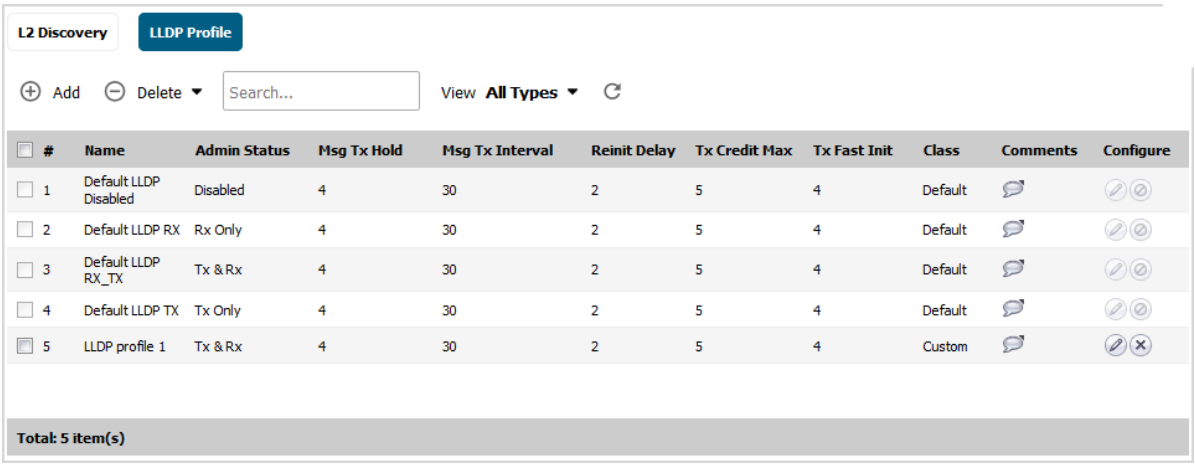
The information for all interfaces is updated.

Switching > Port Mirroring > LLDP Profile

Topics:

- Viewing LLDP Profiles on page 621
- Adding a LLDP Custom Profile on page 623
- Editing a Custom LLDP Profile on page 625
- Deleting Custom Profiles on page 625

Viewing LLDP Profiles



Name	Default or custom profile name.
Admin Status	LLDP mode of the LLDP profile: • **Disabled** • **Rx Only** • **TX & RX** • **Tx Only**
Msg Tx Hold	Multiplier of the time interval between transmissions during normal transmission periods to determine the time-to-live of LLDP frames for the LLDP profile.
Msg TX Interval	Time interval, in timer ticks, between transmissions during normal transmissions periods for the LLDP profile.
Reinit Delay	Indicates the amount of delay from when **Admin Status** becomes **Disabled** until reinitialization of the profile can be attempted again.
TX Credit Max	Maximum transmission credit for the LLDP profile.
Tx Fast Init	Determines the number of LLDPDUs sent during fast transmission periods.
Class	Type of profile: • **Default** • **Custom**
Comments	Information icon that, when moused over, displays either: • **Auto-added LLDP Profile** for default profiles. • Comment specified when the custom profile as added or edited; if no comment was specified, nothing is displayed
Configure	Contains the **Edit** and **Delete** icons for custom profiles. Default profiles cannot be edited or deleted, so the icons are dimmed.

Topics:

- Refreshing the LLDP Profile Table Information on page 622
- Searching the LLDP Profile Table on page 622
- Viewing only Certain Types of LLDP Profiles on page 623

Refreshing the LLDP Profile Table Information

To refresh LLDP profile information:

1 Navigate to **MANAGE | System Setup > Switching > Port Mirroring**.

1 Click **LLDP Profile**.

2 Click the **Refresh** icon above the **LLDP Profile** table.

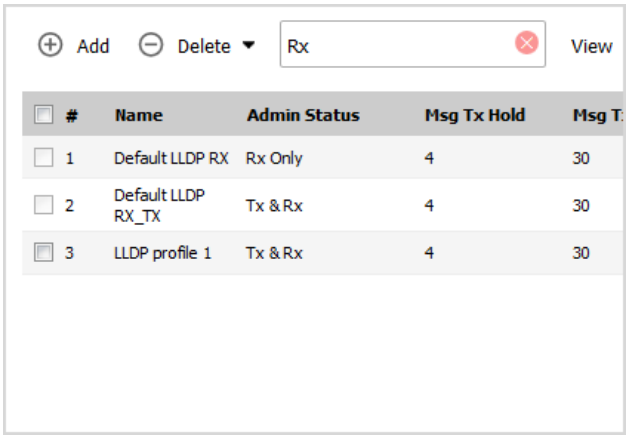
Searching the LLDP Profile Table

To limit the profiles displayed in the LLDP Profile table:

1 Navigate to **MANAGE | System Setup > Switching > Port Mirroring**.

2 Click **LLDP Profile**.

3 In the **Search** field, enter the search criterion. The display changes.



4 To clear the search and redisplay the entire table. click the red **Delete** icon in the **Search** field.

Viewing only Certain Types of LLDP Profiles

To display only certain types of LLDP Profiles:

1 Navigate to **MANAGE | System Setup > Switching > Port Mirroring**.

2 Click **LLDP Profile**.

3 From **View**, select the type of LLDP Profile to view:

- **All Types** (default)
- **Default**
- **Custom**

Adding a LLDP Custom Profile

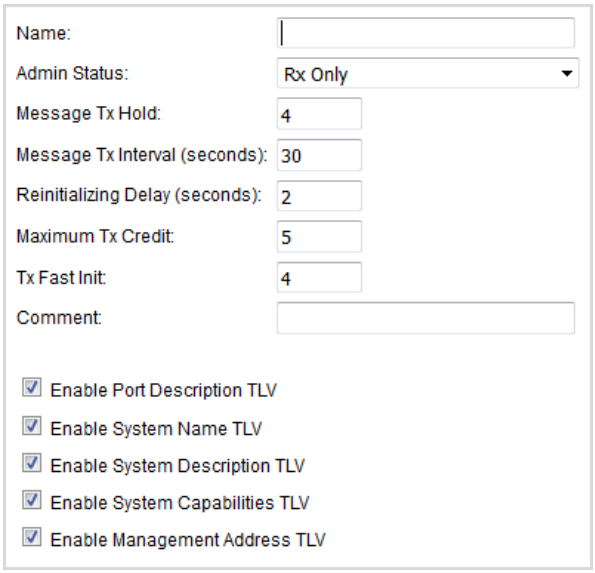
IMPORTANT: Changing default values affects the duration and the number of frames transmitted during each cycle.

To add an LLDP custom profile:

1 Navigate to **MANAGE | System Setup > Switching > Port Mirroring**.

2 Click **LLDP Profile**.

3 Click **Add**. The **Add LLDP Profile** dialog displays.



4 In the **Name** field, enter a descriptive name for the LLDP profile.

5 From **Admin Status**, select transmission mode for the LLDP profile:

- **Disabled**
- **Rx Only** (default)
- **Tx Only**
- **Tx & Rx**

6 To change the message transmission interval to determine the value of the TTL value of the LLDP frames transmitted by the LLDP agent, enter a multiplier in the **Message Tx Hold** field. The minimum value is 1, the maximum is 100, and the default value is **4**.

7 To define the time interval, in timer ticks, between transmissions during normal transmission periods, enter the interval, in seconds, in the **Message Tx Interval** field. The minimum is 1 second, the maximum is 3600 seconds, and the default is **30** seconds.

8 To specify the amount of delay from when Admin Status becomes Disabled until reinitialization can be attempted again for the profile, enter the delay, in seconds, in the **Reinitializing Delay** field. The minimum is 1 second, the maximum is 10 seconds, and the default, and recommended delay, is **2** seconds.

9 To specify the maximum number of LLDPDUs that can be transmitted at any time for the LLDP profile, enter the number in the **Maximum Tx Credit** field. The minimum is 1, the maximum is 10, and the default is **5**.

10 To specify an initial number of LLDPDUs transmitted during a fast transmission period, enter the number in the **Tx Fast Init** field. The minimum is 1, the maximum is 8, and the default is **4**.

11 Enter an optional comment in the **Comment** field. What you enter here displays when you mouse over the Information icon in the **Comments** column of the **LLDP Profile** table.

12 To include the port description of the Security Appliance in the optional TLV of a LLDPDU message, select **Enable Port Description TLV**. This option is selected by default.

13 To include the configured Firewall Name of the Security Appliance in the optional TLV if a LLDPDU message, select **Enable System Name TLV**. This option is selected by default.

14 To include **Firewall** as the identification of the Security Appliance in the optional TLV if a LLDPDU message, select **Enable System Description TLV**. This option is selected by default.

15 To include an IPv4 or MAC address used for managing an interface of the Security Appliance in the optional TLV if a LLDPDU message, select **Enable Management Address TLV**. This option is selected by default.

16 Click **OK**.

Editing a Custom LLDP Profile

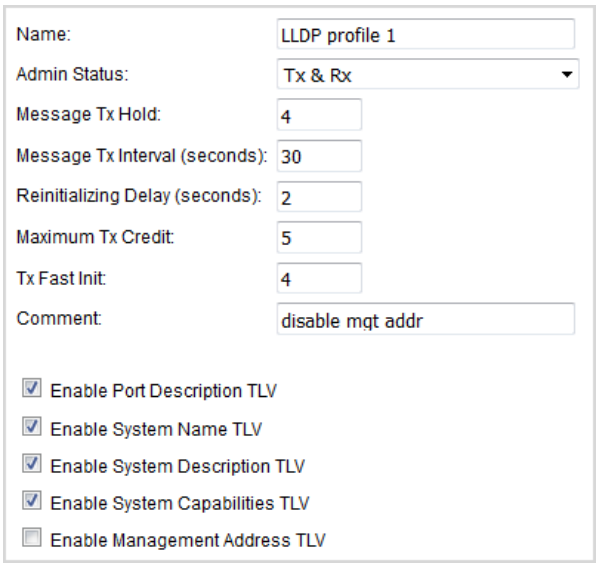
TIP: Default LLDP profiles cannot be edited.

To edit a custom LLDP profile:

1 Navigate to **MANAGE | System Setup > Switching > Port Mirroring**.

2 Click **LLDP Profile**.

3 Click the **Edit** icon for the profile. The **Edit LLDP Profile** dialog displays.



4 Make changes as needed. For information about the options, see Adding a LLDP Custom Profile on page 623.

5 Click **OK**.

Deleting Custom Profiles

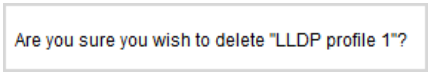
TIP: Default profiles cannot be deleted.

To delete a custom profile:

1 Navigate to **MANAGE | System Setup > Switching > Port Mirroring**.

2 Click **LLDP Profile**.

3 Click the **Delete** icon for the profile. A confirmation message displays.



4 Click **OK**.

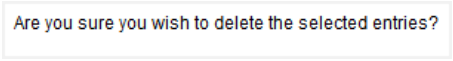
To delete one or more custom profiles:

1 Navigate to **MANAGE | System Setup > Switching > Port Mirroring**.

2 Click **LLDP Profile**.

3 Select the profiles to be deleted.

4 Click the **Delete Selected** from **Delete**. A confirmation message displays.



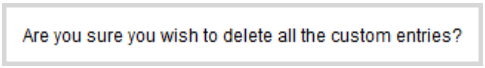
5 Click **OK**.

To delete all custom profiles:

1 Navigate to **MANAGE | System Setup > Switching > Port Mirroring**.

2 Click **LLDP Profile**.

3 Click the **Delete All** from **Delete**. A confirmation message displays.



4 Click **OK**.

42

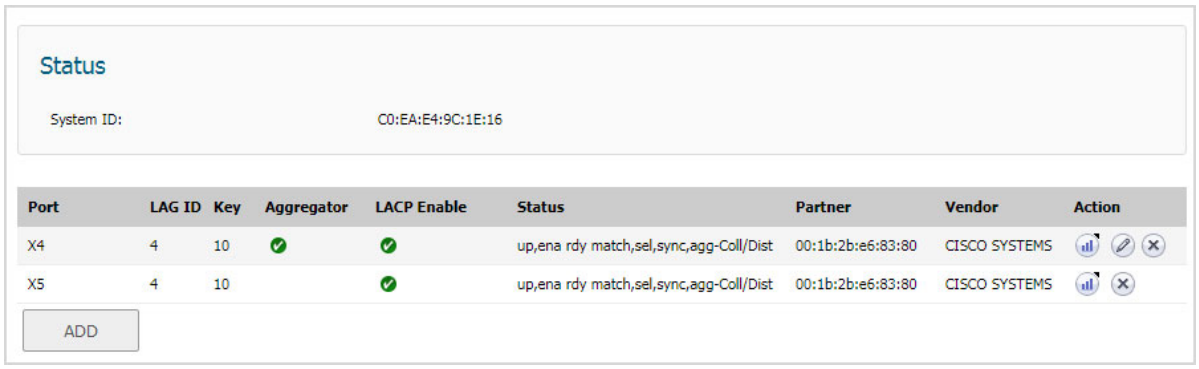
Configuring Link Aggregation

NOTE: Switching is available on all NSA 2650 and above and SuperMassive appliances.

Topics:

- **Switching > Link Aggregation** on page 627
  - About Link Aggregation on page 627
  - Viewing Link Aggregation on page 629
  - Creating a Logical Link (LAG) on page 631
  - Deleting a LAG on page 632

Switching > Link Aggregation



Topics:

- About Link Aggregation on page 627
- Viewing Link Aggregation on page 629
- Creating a Logical Link (LAG) on page 631

About Link Aggregation

NOTE: Link Aggregation (LAG) is supported for NSA 3600 and higher firewalls.

Link Aggregation allows port redundancy and load balancing in Layer 2 networks by allowing you to inter-connect SonicWall Security Appliances with two or more links between them in such a way that the multiple links are combined into one larger virtual pipe that can carry a higher combined bandwidth. As multiple links are

present between two devices, if one link fails, the traffic is transferred through other links without disruption. With multiple links present, traffic also can be load balanced in such a way to achieve even distribution. Load balancing is controlled by the SonicWall Security Appliance, based on source and destination MAC address pairs. The **Switching > Link Aggregation** page provides information and statistics about and allows configuration of interfaces for aggregation.

SonicOS supports the two types of LAG:

- Static LAG on page 628
- Dynamic LAG on page 628

Static LAG

In Static Link Aggregation, ports that are in the same VLAN (same PortShield Group) or are VLAN trunk ports are eligible for link aggregation. Up to four ports can be aggregated in a logical group, and there can be four Logical Links (LAGs) configured. With Static Link Aggregation, all configuration settings are set on both participating LAG components.

Two main types of usage are enabled by this feature:

Firewall to Server	Implemented by enabling Link Aggregation on ports within the same VLAN (same PortShield Group). This configuration allows port redundancy, but does not support load balancing in the appliance-to-Server direction due to a hardware limitation on the Security Appliance.
Firewall to Switch	Allowed by enabling Link Aggregation on VLAN trunk ports. Load balancing is performed automatically by the hardware. The Security Appliance supports one load balancing algorithm based on source and destination MAC address pairs.

Similarly to PortShield configuration, you select an interface that represents the aggregated group. This port is called an *aggregator*. The aggregator port must be assigned a unique key. Non-aggregator ports can be optionally configured with a key, which can help prevent an erroneous LAG if the switch connections are wired incorrectly.

NOTE: The key is not the same as the LAG ID, which is the same as the interface number and cannot be changed. The key must be assigned when the LAG group is configured. All the non-aggregator ports should have the same key as the aggregator port.

Ports bond together if connected to the same link partner and their keys match. A link partner cannot be discovered for Static link aggregation. In this case, ports aggregate based on keys alone.

Like a PortShield host, the aggregator port cannot be removed from the LAG as it represents the LAG in the system.

NOTE: After link aggregation has been enabled on VLAN trunk ports, additional VLANs cannot be added or deleted on the LAG.

Dynamic LAG

SonicOS supports Dynamic Link Aggregation using Link Aggregation Control Protocol (LACP defined by IEEE 802.3ad) on all SonicWall Security Appliances that support Advanced Switching features.

About Dynamic LAG using LACP

LACP allows the exchange of information related to link aggregation between the members of the LAG group in protocol packets called Link Aggregation Control Protocol Data Units (PDUs). with LACP, errors in configuration, wiring, and link failures can be detected quickly.

The two major benefits of LAG such as increased throughput and link redundancy can be achieved efficiently using LACP. LACP is the signaling protocol used between members in a LAG. It ensures links are only aggregated into a bundle if they are correctly configured and cabled. LACP can be configured in one of two modes:

- **Active mode** - Device immediately sends LACP PDUs when the port comes up.

  NOTE: SonicOS 6.5 only supports the Active mode of LACP.

- **Passive mode** -Port is placed in a passive negotiating state, in which the port only responds to LACP PDUs it receives, but does not initiate LACP negotiation.

If both sides are configured as active, LAG can be formed assuming successful negotiation of the other parameters. If one side is configured as active and the other one as passive, LAG can be formed as the passive port responds to the LACP PDUs received from the active side. If both sides are passive, LACP fails to negotiate the bundle. Passive mode is rarely used in deployments.

In the configuration, all member ports of the same LAG must be set up on the same VLAN as the Aggregator port. Data packets received on the LAG members are associated with the parent Aggregator port using the VLAN. When the state of the Aggregator/member ports of a LAG reaches a stable Collection/Distribution state, the ports are ready to transmit and receive data traffic.

All information related to LAG, such as the Aggregator ports configured, this information is displayed on the **Switching > Link Aggregation** page:

- Member ports that are part of the LAG.
- Status of each of the ports that form the LAG.
- The Partner MAC address received through LACP.

Six load balancing options are available for configuration. The load balancing option must be chosen when creating a LAG along with the Aggregator port.

IMPORTANT: You cannot modify the load balancing option after the LAG is created.

VLAN Enhancements for LAG

NOTE: This enhancement is not supported on the NSA 2600, TZ Series, or SOHO W firewalls.

Link Aggregation (LAG) allows you to inter-connect devices with two or more links between them in such a way that the multiple links are combined into one larger virtual pipe that can carry a higher combined bandwidth. As multiple links are present between two devices, if one link fails, the traffic is seamlessly transferred through other links without disruption. With multiple links being present, traffic also can be load balanced in such a way to achieve even distribution.

With this enhancement;

- LAG does not have to be dismantled or removed before the VLAN is added/deleted. The configuring allows you to add the VLAN to an existing LAG or delete the VLAN from an existing LAG without disrupting the current traffic related to the LAG or other VLANs configured on the LAG.
- VLAN can be added to/deleted from any member of the LAG and it gets applied to all the other members of the LAG automatically without the need to explicitly add to/delete from other members of the LAG.

Viewing Link Aggregation

Topics:

- Viewing Status on page 630

- Viewing Link Aggregation Ports on page 630

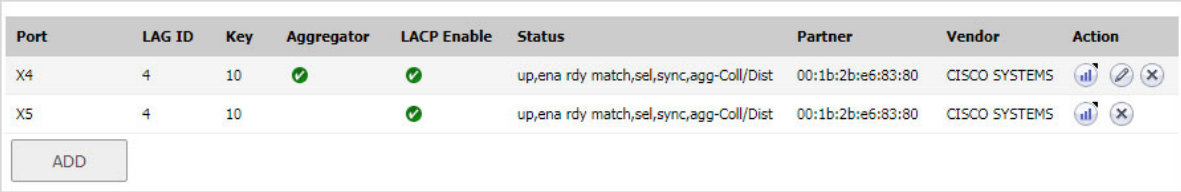
Viewing Status

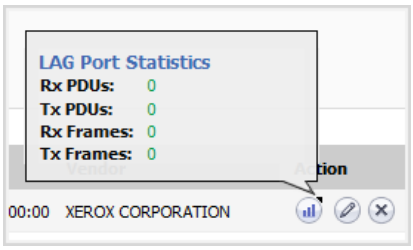


The **Status** table displays the MAC address System ID for the firewall.

Viewing Link Aggregation Ports

To view Link Aggregation Ports, navigate to **System Setup | Switching > Link Aggregation**



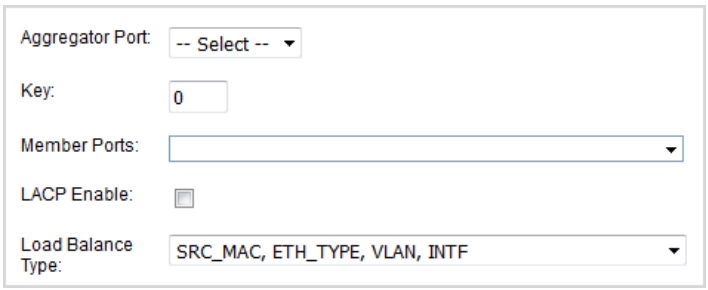
Port	Interface used as an aggregator port or a member port
LAG ID	System-configured link aggregator. A port that is not an aggregator has a LAG ID of the aggregator of which it is a member.
Key	Indicates port membership from the **Add LAG Port** dialog.
Aggregator	Indicates an aggregator port by a green checkmark; otherwise, it is blank.
LACP Enable	Indicates whether LACP is enabled.
Status	Indicates whether the port is **up** or **down**.
Partner	MAC addresses of the link partners after they are physically connected; for • Static LAG, displays **00 : 00 : 00 : 00 : 00 : 00** • Dynamic LAG, displays the partner's MAC address
Vendor	Displays the name of the equipment manufacturer.
Action	Displays these icons: • **Statistics** – when moused over, displays the **LAG Port Statistics** popup:  • **Edit** (only an Aggregator port can be edited) • **Delete**

Creating a Logical Link (LAG)

To create a Logical Link (LAG):

1 Navigate to **Switching > Link Aggregation**.

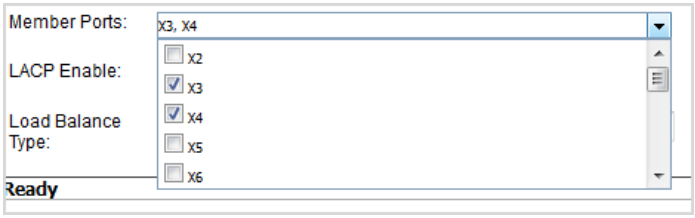
2 Click **Add**. The **Add LAG Port** dialog displays.



3 Select the interface from **Aggregator Port**.

4 Specify the port membership to an LAG group by entering the desired key into the **Key** field. The minimum value is 1, and the maximum value is 255. The field has a default value of **0**, which must be replaced.

5 Select the ports to be aggregated from the **Member Ports** drop-down menu. You can select any number of ports in the list by selecting the checkbox for each port to be aggregated.



NOTE: The listed ports depend on the interface chosen in Step 3.

6 To enable Link Aggregation Control Protocol (LACP) for this port, select **LACP Enable**. This option is not selected by default.

7 From **Load Balance Type**, select the how load balancing is performed:

IMPORTANT: You cannot modify the load balancing option after the LAG is created.

- SRC_MAC, ETH_TYPE, VLAN, INTF (default)
- DST_MAC, ETH_TYPE, VLAN, INTF
- SRC_MAC, DST_MAC, ETH_TYPE, VLAN, INTF
- SRC_IP, SRC_PORT
- DST_IP, DST_PORT
- SRC_IP, SRC_PORT, DST_IP, DST_PORT

8 Click **OK**.

Deleting a LAG

To delete a member of a LAG:

1 Navigate to **System Setup | Switching > Link Aggregation**.

2 Delete the member port of the lag by clicking its **Delete** icon.

To delete an aggregator port:

1 Navigate to **System Setup | Switching > Link Aggregation**.

2 Delete all the member ports by clicking their **Delete** icons.

> **NOTE:** All member ports must be deleted from the LAG before deleting the Aggregator port.

3 Delete the aggregator port by clicking its **Delete** icon.

43

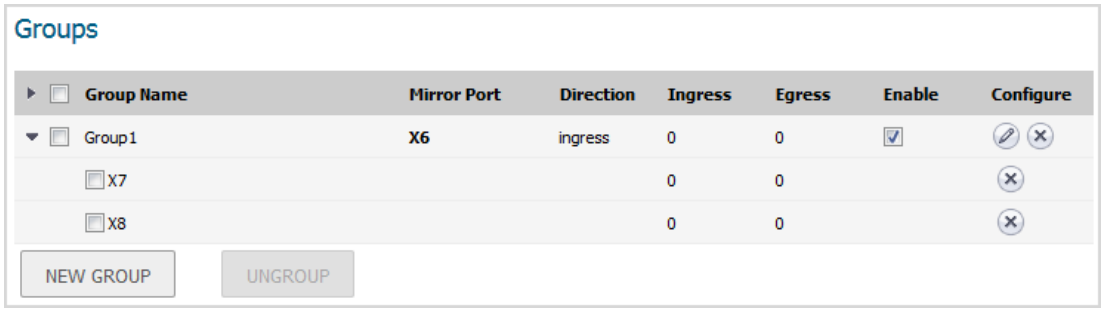
Configuring Port Mirroring

NOTE: Switching is available on all NSA 2650 and above and SuperMassive appliances.

Topics:

- **Switching > Port Mirroring** on page 633
  - About Port Mirroring on page 633
  - Viewing Mirrored Ports on page 634
  - Configuring a Port Mirroring Group on page 634
  - Editing a Port Mirroring Group on page 636
  - Deleting Port Mirroring Groups on page 636

Switching > Port Mirroring



Topics:

- About Port Mirroring on page 633
- Viewing Mirrored Ports on page 634
- Configuring a Port Mirroring Group on page 634
- Editing a Port Mirroring Group on page 636
- Deleting Port Mirroring Groups on page 636

About Port Mirroring

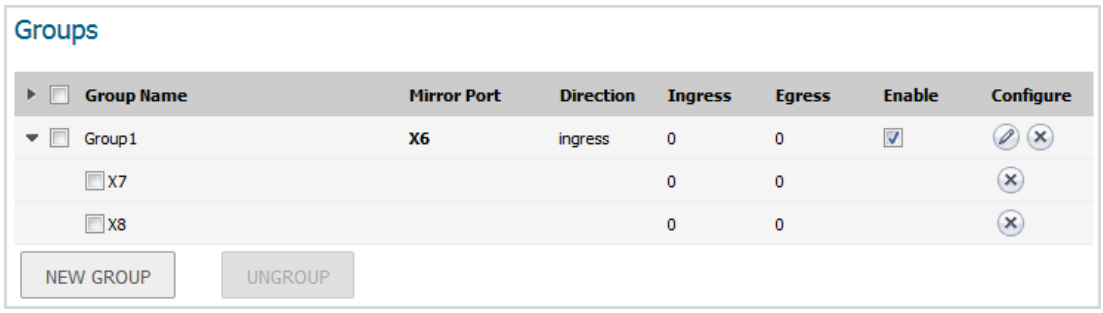
You can configure Port Mirroring on SonicOS to send a copy of network packets seen on one or more switch ports (or on a VLAN) to another switch port called the mirror port. By connecting to the mirror port, you can monitor the traffic passing through the mirrored port(s).

On the NSA 2650, a VLAN trunk port can be a mirror or a mirrored port. For all other platforms, a VLAN trunk port can be mirrored, but cannot act as a mirror port itself.

MANAGE | System Setup | Switching > Port Mirroring allows you to assign mirror ports to mirror ingress, egress or bidirectional packets coming from and/or to a group of ports.

Viewing Mirrored Ports

You monitor traffic on the mirrored port(s) by connecting to the mirror port.



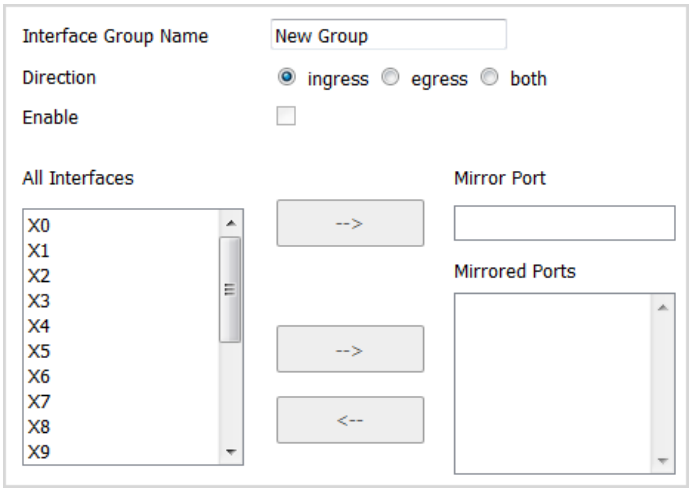
Group Name	Name of the interface group.
Mirror Port	Interface used as the mirror port, that is, the port that monitors other ports on the selected direction.
Direction	Direction of the traffic being mirrored: • **both** (bidirectional) • **ingress** • **egress**
Ingress	Number of packets arriving on the mirrored port(s). For egress-only ports, this is always 0.
Egress	Number of packets sent out on the mirrored port(s). For ingress-only ports, this is always 0.
Enable	Indicates whether mirroring is enabled – a checkmark is in the checkbox – or disabled – the checkbox is blank – for the group.
Configure	Contains the **Edit** and **Delete** icons for the group entry and a **Delete** icon for each port in the group.

Configuring a Port Mirroring Group

To create a new port mirroring group:

1 Navigate to **Switching > Port Mirroring**.

2 Click **New Group**. The **Edit Mirror Group** dialog displays.



3 Enter a descriptive name for the group into the **Interface Group Name** field. The default name is **New Group**.

4 For the **Direction**, select one of the following:

- **ingress** – Monitors traffic arriving on the mirrored port(s).
- **egress** – Monitors traffic being sent out on the mirrored port(s).
- **both** – Monitors traffic in both directions on the mirrored port(s).

5 From the **All Interfaces** list:

a Select the port to mirror the traffic to. You must use an unassigned port as the mirror port.

b Click the top **Right Arrow** to move the port to the **Mirror Port** field.

6 From the **All Interfaces** list:

a Select one or more ports to be monitored. You monitor traffic on the mirrored port(s) by connecting to the mirror port.

b Click the lower **Right Arrow** to move it/them to the **Mirrored Ports** list.

7 To enable port mirroring for these ports, select **Enable**.

NOTE: Only one ingress group and one egress group can be enabled at one time. If a group has both directions and it is enabled, the individual ingress and egress groups or another group with both directions cannot be enabled. The individual ingress and egress groups can be enabled separately.

This option is dimmed until you specify the mirror port and its mirrored ports.

8 Click **OK**.

Enabling a Mirrored Group

If you did not enable the mirrored group when you created it, you can enable mirroring on **Groups** table by selecting **Enable** for the mirrored group.

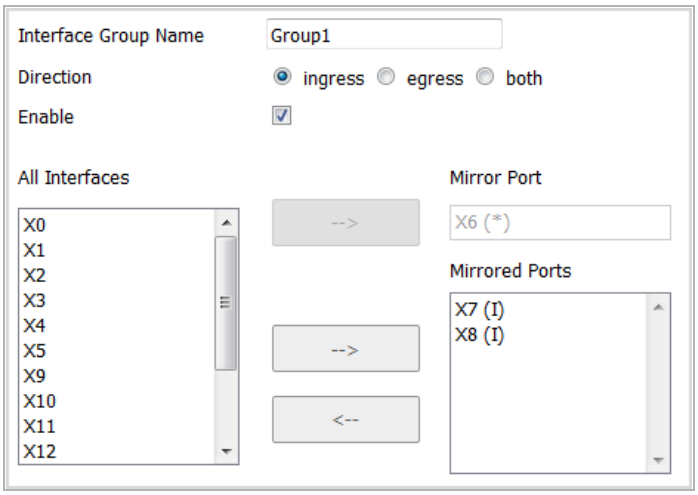
Editing a Port Mirroring Group

You can edit all attributes of a mirrored group except the mirror port, which is dimmed.

To edit a port mirroring group:

1 Navigate to **Switching > Port Mirroring**.

2 Click the **Edit** icon of the mirror port. The **Edit Mirror Group** dialog for the group displays.



3 Make the changes to any of the options.

> **NOTE:** You can add or delete mirrored ports, but not the mirror port itself. If you delete a member of the group, no confirmation message is displayed.

4 If mirroring has been enabled for the group, **Enable** is selected. To disable port mirroring for these ports, deselect **Enable**.

> **NOTE:** Only one ingress group and one egress group can be enabled at one time. If a group has both directions and it is enabled, the individual ingress and egress groups or another group with both directions cannot be enabled. The individual ingress and egress groups can be enabled separately.

5 Click **OK**.

Deleting Port Mirroring Groups

You can delete members of a mirror group, a mirror group, multiple groups, or all groups.

Topics:

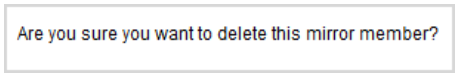
- Removing Port Group Members on page 637
- Removing a Port Mirror Group on page 637
- Removing Multiple Port Mirror Groups on page 637
- Removing All Port Mirror Groups on page 638

Removing Port Group Members

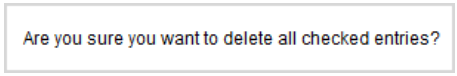
You can delete a member of a port group as described in Editing a Port Mirroring Group on page 636 or you can delete it in the **Groups** table.

To remove a member of a Port Group in the Groups table:

1. Navigate to **Switching > Port Mirroring**.
2. Display the group members by clicking the group's **Expand** button.
3. Either:
   - Click the **Delete** icon for the member(s) to be deleted. A confirmation message displays.

   - Click one or more checkboxes of the members to be deleted, and then click **Ungroup**. A confirmation message displays.

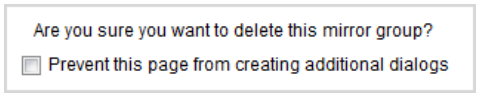
     

4. Click **OK**.

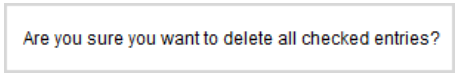
Removing a Port Mirror Group

To remove a port mirror group in the Groups table:

1. Either:
   - Click the **Delete** icon for the group to be deleted. A confirmation message displays:

   - Select the checkbox for the group and then click **Ungroup**. A confirmation message displays:

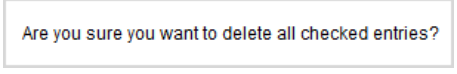
     

2. Click **OK**.

Removing Multiple Port Mirror Groups

To remove multiple port mirror groups:

1. In the **Groups** table, select the checkbox next to the port mirror groups you want to delete.
2. Click **Ungroup**. A confirmation dialog displays.

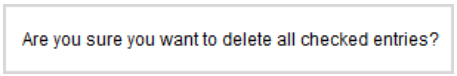
3. Click **OK**.

Removing All Port Mirror Groups

To remove all port mirror groups:

1 In the **Groups** table, select the checkbox in the table heading.

2 Click **Ungroup**. A confirmation dialog displays.



3 Click **OK** in the confirmation dialog.

44

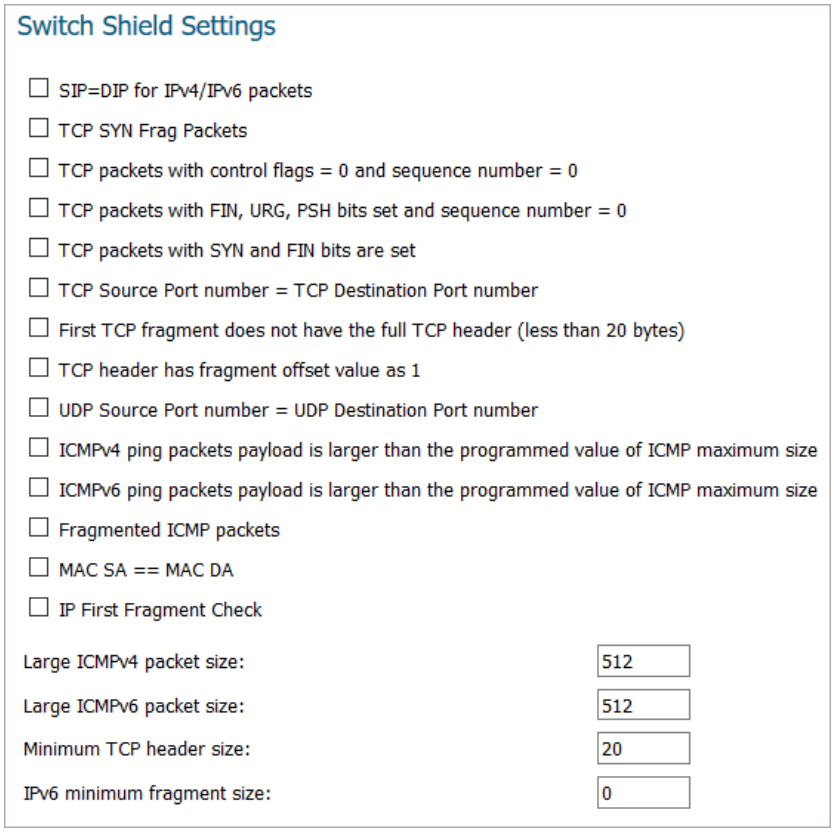
Configuring Switch Shield

NOTE: Switching is available on all NSA 2650 and above firewalls.

Topics:

- Switching > Switch Shield on page 639

Switching > Switch Shield



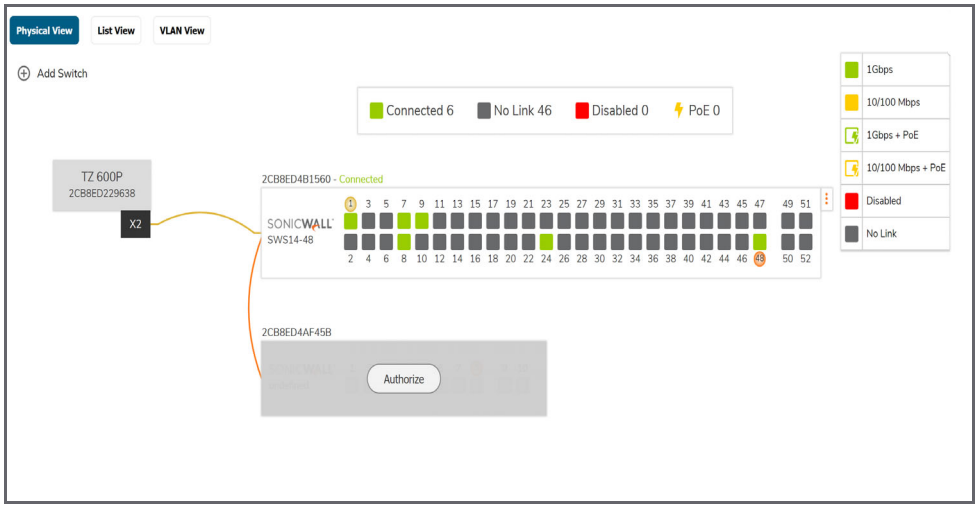
To configure Switch Shield settings:

1. Navigate to **MANAGE | System Setup > Switching > Switch Shield**.
2. Select the options for DDOS protection using switch capabilities. None of the options are selected by default.
3. Click **ACCEPT**.

45

Configuring the SonicWall Switch

Starting with release 6.5.4.6, SonicOS supports firewall management of SonicWall Switches. The Switch Controller user interface consists of pages and sub-pages.



Topics:

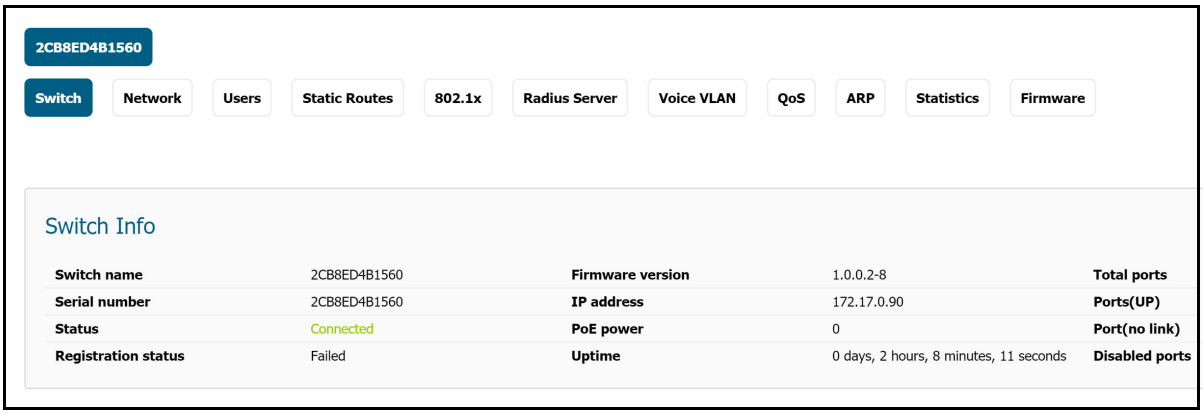
- Before Adding a Switch on page 641
- Checking Switch Details on page 641
- Adding a Switch to a Firewall with Zero-Touch on page 642
- Adding a Switch to a Firewall Manually on page 645
- Check the Status and Link Details on page 647
- Enabling the Switch on page 648
- Shutting Down the Switch on page 648
- Restarting the Switch on page 649
- Adding a VLAN on page 650
- Adding Static Routes on page 653
- Editing DNS on page 654
- Setting Up the Ports on page 655
- Setting Up QoS on page 656
- Setting Up PoE on page 658
- Setting Up Users on page 659
- Setting Up 802.1x Authentication on page 660
- Daisy-Chaining Switches on page 661
- Connecting Access Points on page 663
- Modifying the MAC Address Table on page 670
- Checking Port Statistics on page 671
- Setting Spanning Tree Protocol on page 672

Before Adding a Switch

- When adding a Switch manually, first check that it is configured to factory defaults. This can be ensured by depressing the reset Switch for 10 seconds or from the Switch Local UI, or the Command Line Interface.
- When adding a management link to a Switch manually, ensure that the DHCP lease range supports the default management IP address.
- The firewall interface linking to the Switch interface cannot be a PortShield host and no other firewall interface can be portshielded to it. The firewall interface linking to the Switch cannot be a PortShield group member, that is, it cannot be portshielded to another firewall interface.
- If the management link between the switch and firewall is isolated from data traffic, the switch must configured at a static IP address.
- Switches may be added into daisy-chained configurations manually or by using Zero-Touch.
- For daisy chaining Switches, consider setting up a common link (management and data) with sufficient capacity and do not make further connections from firewall to parent switch without configuring them. Make any other connections from the firewall to the Switch when you add the Switch.
- Make any changes in the Reserved VLAN range for the firewall interface before adding the SonicWall Switch. If the Reserved VLAN range changes after connecting the Switch, then the Switch must be removed and re-added.
- Logging into switch with SSHv1 client is not supported.
- If adding Switches to a High Availability (HA) pair:
    - Switches cannot be added to HA pairs with Zero-Touch.
    - To use the Switch with HA, you must first create an HA pair, and then manually add the Switch.

Checking Switch Details

Navigate to **Switches > Switch** to get a summary on switches connected to the firewall.



Adding a Switch to a Firewall with Zero-Touch

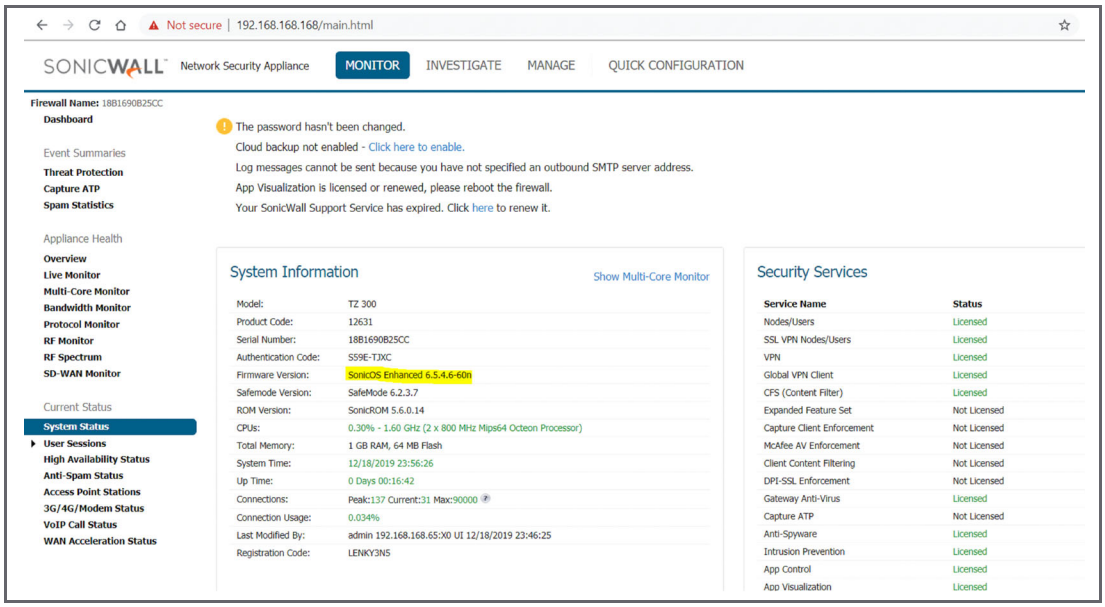
IMPORTANT: Please register your Switch before trying to add it to a firewall.

NOTE: In order for the firewall to sense the presence of the Switch, its firmware must be at SonicOS 6.5.4.6 or higher. The Switch should be at 1.0.0.0-19 or higher.

The firewall must be setup to add Switches with Zero Touch.

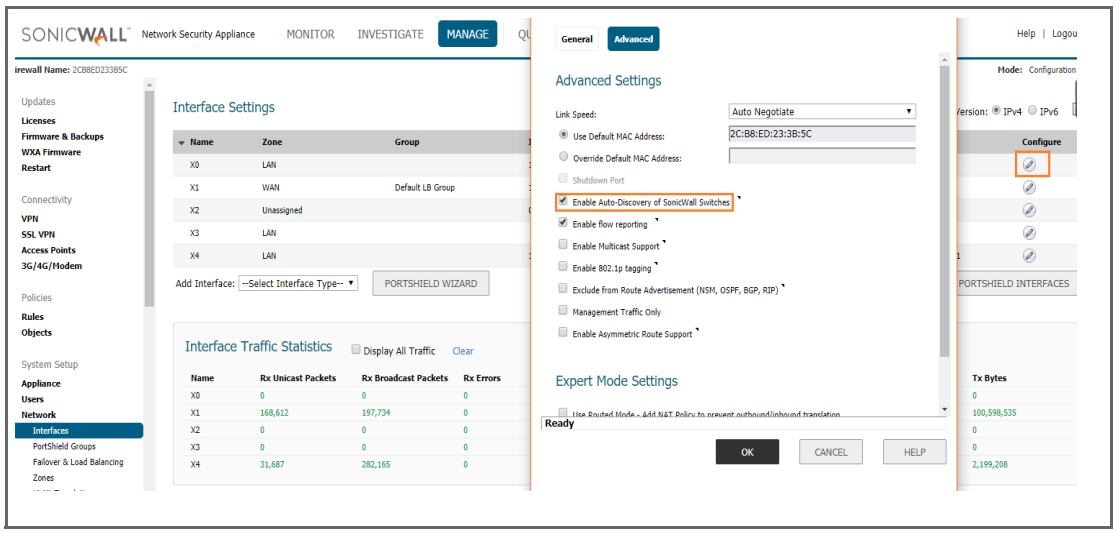
To prepare firewall:

1 Check that the firewall firmware is at the most recent level.



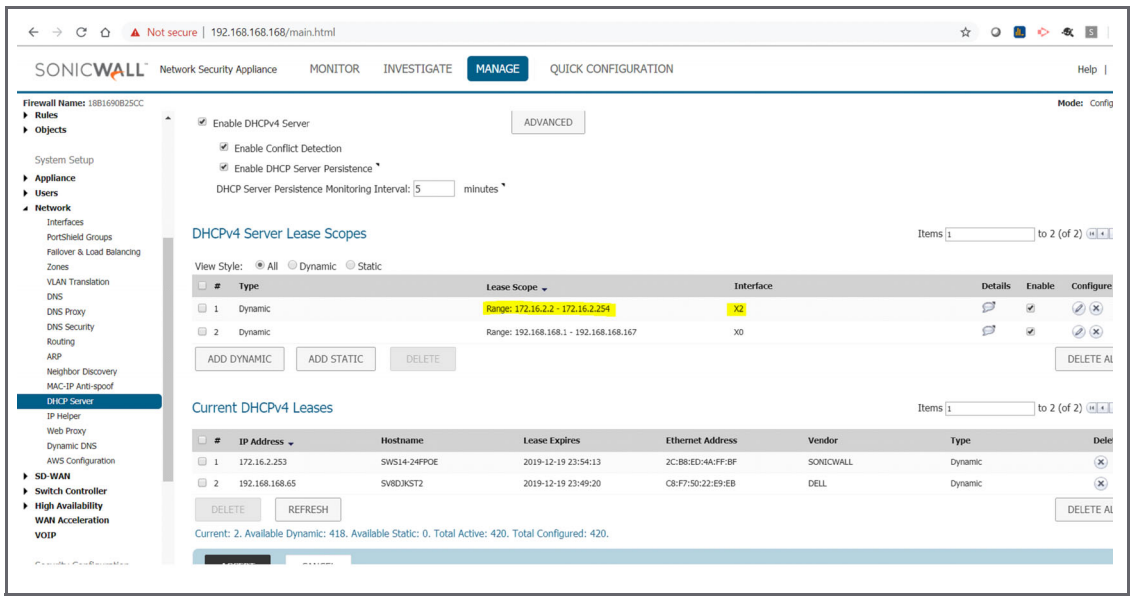
2 Select an interface on the firewall to connect to the Switch. Navigate to **Manage > Network > Interfaces** and select an interface, then click on **Configure** and select the **Advanced** tab.

Select **Enable AutoDiscovery of SonicWall Switches:**

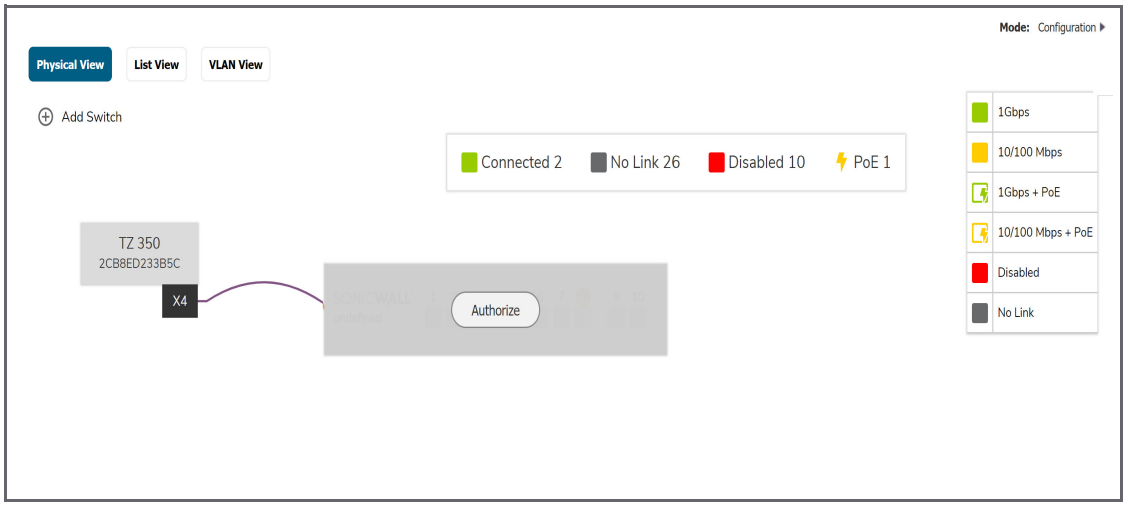


3 Connect the Switch to the selected firewall interface.

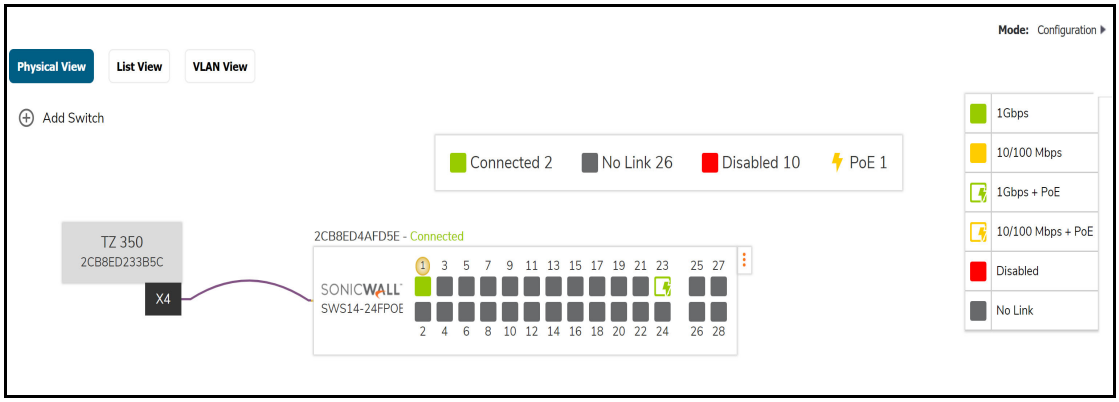
4 Navigate to **Network>DHCP Server** and verify **DHCPv4 Server Lease Scope** is enabled for the selected network.



5 Navigate to **Manage > Switch Controller > Overview**, Click on **Authorize** button to add the Switch to firewall:



6 The network topology will now appear on the **Overview.**



Adding a Switch to a Firewall Manually

IMPORTANT: Please register your Switch before trying to add it to a firewall.

Once an IP address is established as described in Adding a Switch to a Firewall Manually on page 645, only a few steps are necessary to set it up for management from a firewall.

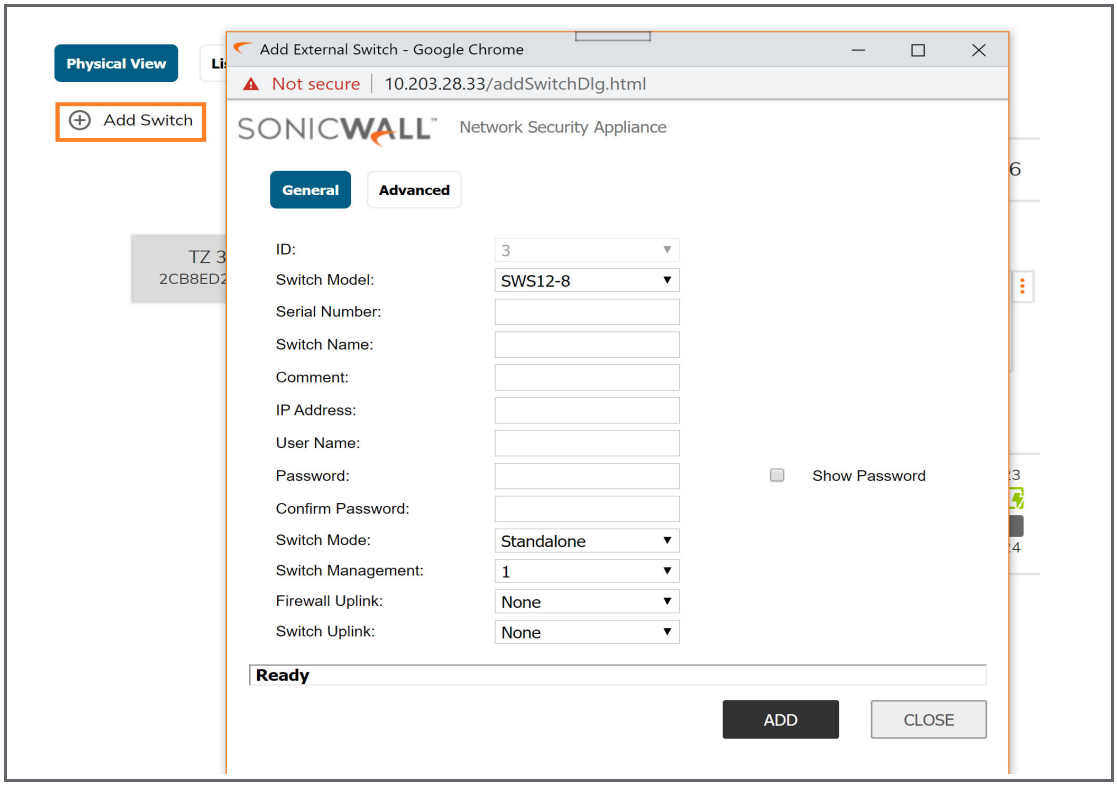
To Connect Switch to Firewall:

1 Connect an Ethernet cable (that is, RJ45 to RJ45) from a port on the Sonic Switch to an available port on the firewall.

NOTE: When adding a Switch manually, first check that it is configured to factory defaults. This can be ensured by depressing the reset switch for 10 seconds or more. The Switch can also be factory defaulted from the Switch Local UI, or the Command Line Interface accessible through the console port.

NOTE: To change the Reserved VLAN range on the firewall, do so before adding the SonicWall Switch. If the Reserved VLAN range changes after connecting the Switch, then the Switch must be removed and re-added.

2 Login to the management console and navigate to **MANAGE | System Setup > Switch Controller** and note the Switch-related pages. From Overview click on **Add Switch** as shown below.



NOTE: If the above options do not appear, check that your firewall and Switch are at the correct firmware release levels.

3 Click on **Switch** and when the Switch List appears click on **Add Switch**. The dialog box will appear:

- **ID**: The system will auto-assign a consecutive number here.

- **Switch Model**: Click on the down-pointing selection arrow and choose. Check the label on the bottom of your Switch if necessary.
- **Serial Number**: Check the label on the bottom of the Switch.
- **Switch Name**: Your choice.
- **Comment**: For optional reference.
- **IP Address:** What was established when the switch was added to the network.
- **User Name**: In most test cases, Admin. To add additional users, see Setting Up Users on page 659.
- **Password:** The password is configured on the Switch.
- **Switch Mode:** The choices are **Standalone** for one Switch per port or **Daisy Chain** when multiple Switches are added such that no Switch connects with more than two others.
- **Switch Management:** Management traffic flows on this interface.
- **Firewall Uplink:** This is the port on the firewall to which the Switch will connect.
- **Switch Uplink:** This is the port on the Switch connecting to the firewall.

Click on **Advanced** to access the Spanning Tree Protocol settings:

NOTE: Jumbo frames are not supported on TZ model firewalls.

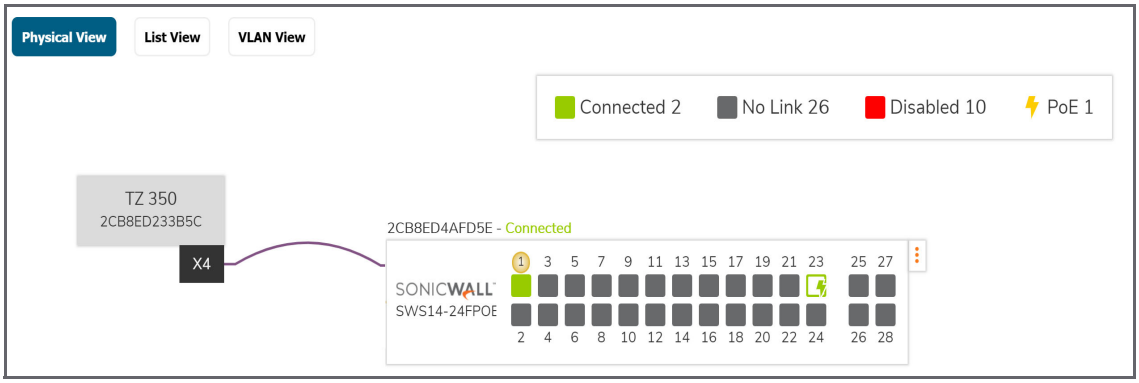
- **STP Mode**: **Rapid** or **Multiple**
- **STP State**: **Enable** or **Disabled**

  For more information on STP, see Setting Spanning Tree Protocol on page 671.
- **Jumbo Frame Size:** The default is the maximum standard transmission unit size in bytes. The default is 1522. Frame sizes larger than this are jumbo.

  NOTE: Jumbo frames are not supported on TZ model firewalls.

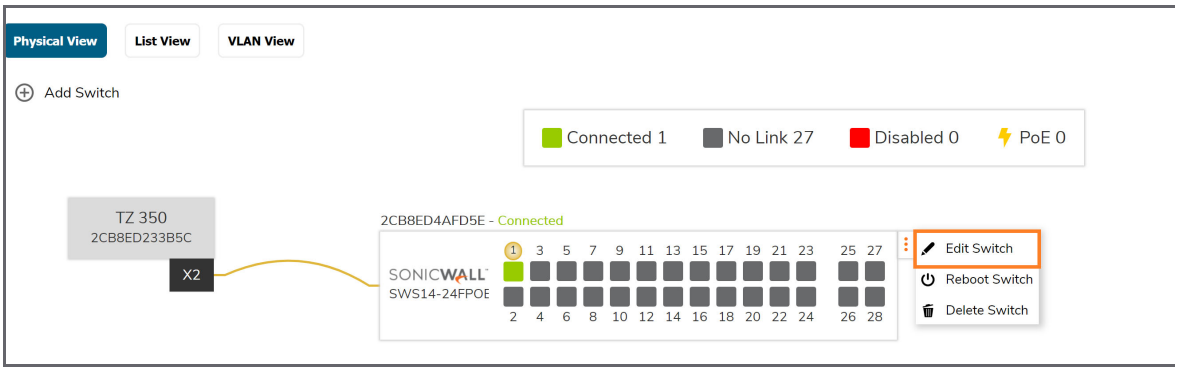
4 Once the dialog boxes are complete, click on **Add**.

5 Go to the **Overview**, the new Switch will appear graphically with the ports linking the Switch and the firewall indicated.



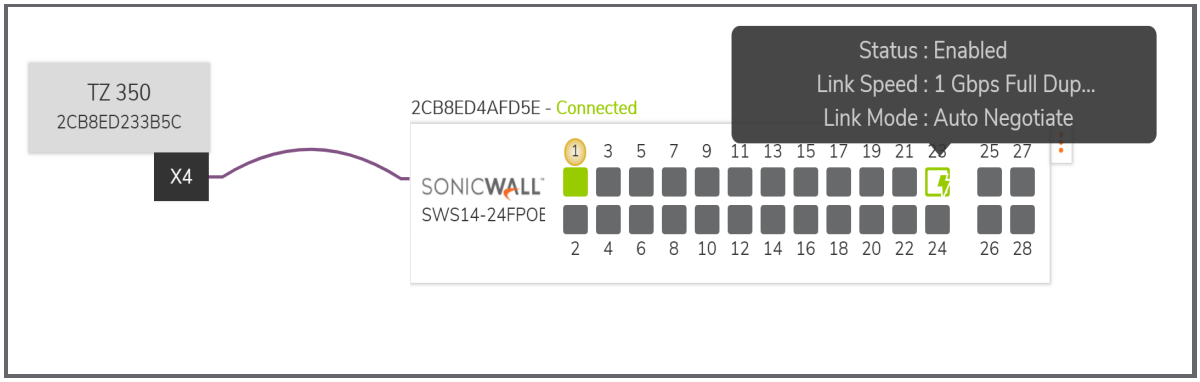
Changing the Switch Configuration

To edit the Switch configuration, go to the **Overview > Physical View** display and click on the “three-dot” symbol to the right of the switch graphic. Then click on b.



Check the Status and Link Details

Navigate to **Switch Controller > Overview** and hover over the Switch port.



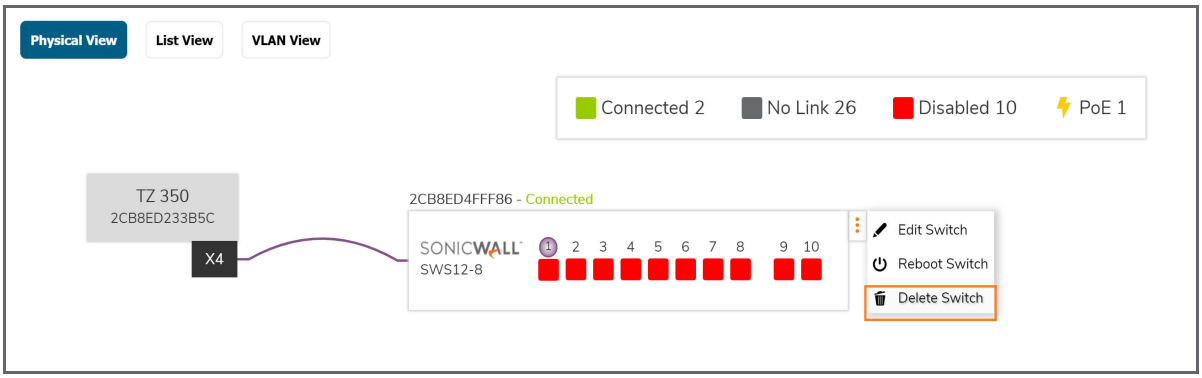
Enabling the Switch

If the firewall to Switch link is down, navigate to **Switch Controller > Overview** and click on 3 dot menu of the Switch which is off-line and then click on **Edit Switch** to bring up the Switch configuration dialog box. Check if the Switch configuration details are correct including: IP address, serial number, and Switch Management interface.



Shutting Down the Switch

To remove a Switch from a firewall, navigate to **Switch Controller > Overview > Physical View** and click on the **Delete Switch**.



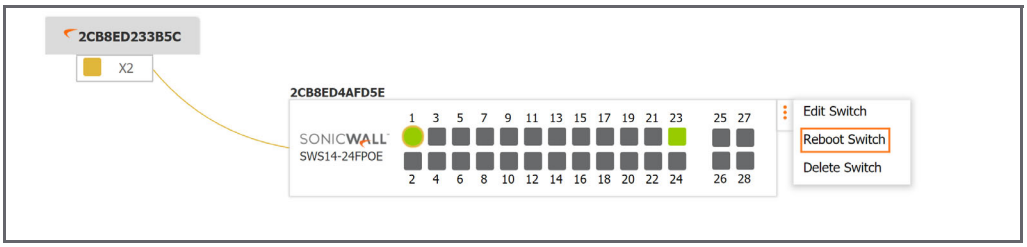
Restarting the Switch

To reboot the Switch:

1 Simply depress the recessed reset Switch on the front panel for a second.

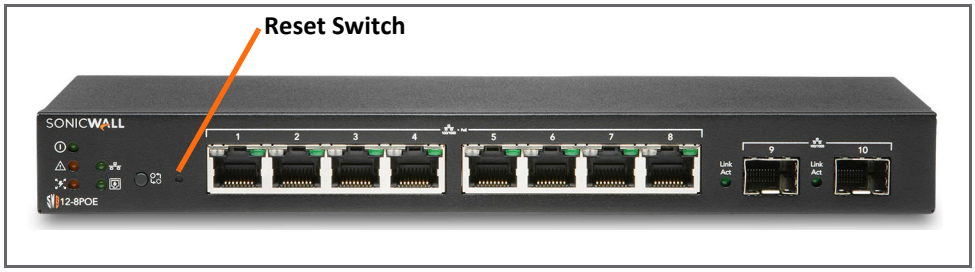
OR:

1 Click on the 3 dot menu on the Switch image on the Overview page and click on **Reboot Switch**.



To reboot the Switch to factory defaults:

1 Use a paper clip to depress the front panel reset Switch for over 10 seconds.



Adding a VLAN

A Virtual LAN (VLAN) is a group of ports that form a logical Ethernet segment on a Layer 2 Switch to provide better administration, security, and management of traffic. A VLAN is a network topology configured according to a logical scheme rather than a physical layout. When you use a VLAN, users can be grouped by logical function instead of physical location. All ports that frequently communicate with each other are assigned to the same VLAN, regardless of where their location in the network. VLANs let you logically segment your network into different broadcast domains allowing the grouping of ports with related functions into their own separate, logical LAN segments on the same Switch. This allows broadcast packets to be forwarded only between ports within the VLAN, thus avoiding broadcast packets being sent to all the ports on a single Switch. A VLAN also increases network performance by limiting broadcasts to a smaller, more manageable logical broadcast domain. By limiting traffic to specific broadcast domains, VLANs improve security.

Each VLAN in a network has an associated VLAN ID, which appears in the IEEE 802.1Q tag in the Layer 2 header of packets transmitted on a VLAN. The IEEE802.1Q specification establishes a standard method for tagging Ethernet frames with VLAN membership information. The key for IEEE802.1Q to perform its functions is in its tags. 802.1Q-compliant Switch ports can be configured to transmit tagged or untagged frames. A tag field containing VLAN information can be inserted into an Ethernet frame. When using 802.1Q VLAN configuration, you configure ports to be a part of a VLAN group. When a port receives data tagged for a VLAN group, the data is discarded unless the port is a member of the VLAN group.

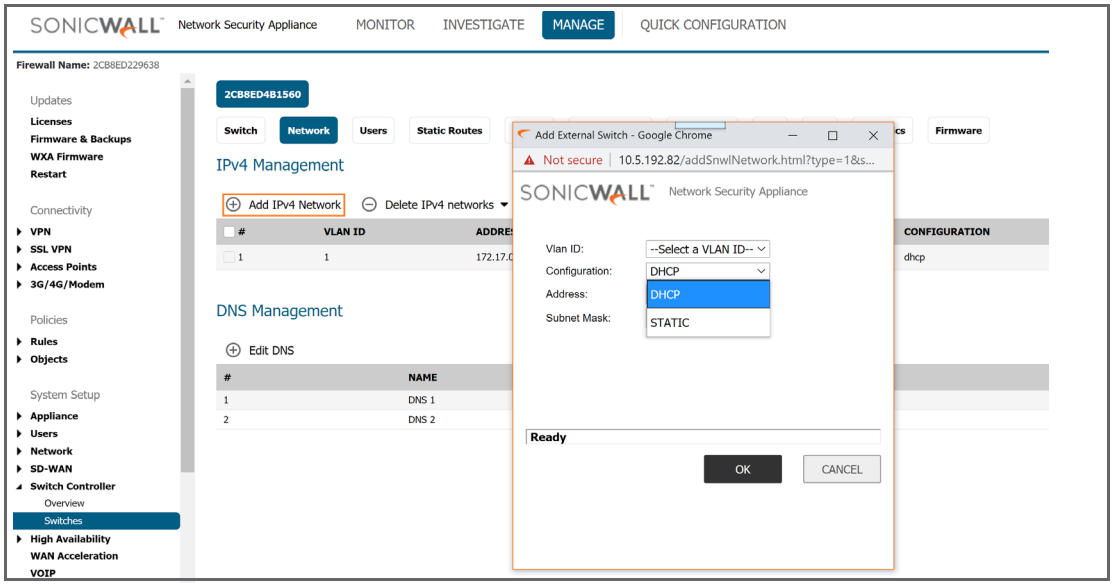
IMPORTANT: To change the Reserved VLAN range on the firewall, do so before adding the SonicWall Switch. If the Reserved VLAN range changes after connecting the Switch, then the Switch must be removed and re-added.

Adding a VLAN Interface:

Add a VLAN by adding a virtual interface under the uplink to the firewall.

1 Navigate to **Switch Controller > Switches**.

2 Click on **Add IPv4 Network**.



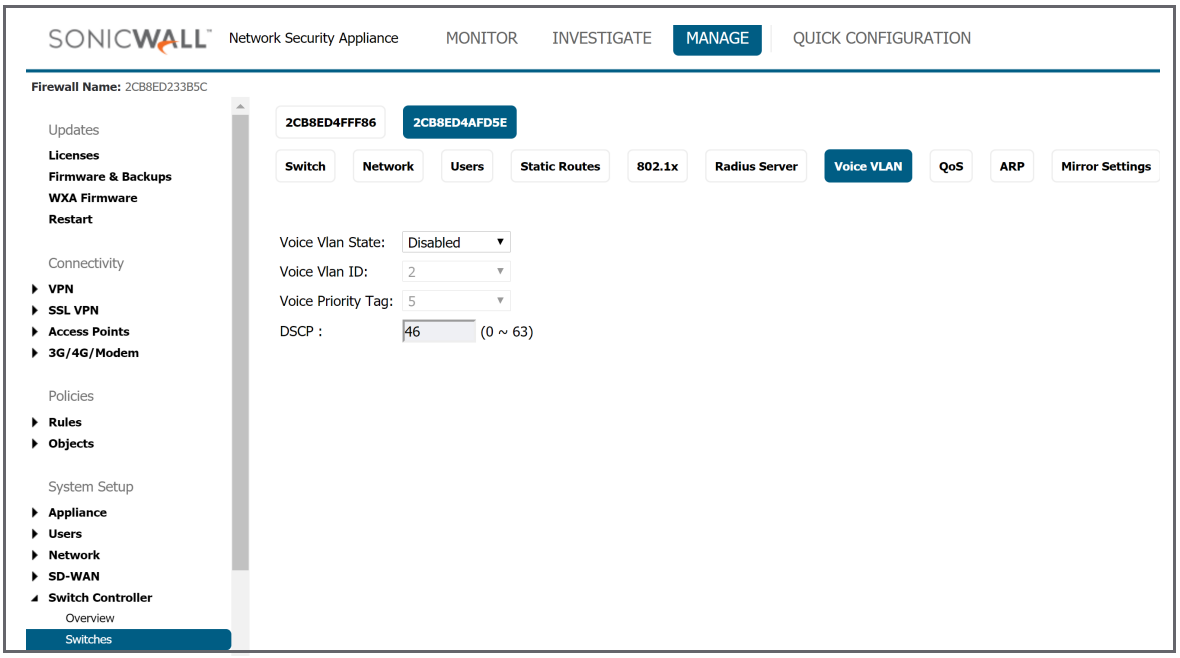
3 Define **Vlan ID**, **Address**, **Subnet Mask** and choose address assignment method: **Static** or **DHCP**.

4 Click on **OK**.

Configuring Voice VLAN:

NOTE: Voice VLANs can be enabled/disabled per port in the **Switch Controller > Overview > Physical View** display.

1 To configure a voice VLANs navigate to **Switch Controller > Switches** and then click on **Voice VLAN.**



2 Set up a voice VLAN by moving the state from **Disabled** to **Auto** and set the other parameters before clicking on **Accept** as it appears at the bottom of the display.

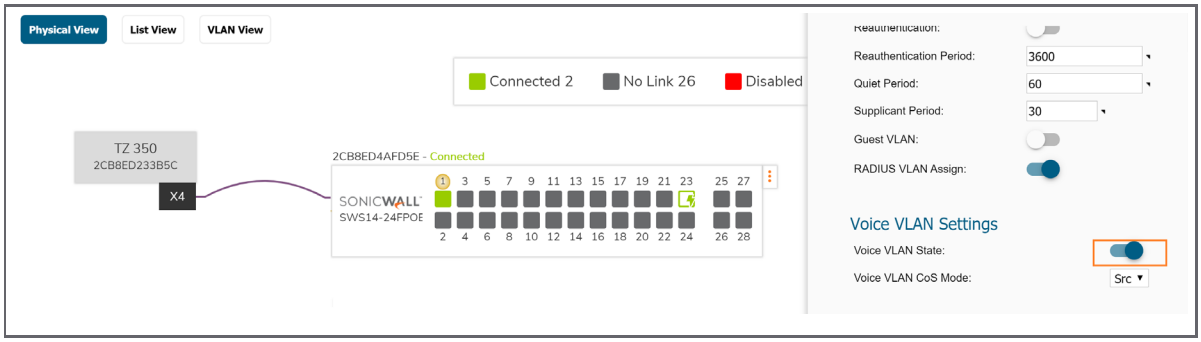
- **Voice Vlan ID —** identifies LAN.
- **Voice Priority Tag —** determines priority among active voice streams.
- **Differentiated Service Code Point —** defines QoS.

Use the Voice VLAN Settings to enable Voice traffic management and determine if Class of Service (CoS) queues will be defined for all ports or only those sourcing voice traffic. For more on CoS definition, see Setting Up QoS on page 656.

NOTE: The Switch remarks incoming voice VLAN traffic tags for voice priority and DSCP as defined by these settings.

To Enable/Disable Voice VLAN from the Physical View:

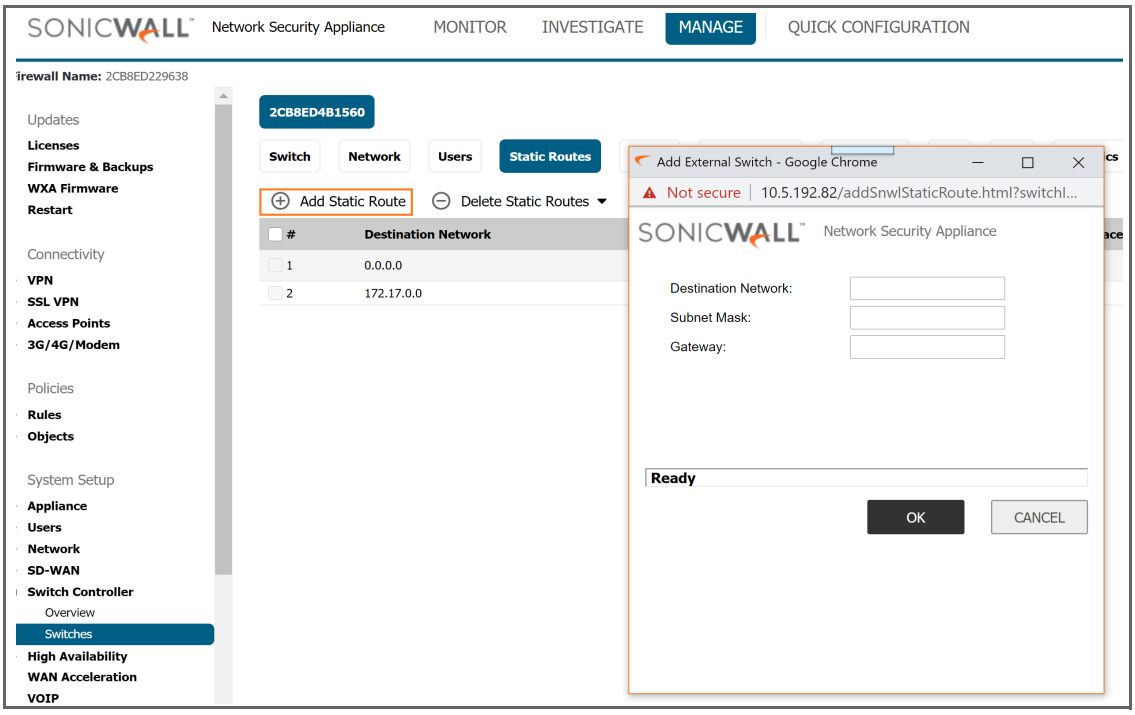
Simply go to **MANAGE > Switch Controller > Overview** and click on the port. When the sideband display appears scroll to Voice VLAN state as shown below.



Adding Static Routes

To add a static route to a Switch:

1 Navigate to **Switch Controller > Switches** then select **Static Routes** and click on **Add Static Route**.



2 Fill out the dialog box.

Destination IP address with '0' as the last octet: x.x.x.0.

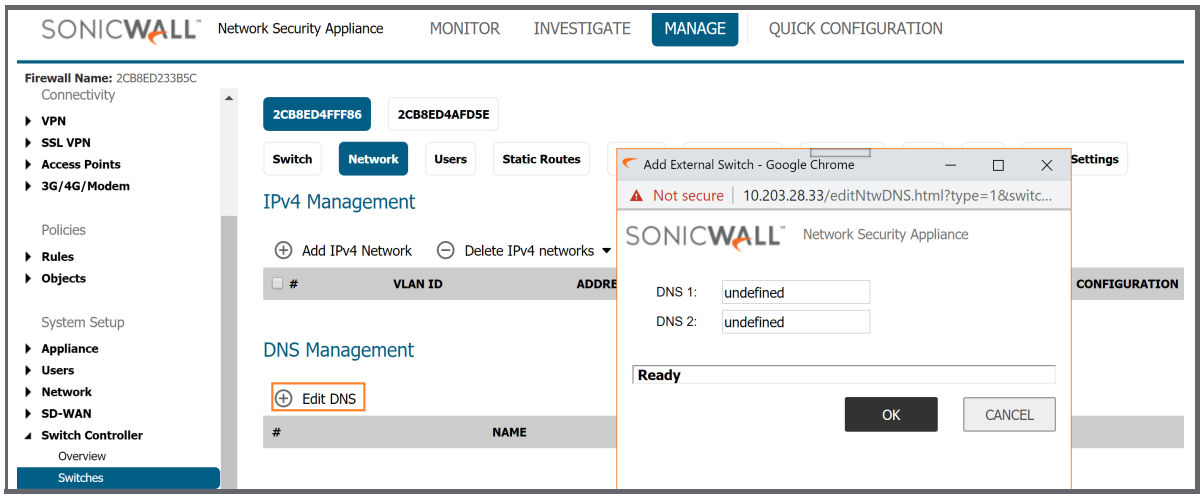
Subnet Mask for the destination.

Gateway: IP address gateway between Switch and destination.

3 Click on **OK**.

Editing DNS

To set DNS addresses go to **Switch Controller > Switches** and select **Network,** then click on **Edit DNS**.



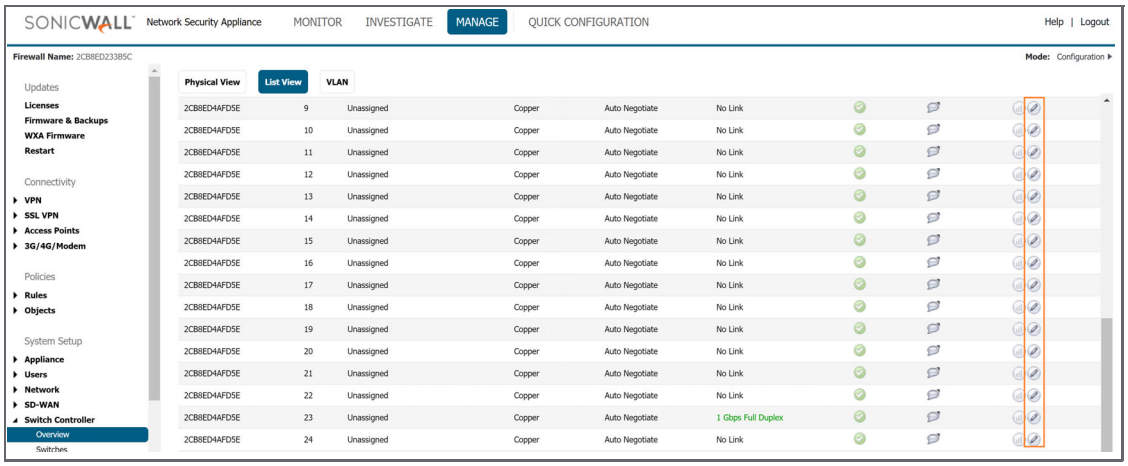
Setting Up the Ports

To configure specific ports:

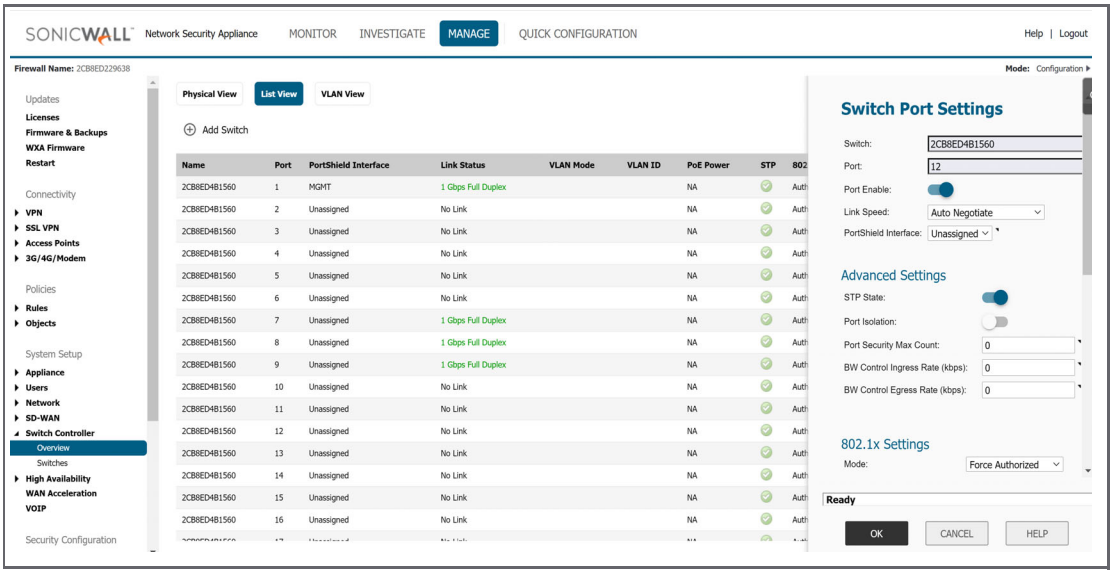
1 Go to **Switch Controller > Overview** and click on **List View**.

This can also be done from the **Physical View**.

2 When the list appears click on the edit button for the specific port.



3 The port setup dialog for the specific port will now appear to the right of the screen.

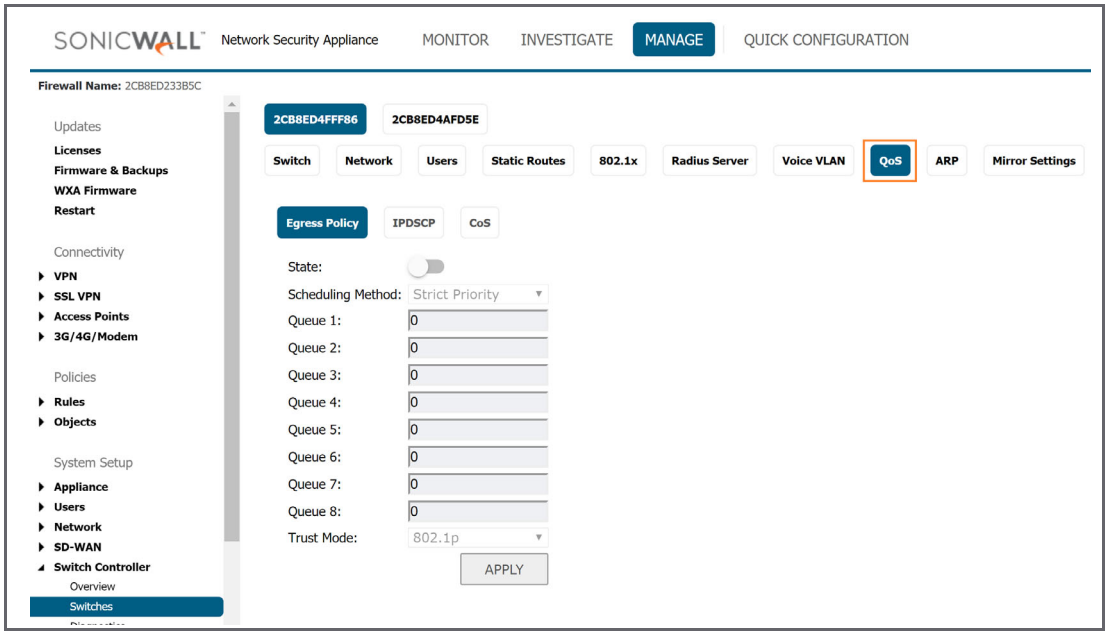


Setting Up QoS

Quality of Service (QoS) provides the ability to implement priority queuing within a network. QoS enables traffic to be prioritized, while minimizing excessive broadcast and multicast. Traffic such as voice and video streaming which requires a minimal delay can be assigned to a high priority queue, while other traffic can be assigned to a lower priority queue resulting in uninterrupted actions.

To set up QoS for a Switch:

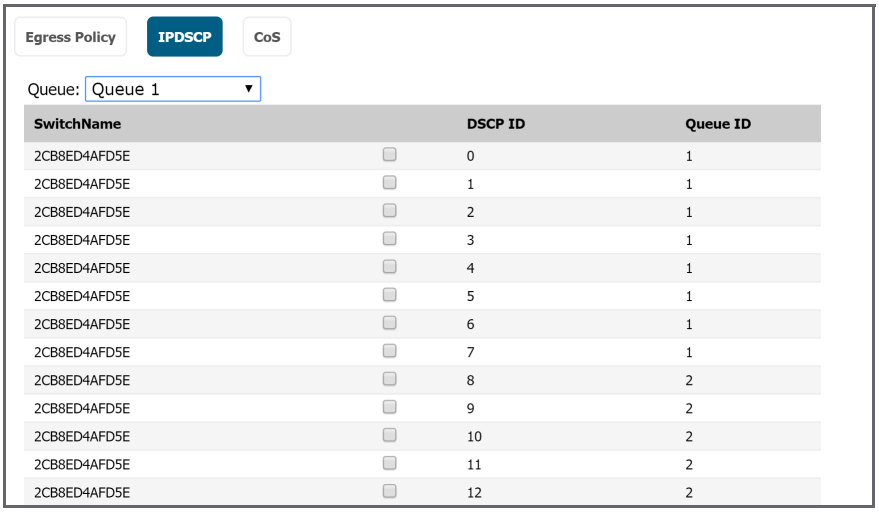
1 Navigate to **Switch Controller > Switches** and click on **QoS**.



2 Set Egress Policy.

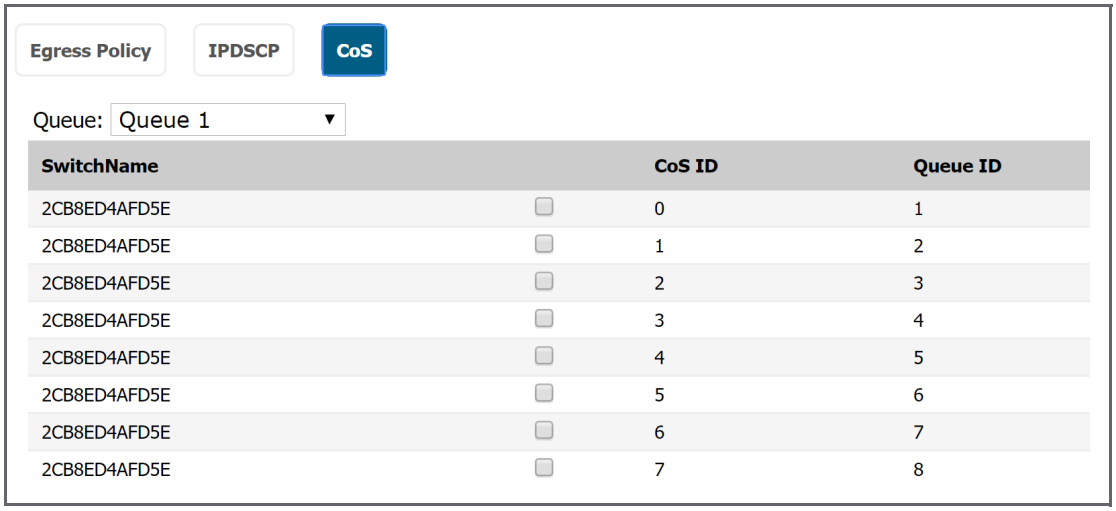
The first screen details Egress Policy which applies for all approaches to packet and traffic classification.In the preceding UI screen, the **State** slider determines whether QoS is enabled (to the right) or disabed (to the left). Scheduling method can be set as **Strict Priority** based on Queue number or as Weighted Round Robin (**WRR**). The classification of packets can be set as 802.1p or DSCP (Differentiated Sevices Code Point), or as both.

3 Select the **IPDSCP** screen to set DSCP codes to specific Queues.



4 To set class of service, click on **CoS**..

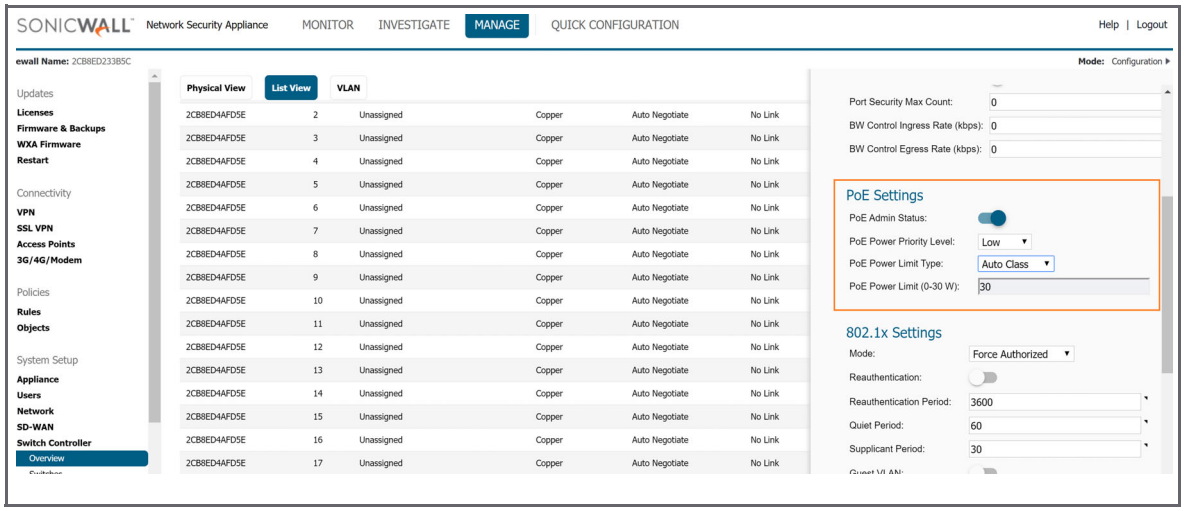
In the CoS (Class of Service) screen, the CoS priority tag values, where 0 is the lowest and 7 is the highest are related to eight traffic priority queues from 1 to 8, where one is the lowest priority and eight is the highest priority.



SwitchName		CoS ID	Queue ID
2CB8ED4AFD5E	☐	0	1
2CB8ED4AFD5E	☐	1	2
2CB8ED4AFD5E	☐	2	3
2CB8ED4AFD5E	☐	3	4
2CB8ED4AFD5E	☐	4	5
2CB8ED4AFD5E	☐	5	6
2CB8ED4AFD5E	☐	6	7
2CB8ED4AFD5E	☐	7	8

Setting Up PoE

To set up PoE limits per port, navigate to **Switch Controller > Overview** and click on **List View.** Select the click on the edit button for the port for PoE setup. Scroll down in the port configuration panel until the PoE settings appear.



The PoE+ Switches support Power over Ethernet (PoE) as defined by the IEEE 802.3af and 802.3at. The SWS12-8 PoE-enabled Switches support the -af standard and up to 15.4 Watts per port. The SWS12-10 and SWS14 series PoE-enabled Switches support the 30 Watts per port.

The Switches follow the standard PSE (Power Sourcing Equipment) pinout, whereby power is sent out over pins 1, 2, 3 and 6.

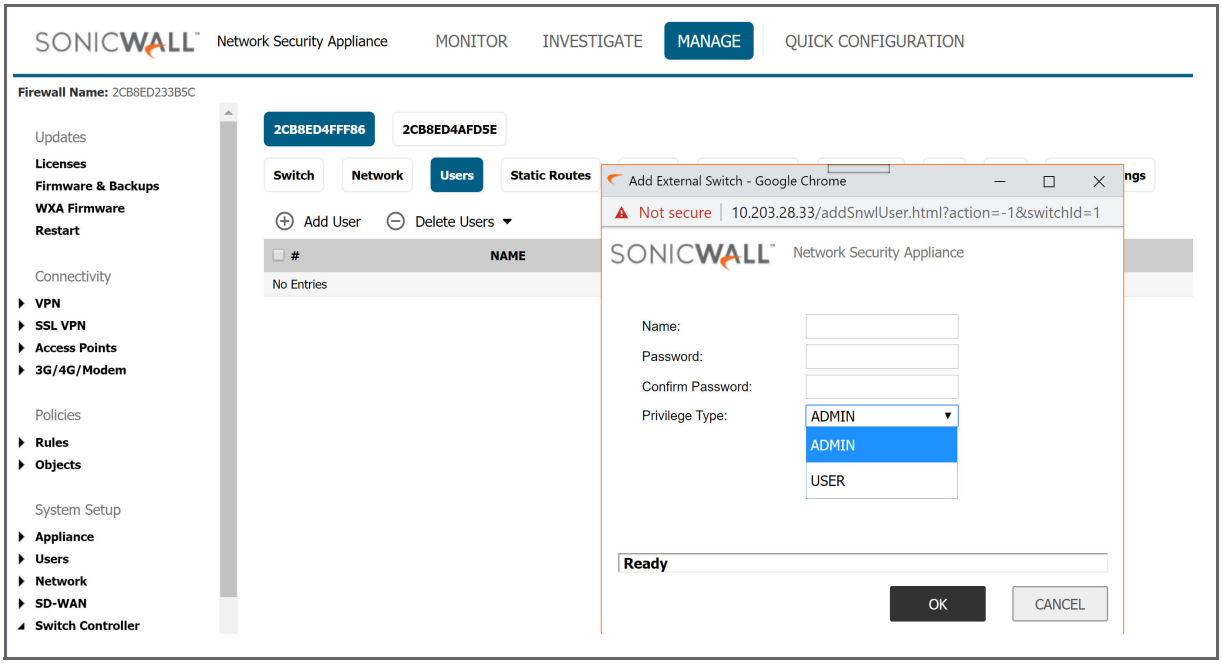
- **PoE Admin Status**
  - **Enabled -** Enables the Device Discovery protocol and provides power to the device using the PoE module. The Device Discovery Protocol lets the device discover powered devices attached to device interfaces and learns their classification.
  - **Disabled -** Disables the Device Discovery protocol and halts the power supply delivering power to the device using the PoE module.
- **PoE Priority**
  Select the port priority if the power supply is low. The field default is Low. For example, if the power supply is running at 99% usage, and port 1 is prioritized as high, but port 6 is prioritized as low, port 1 is prioritized to receive power and port 6 may be denied power. The possible field values are: 4.
  - **Low** – Sets the PoE priority level as low.
  - **Medium** – Sets the PoE priority level as medium.
  - **High** – Sets the PoE priority level as high.
  - **Critical** – Sets the PoE priority level as critical.
- **PoE Power Limit Type**
  - **Auto Class** - 15.4 or 30 W per port.
  - **User Defined** - Sets the maximum amount of power that can be delivered by a port.

    NOTE: The User Power Limit can only be implemented when the Auto Class value is set to User-Defined.

Setting Up Users

Users with different access levels, **admin** and **user**, can be defined by navigating to **Switch Controller > Switches** and clicking on **Users**.

Users are limited to Non-Configuration Mode.



Setting Up 802.1x Authentication

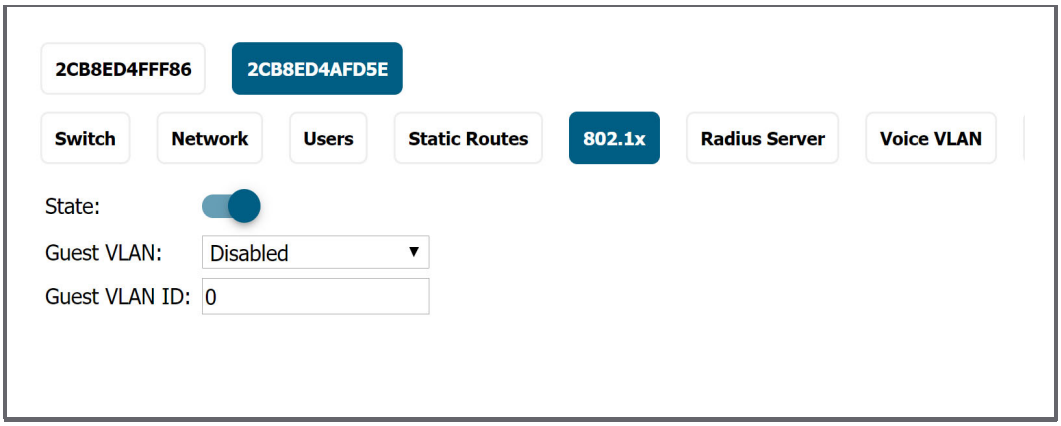
The IEEE-802.1X authentication provides a security standard for network access control with RADIUS servers and holds a network port disconnected until authentication is completed. With 802.1X authentication, the supplicant provides credentials, such as user name, password, or digital certificate to the authenticator, and the authenticator forwards the credentials to the authentication server for verification. If the authentication server determines the credentials are valid, the supplicant (client device) is allowed to access resources located on the protected side of the network. The Switch uses 802.1X to enable or disable port access control, to enable or disable the Guest VLAN, and to enable or disable the forwarding EAPOL (Extensible Authentication Protocol over LANs) frames.

To enable 802.1 Authentication

1 Go **MANAGE > Switch Controller > Switches** and click on 802.1 x.

2 Set the **State** slider to the right to enable authentication.
Other settings are:

- **Guest VLAN** — Select whether Guest VLAN is enabled or disabled on the Switch. The Default is disabled.
- **Guest VLAN ID** — Select the Guest VLAN from the list for currently defined VLANs.



To enable RADUIS server

1 In **MANAGE** > **Switch Controller > Switches** click on **Radius Server.** When the server list appears click on **Add Radius Server.**

2 To enable the Radius server, set the **Authorized Port** as **1812.**

Daisy-Chaining Switches

Switches can be setup with firewalls in standalone or daisy-chained configurations.

- **Standalone mode** — Up to four Switches can interface to a single firewall over separate ports.
- **Daisy Chain mode** — Up to four Switches can be supported in three configurations.
    a 1 Switch in standalone mode and three Switches connected to it in daisy chain mode.
    b 2 Switches in standalone mode and with 1 Switch connected to each in daisy chain mode.
    c 3 Switches in standalone mode and 1 Switch connected to any of the Switches in daisy chain mode.

NOTE: Switches may be added into daisy-chained configurations manually or by using Zero-Touch.

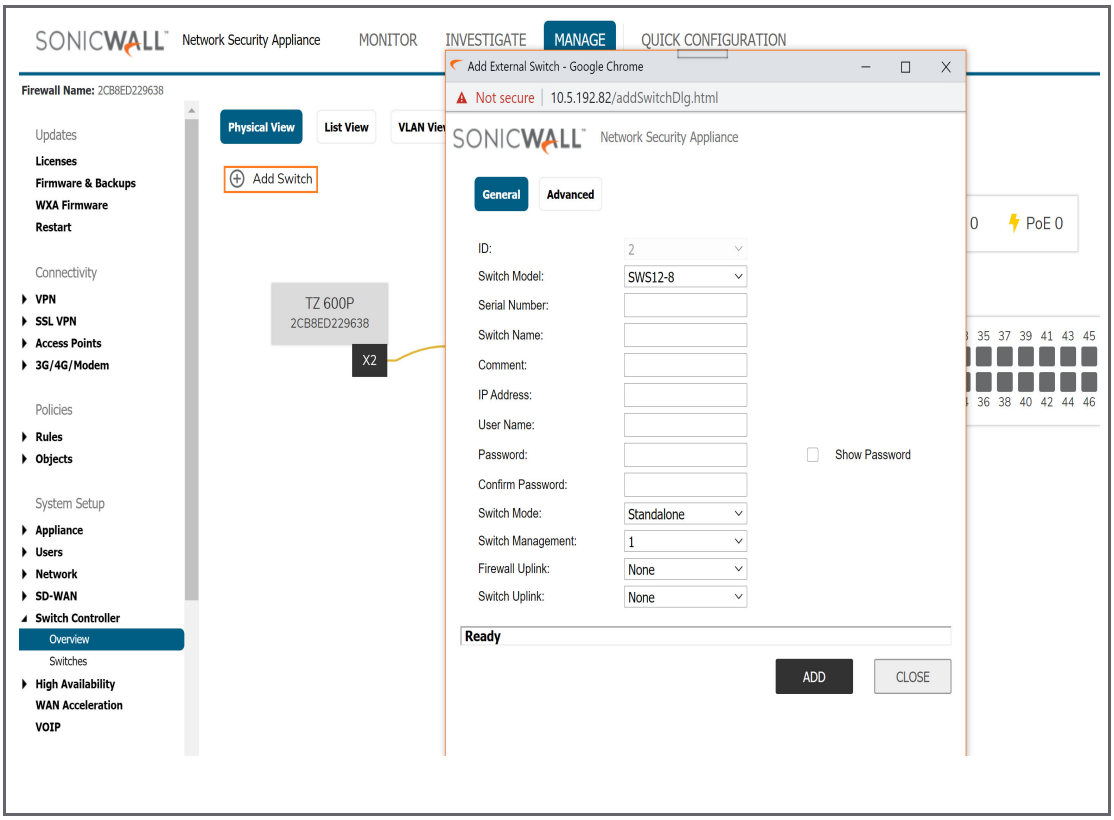
NOTE: Adding un-configured connections between the firewall and parent Switch will bring down the link between the parent Switch and a child Switch. To avoid this, configure additional links between the firewall and parent Switch before making the physical connection.

After connecting the child Switch to the parent Switch, the Switch will be visible in the **Switch Controller > Overview** page. Simply click the Authorize option and the Switch will be added in daisy chain manner.

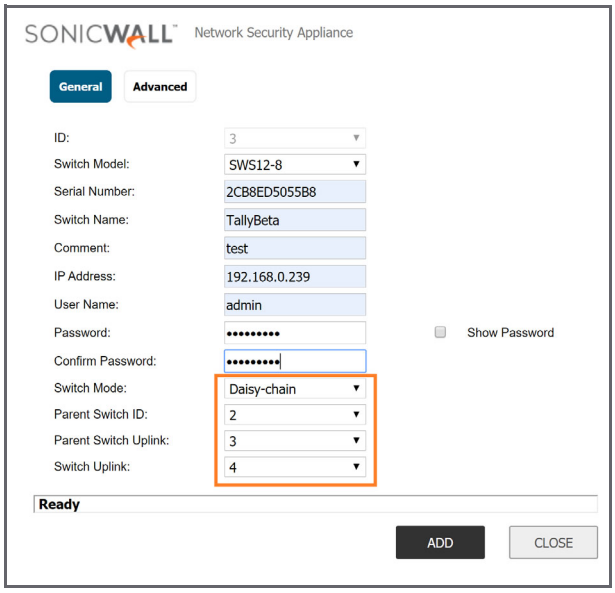
To add a Switch in daisy chain mode:

1 Select a Switch in standalone configuration to daisy-chain the additional Switch to. Then determine which ports to use to connect the additional Switch.

2 Go to **MANAGE | Switch Controller > Overview** and click on **Add Switch**.



3 When the **Add Switch** dialog box appears, make the entries outlined below.



- **IP Address —** This is an address within the leasehold of the DHCP server for Parent Switch. To identify this address range, go to **MANAGE | Network > DHCP Server.**
- **Switch Mode —** Select **Daisy-chain**.
- **Parent Switch ID —** For the ID of the parent Switch, refer to screen shot in Step 2. It is the second column in the row for this Switch.
- **Parent Switch Uplink** — Interface on parent Switch which is connected to the child Switch.
- **Switch Uplink —** This is the port through which the daisy-chained Switch connects to the Parent Switch.

4 When complete with the dialog box click on **ADD**.

NOTE: Define the first Switch connected to the firewall as **Standalone.** Setup the Switch connected to that Switch as **Daisy-chain**.

5 Navigate to **Switch Controller > Overview** and **click** on **Physical View**. The new Switch will appear graphically with the ports linking the Switch and the firewall indicated.

Connecting Access Points

With the firewall user interface, administrators may manage SonicWave access points connected to Switches.

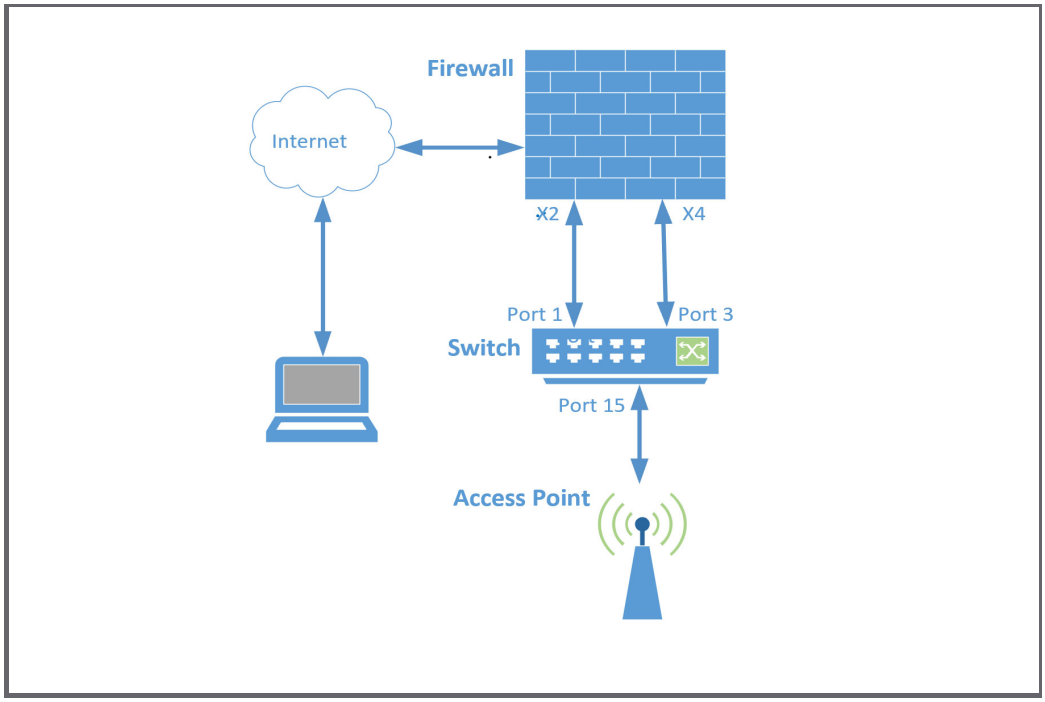
Adding access points to a Switch involves three steps beyond making the physical connection.

- Configure the network interface to the Switch supporting the access point to support the WLAN.
- Configure the WLAN zone for trust and security services.
- Configure the SonicWave access point entry for the desired radio frequency, mode, and authentication type.

The following graphic exemplifies a ***firewall — Switch — access point*** configuration.

To manage an access point through a Switch:

(this procedure refers to the following diagram)



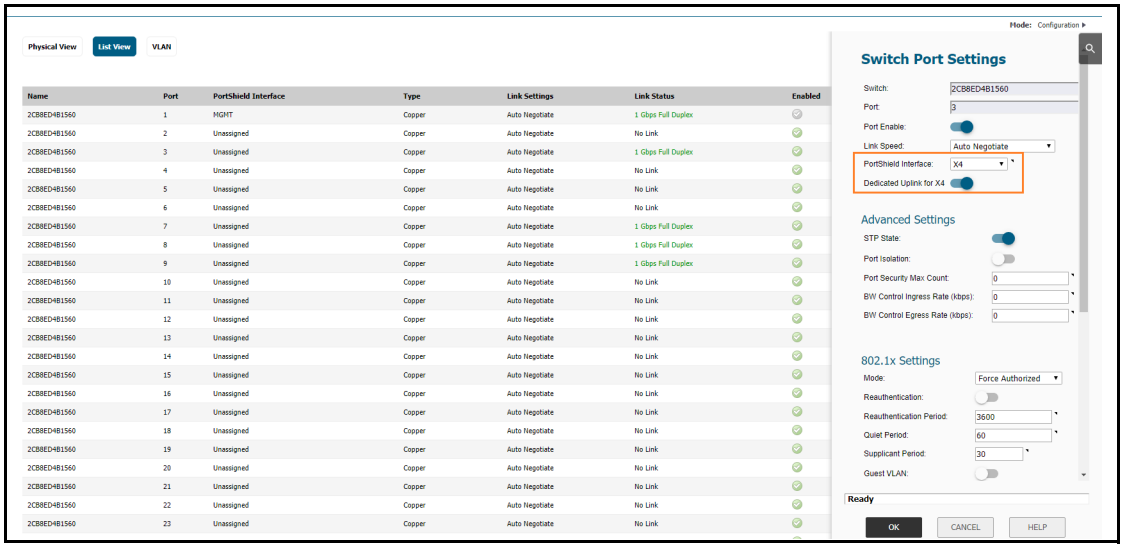
1 Connect Port-1 of Switch to X2 interface and enable auto discovery on X2 interface.
For details refer to Configuring from the Firewall on page 640.

2 Add the Switch.

3 Configure X4 in WLAN zone with VLANs.
Refer to Connect the SonicWave access point to port 15 on the Switch. on page 664 and To Configure the WLAN Zone: on page 665.

4 Connect Switch Port 3 to X4 interface.

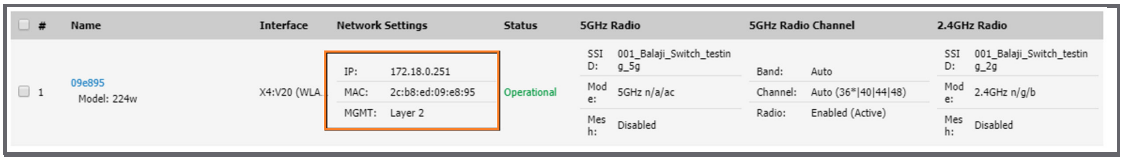
5 In the firewall GUI, navigate to **Switch Controller > Overview**, click **List View**. Click on the pencil icon to configure port 3. To create a dedicated uplink, set the **Portshield Interface** to X4, and move the **Dedicated Uplink** Switch to the right.



6 Connect the SonicWave access point to port 15 on the Switch.

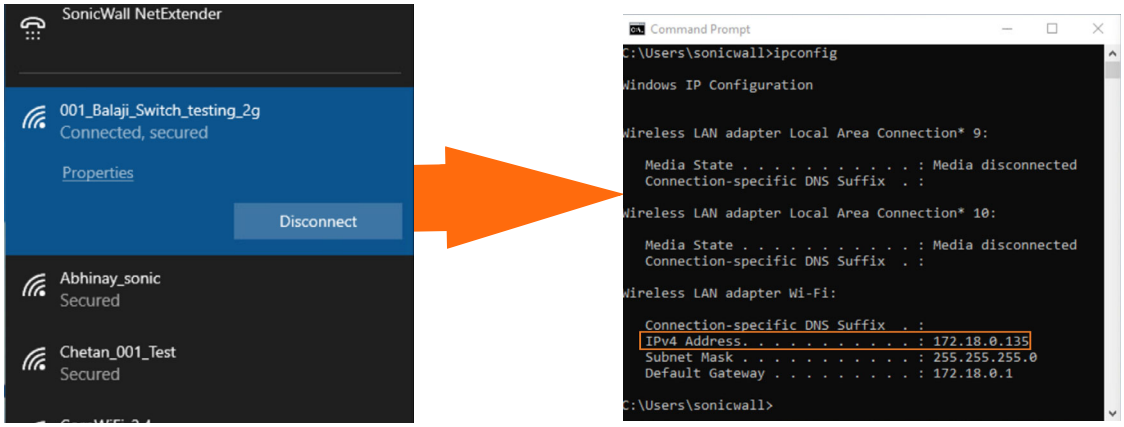
7 Go to the Switch Port Settings for port 15 (as in step 5) and set the Portshield Interface to X4.
You may instead set the port to any VLAN in the X4 interface which is in the WLAN zone, see Adding a VLAN on page 650.

8 After connecting and Port-Shielding the interface where SonicWave connected to firewall interface, verify that the Sonicwave gets an IP address from the configured network.
To do this, in the firewall GUI, go to **Access Points > Base Settings** and select **SonicWave Object.**



For details on configuring the SonicWave object, see Configuring the SonicWave Settings: on page 666.

9 Connect a WiFi client and check that it gets an IP address from in the X4 Portshield leasehold.

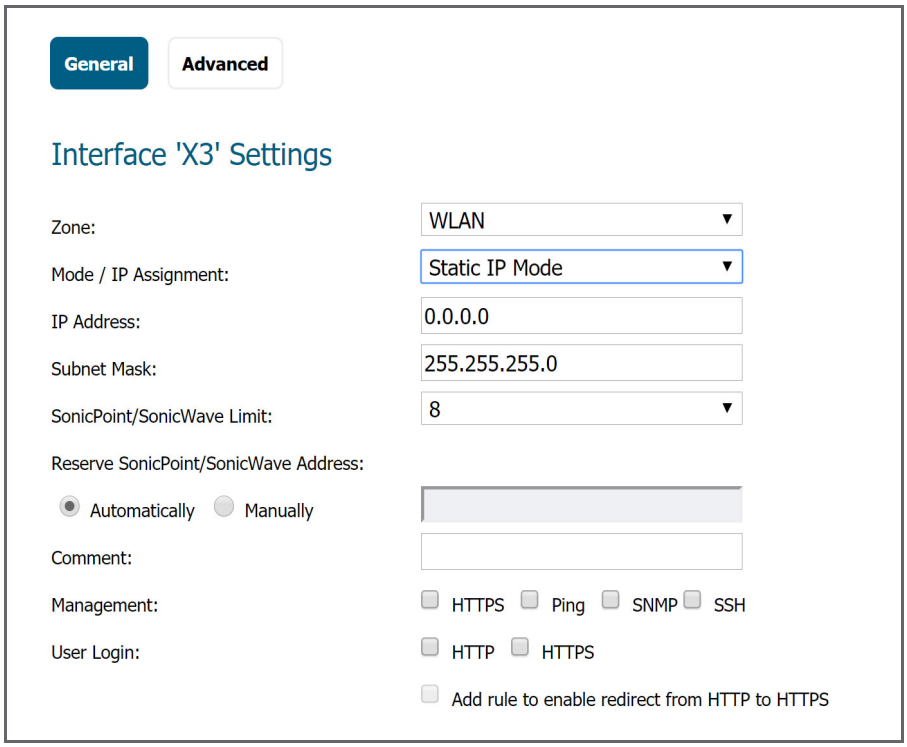


Configure the network interface to the Switch supporting the access point to support to WLAN.

1 Login to the firewall as an adminstrator and go to **MANAGE | Network > Interfaces** page and click on the configure icon for the inteface the Switch is supported on.

2 Select **WLAN** for the **Zone** type.

3 Select the **Static IP Mode** for the **Mode/IP Assignment**.



4 In the **IP Address** field, type in any private IP address that does not interfere with the IP address range of any other interfaces on the appliance.

5 Enter a **Subnet Mask**. The default is 255.255.255.0.

6 Use the default settings or select appropriate settings for the other fields and then click **OK**.

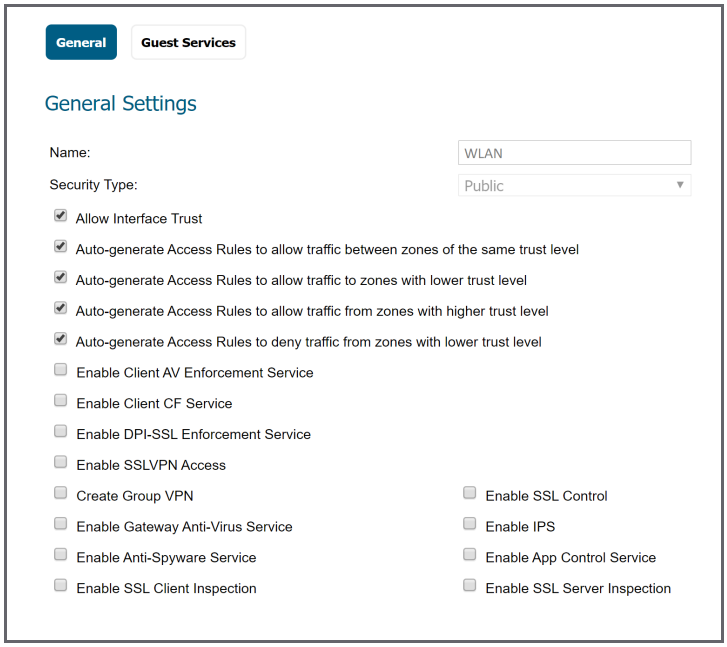
CAUTION: Allowing Management and User Login to the appliance from a wireless zone can pose a security threat if you or your users have not set strong passwords.

To Configure the WLAN Zone:

1 In the **MANAGE** view on the **System Setup | Network > Zones** page, click the **Edit** icon in the **Configure** column of the **WLAN** row.

2 On the **General** page, under **General Settings**, select the **Allow Interface Trust** option to automate the creation of Access Rules to allow traffic to flow between the interfaces within the zone, regardless of the interfaces to which the zone is applied.

For example, if the WLAN zone has both the **X1** and **X2** interfaces assigned to it, selecting **Allow Interface Trust** creates the necessary access rules to allow hosts on these interfaces to communicate with each other.



3 Select the checkboxes to enable security services on this zone. Minimally, you would select **Enable Gateway Anti-Virus Service**, **Enable IPS**, and **Enable Anti-Spyware Service**, if your wireless clients are all running **Spyware Service**. If your wireless clients are all running **SonicWall Client Anti-Virus**, select Enable Client .

4 In the **Guest Services** page, optionally configure guest Internet access. For information about Guest Services, see *SonicOS 6.5.4 Connectivity Administation* the documentation.

5 In **Wireless** under **SonicPoint/SonicWave Settings**, select **Only allow traffic generated by a SonicPoint/SonicWave** to allow only traffic from SonicPoints/SonicWaves to enter the WLAN zone interfaces, providing maximum security.

6 When finished, click **OK**.

Configuring the SonicWave Settings:

When a SonicWave AP is initially connected to an interface, the firewall uses a default provisioning profile to create a SonicWave AP entry. It can take up to five minutes for the entry to be created.

You can modify the SonicWave AP entry to configure the access point name, radio frequency mode, authentication type, and other settings specific to your SonicWave AP.

TIP: For deployments of multiple SonicWaves that need the same provisioning settings, you can create a custom provisioning profile in the upper section of **Access Points > Base Settings** page in the **MANAGE** view. In **System Setup | Network > Zones** page, you can edit the WLAN zone and specify this profile on the **Wireless** page. Any SonicWaves connecting to an interface in the WLAN zone can then be provisioned with the assigned profile.

You might want to use the new **Floor Plan View** and **Topology View** features as well. See the *SonicOS 6.5.4 Connectivity Administration* documentation for more information.

To modify the SonicWave AP entry in SonicOS:

1 In the **MANAGE** view, navigate to **Access Points > Base Settings**.

2 In the **SonicPoint/SonicWave Objects** table, click the **Configure** icon in the row for the SonicWave AP entry you wish to modify.

General page settings:

1 On the **General** page, select **Enable SonicPoint**.

2 In the **Name** field, optionally type in a new name for this SonicWave AP. The existing name is assigned by the provisioning profile based on the name prefix in the profile with a unique number appended.

This is the access point name that appears in clients' lists of available wireless connections.

3 Verify the **Country Code** for the area of operation.

4 Configure the remaining options as necessary. For more information, see the *SonicOS Connectivity Administration* documentation.

Radio 0 Basic / Radio1 Basic Settings:

1 Click **Radio 0 Basic**, or **Radio 1 Basic**.

The configuration is very similar for both Radio 0 Basic and Radio 1 Basic. The main differences are the radio frequencies:

Radio	Frequency	Default Mode
Radio 0	5 GHz	5GHz 802.11ac/n/a Mixed
Radio 1	2.4 GHz	2.4GHz 802.11n/g/b Mixed

2 Select **Enable Radio**.

3 Select a **Mode** or use the default.

4 Under Wireless Security, select the **Authentication Type** for your wireless network. SonicWall recommends using **WPA2** as the authentication type if all client devices support it.

TIP: *PSK* uses a personal passphrase for authentication, *EAP* uses an Enterprise RADIUS server.

5 Select the **Cipher Type**. When using WPA and WPA2, SonicWall recommends **AES** for maximum security.

NOTE: Older client devices might not support AES.

6 Fill in the fields specific to the authentication type that you selected. The remaining fields change depending on the selected authentication type.

7 Optionally, under **ACL Enforcement**, select **Enable MAC Filter List** to enforce Access Control by allowing or denying traffic from specific devices. Select a MAC address object group from the **Allow List** or **Deny List** to automatically allow or deny traffic to and from all devices with MAC addresses in the group. The **Deny List** is enforced before the **Allow List**.

Virtual Access Point Encryption Settings:

NOTE: This section displays only if a VAP was selected from the Radio 0 Basic/1 Virtual AP Group drop-down menu in the **Virtual Access Point Settings** section of the **General** page.

The **Virtual Access Point Encryption Settings** section of both **Radio 0 Basic** and **Radio 1 Basic** are the same for the **802.11n Radio**.

Radio 0 Advanced / Radio1 Advanced Settings:

1 Click **Radio 0 Advanced** or **Radio 1 Advanced**.

The configuration is very similar for Radio 0 Advanced and Radio 1 Advanced. For most advanced options, the default settings give optimum performance. For a full description of the fields on this page, see the *SonicOS Connectivity Administration* documentation.

2 Optionally select the **Hide SSID in Beacon** checkbox.

The *SSID* refers to the access point name that appears in clients' lists of available wireless connections.

Hiding the SSID provides additional security because it requires that you know the access point name before connecting.

3 When finished configuring all options, click **OK**.

Sensor page

On the Sensor page, enable or disable **Wireless Intrusion Detection and Prevention (WIDP) mode**.

NOTE: If this option is selected, Access Point or Virtual Access Point(s) functionality is disabled automatically.

1 Select **Enable WIDF sensor** to have the SonicWave operate as a dedicated WIDP sensor. This option is not selected by default.

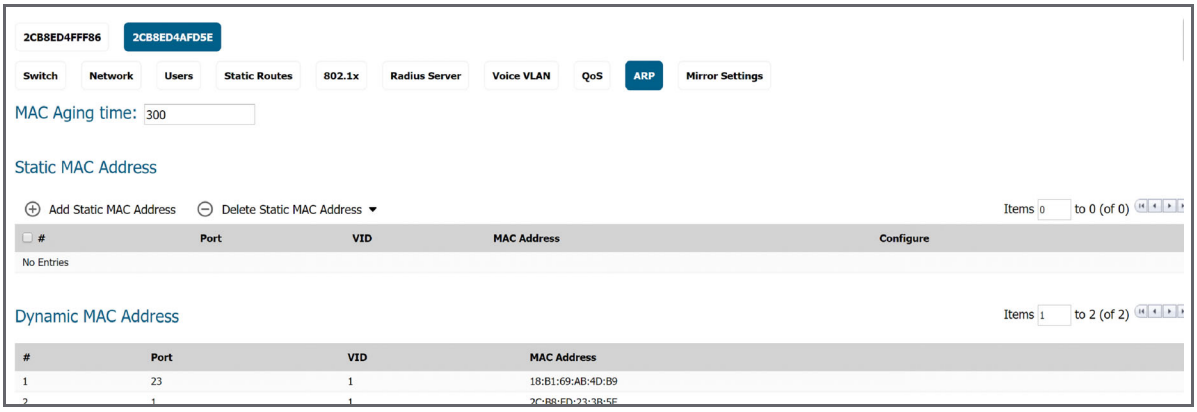
2 From the drop-down menu, select the schedule for when the SonicWave operates as a WIDP sensor or select **Create new schedule...** to specify a different time; default is **Always on**.

Modifying the MAC Address Table

The MAC address table links the MAC destination address on incoming Ethernet frames with the port closest to the destination based on learning from the transit of earlier frames. This feature allows:

- Defining MAC aging time
- Setting Static MAC table entries
- Checking Dynamic MAC entry learning

Navigate to **MANAGE > Switch Controller > Switches** and then click on **ARP**.

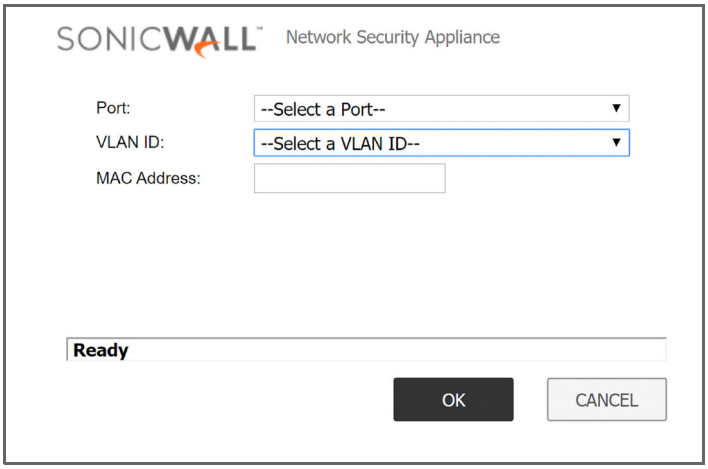


To set MAC Aging Time:

The **MAC Aging time** specifies the **time** before an entry ages and is discarded from the **MAC** address table. The range is from 0 to 1000000; The default value is 300 seconds. Entering the value 0 disables **MAC aging**. This age specification applies to all VLANs.

To add static MAC Addresses:

1 Click on **Add Static MAC Addresses** and the following dialog box will appear.



2 Select the **Port** and **VLAN ID** along with the destination MAC address and click on **OK.**

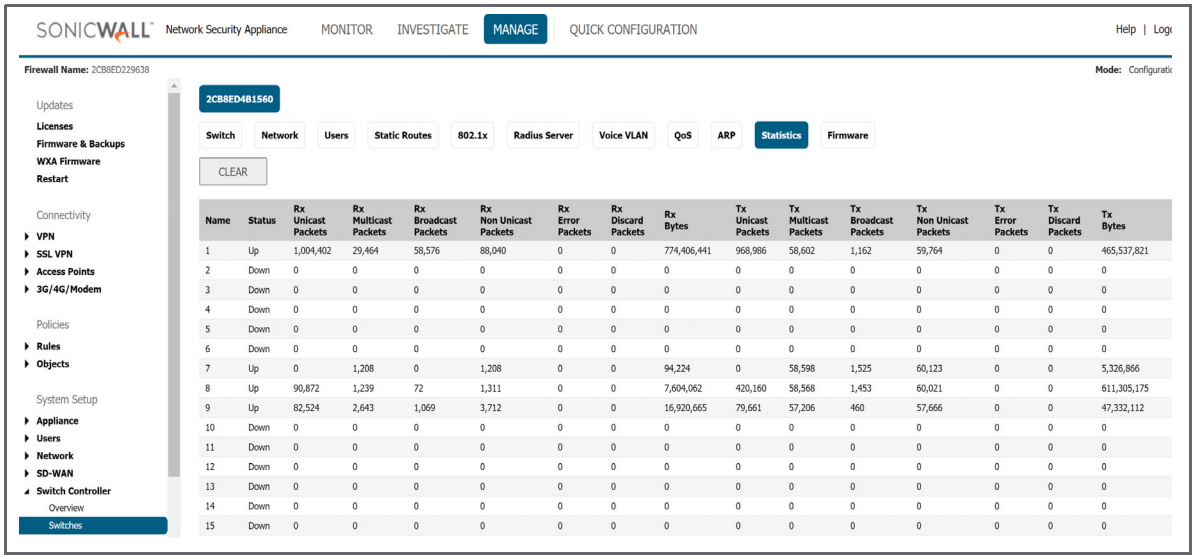
To Check Dynamic MAC Address Learning:

The dynamic MAC address table lists currently learned MAC addresses and accompanying Port and VLAN IDs. The defined **MAC Aging time** deterimines how current this information is. This table provides details on the LAN supported by the Switch.

Checking Port Statistics

The statistics table for a Switch can also be reached through **Switch Controller > Switches > Statistics**.

This table presents details on port-by-port performance.



Name	Status	Rx Unicast Packets	Rx Multicast Packets	Rx Broadcast Packets	Rx Non Unicast Packets	Rx Error Packets	Rx Discard Packets	Rx Bytes	Tx Unicast Packets	Tx Multicast Packets	Tx Broadcast Packets	Tx Non Unicast Packets	Tx Error Packets	Tx Discard Packets	Tx Bytes
1	Up	1,004,402	29,464	58,576	88,040	0	0	774,406,441	968,986	58,602	1,162	59,764	0	0	465,537,821
2	Down	0	0	0	0	0	0	0	0	0	0	0	0	0	0
3	Down	0	0	0	0	0	0	0	0	0	0	0	0	0	0
4	Down	0	0	0	0	0	0	0	0	0	0	0	0	0	0
5	Down	0	0	0	0	0	0	0	0	0	0	0	0	0	0
6	Down	0	0	0	0	0	0	0	0	0	0	0	0	0	0
7	Up	0	1,208	0	1,208	0	0	94,224	0	58,598	1,525	60,123	0	0	5,326,866
8	Up	90,872	1,239	72	1,311	0	0	7,604,062	420,160	58,568	1,453	60,021	0	0	611,305,175
9	Up	82,524	2,643	1,069	3,712	0	0	16,920,665	79,661	57,206	460	57,666	0	0	47,332,112
10	Down	0	0	0	0	0	0	0	0	0	0	0	0	0	0
11	Down	0	0	0	0	0	0	0	0	0	0	0	0	0	0
12	Down	0	0	0	0	0	0	0	0	0	0	0	0	0	0
13	Down	0	0	0	0	0	0	0	0	0	0	0	0	0	0
14	Down	0	0	0	0	0	0	0	0	0	0	0	0	0	0
15	Down	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Setting Spanning Tree Protocol

Loops occur when alternate routes exist between hosts. Loops in an extended network can cause the Switch to forward traffic indefinitely, resulting in increased traffic and reducing network efficiency. Once the STP is enabled and configured, primary links are established and duplicated links are blocked automatically. The reactivation of the blocked links is also accomplished automatically. STP provides a tree topology and other Spanning tree versions supported include STP, Multiple Spanning Tree Protocol (MSTP), and Rapid Spanning Tree Protocol (RSTP). Please note that only one spanning tree can be active on the Switch at a time. The default setting is: MSTP.

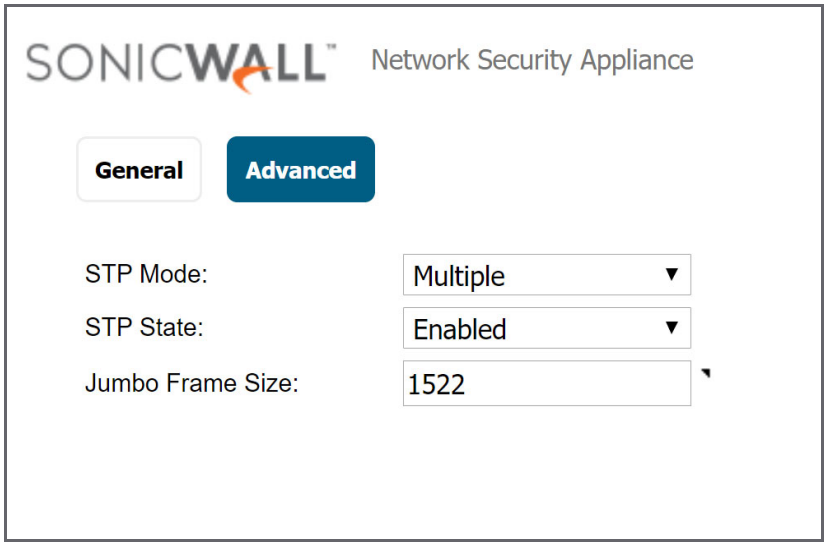
Multiple Spanning Tree Protocol (MSTP) defined in IEEE 802.1s, enables multiple VLANs to be mapped to reduce the number of spanning-tree instances needed to support a large number of VLANs. If there is only one VLAN in the network, a single STP works appropriately.

If the network contains more than one VLAN however, the logical network configured by a single STP would work, but it becomes more efficent to use the alternate paths available by using an alternate spanning tree for different VLANs or groups of VLANs. MSTP (which is based on RSTP for fast convergence) is designed to support independent spanning trees based on VLAN groups. MSTP provides multiple forwarding paths for data traffic and enables load balancing.

STP and RSTP prevent loops from forming by ensuring that only one path exists between the end nodes in your network. RSTP is designed as a general replacement for the slower, legacy STP. RSTP is also incorporated into MSTP. With STP, convergence can take up to a minute to complete in a larger network. This can result in the loss of communication between various parts of the network during the convergence process so STP can subsequently can lose data packets during transmission.

To set Spanning Tree Protocol Configuration:

1 Navigate to **MANAGE | Switch Controller > Overview** and then click on the 3 dot menu to at the right end of the Switch graphic. When the menu appears, click on **Edit**. Then click on **Advanced**. The dialog box will appear:



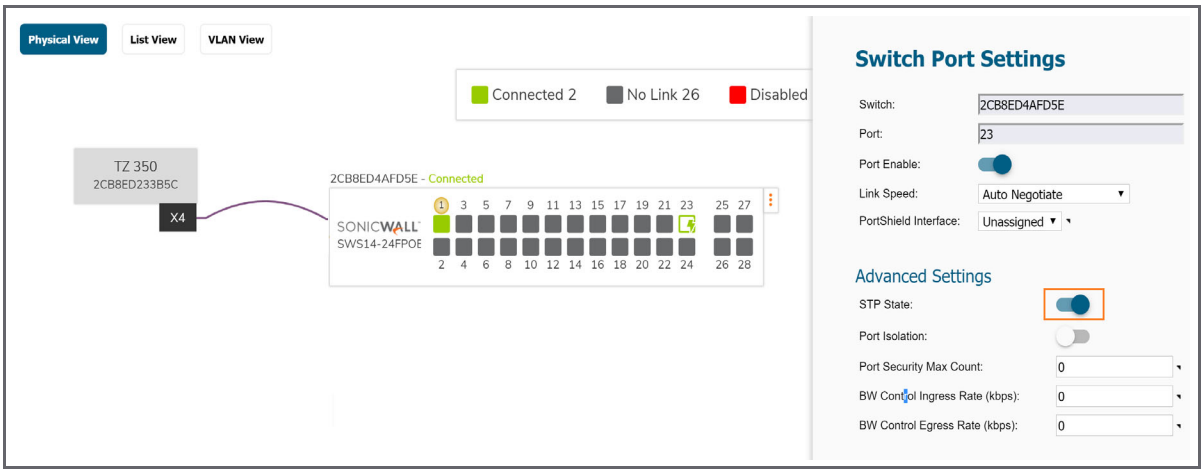
2 For the two STP entries:

- **STP Mode** — Sets the protocol version to Multiple (MSTP) or Rapid (RSTP).
- **STP State** — Enables or Disables spanning tree operation on the Switch.

3 Click on **OK**.

To Enable/Disable STP from Physical View:

Simply go to **MANAGE | Switch Controller > Overview** and click on the port. When the sideband display appears scroll to STP state as shown below.



Part 7

System Setup | High Availability

- About High Availability and Active/Active Clustering
- Configuring High Availability
- Fine Tuning High Availability
- Monitoring High Availability

46

About High Availability and Active/Active Clustering

NOTE: High Availability (HA) is supported on TZ series and above Security Appliances except HA is not supported on TZ series Security Appliances with wireless enabled. Stateful HA and Active/Active DPI are supported on TZ500 Series and above Security Appliances. See Active/Standby and Active/Active DPI Prerequisites on page 683. Active/Active Clustering is supported on NSA 3600 and above Security Appliances. See Licensing Requirements for Active/Active Clustering on page 700.

NAT64 does not support High Availability.

Topics:

- High Availability on page 674
  - About High Availability on page 675
  - About Active/Standby HA on page 679
  - About Stateful Synchronization on page 681
  - About Active/Active DPI HA on page 683
  - Active/Standby and Active/Active DPI Prerequisites on page 683
  - Maintenance on page 686
- Active/Active Clustering on page 688
  - About Active/Active Clustering on page 688

High Availability

This section provides conceptual information about High Availability (HA) in SonicOS and describes how to connect the Security Appliances for HA.

Topics:

- About High Availability on page 675
- About Active/Standby HA on page 679
- About Stateful Synchronization on page 681
- About Active/Active DPI HA on page 683
- Active/Standby and Active/Active DPI Prerequisites on page 683
- Physically Connecting Your Security Appliances on page 685

About High Availability

Topics:

- What Is High Availability? on page 675
- High Availability Modes on page 676
- High Availability Encryption on page 677
- Crash Detection on page 677
- Virtual MAC Address on page 678
- Dynamic WAN Interfaces with PPPoE HA on page 678
- Stateful Synchronization with DHCP on page 678
- About HA Monitoring on page 679

What Is High Availability?

High Availability (HA) is a redundancy design that allows two identical SonicWall Security Appliances running SonicOS to be configured to provide a reliable, continuous connection to the public Internet. One SonicWall SuperMassive is configured as the Primary unit, and an identical Security Appliance is configured as the Secondary unit. If the Primary Security Appliance fails, the Secondary Security Appliance takes over to secure a reliable connection between the protected network and the Internet. Two Security Appliances configured in this way are also known as a High Availability Pair (HA Pair).

High Availability provides a way to share SonicWall licenses between two SonicWall Security Appliances when one is acting as a high-availability system for the other. Both Security Appliances must be the same SonicWall model.

To use this feature, you must register the SonicWall Security Appliances on MySonicWall as Associated Products. For further information, see *SonicOS 6.5 Updates*.

High Availability Terminology

Active	The operative condition of a hardware unit. The Active identifier is a logical role that can be assumed by either a Primary or Secondary hardware unit.
Failover	The actual process in which the Standby unit assumes the Active role following a qualified failure of the Active unit. Qualification of failure is achieved by various configurable physical and logical monitoring facilities described in Configuring High Availability on page 703.
HA	High Availability: non-state, hardware failover capability.
IDV	Interface Disambiguation through VLAN.
PoE	Power over Ethernet is a technology that lets network cables carry electrical power.
PPP	Point-to-point protocol that provides a standard method for transporting multi-protocol diagrams over point-to-point links.
PPPoE	A method for transmitting PPP over ethernet.
PPPoE HA	HA PPPoE support function without State.
Preempt	Applies to a post-failover condition in which the Primary unit has failed, and the Secondary unit has assumed the Active role. Enabling Preempt causes the Primary unit to seize the Active role from the Secondary after the Primary has been restored to a verified operational state.

Primary	The principal hardware unit itself. The Primary identifier is a manual designation and is not subject to conditional changes. Under normal operating conditions, the Primary hardware unit operates in an Active role.
Secondary (Backup)	The subordinate hardware unit itself. The Secondary identifier is a relational designation and is assumed by a unit when paired with a Primary unit. Under normal operating conditions, the Secondary unit operates in a Standby mode. Upon failure of the Primary unit, the Secondary unit assumes the Active role.
SHF	State Hardware Failover, a SonicOS feature that allows existing network flows to remain active when the primary Security Appliance fails and the backup Security Appliance takes over.
Standby (Idle)	The passive condition of a hardware unit. The Standby identifier is a logical role that can be assumed by either a Primary or Secondary hardware unit. The Standby unit assumes the Active role upon a determinable failure of the Active unit.
STP	Spanning Tree Protocol.

High Availability Modes

High Availability has several operation modes, which can be selected on **High Availability > Base Setup**:

- **None**—Selecting **None** activates a standard high availability configuration and hardware failover functionality, with the option of enabling Stateful HA and Active/Active DPI.
- **Active/Standby**—Active/Standby mode provides basic high availability with the configuration of two identical Security Appliances as a High Availability Pair. The Active unit handles all traffic, while the Standby unit shares its configuration settings and can take over at any time to provide continuous network connectivity if the Active unit stops working.

  By default, Active/Standby mode is stateless, meaning that network connections and VPN tunnels must be re-established after a failover. To avoid this, Stateful Synchronization can be licensed and enabled with Active/Standby mode. In this Stateful HA mode, the dynamic state is continuously synchronized between the Active and Standby units. When the Active unit encounters a fault condition, stateful failover occurs as the Standby Security Appliance takes over the Active role with no interruptions to the existing network connections.

  NOTE: Stateful HA is:
  - Included on NSA 4600 and higher NSA platforms and SuperMassive Series platforms.
  - Supported on the NSA 2600 and NSA 3600 platforms with a SonicOS Expanded License or a High Availability License.
  - Supported on the TZ500 and higher TZ platforms with a SonicOS Expanded License or a High Availability (Stateful) Upgrade License.

  For licensing information, see *SonicOS 6.5 Updates*.

- **Active/Active DPI**—The Active/Active Deep Packet Inspection (DPI) mode can be used along with the Active/Standby mode. When Active/Active DPI mode is enabled, the processor intensive DPI services, such as Intrusion Prevention (IPS), Gateway Anti-Virus (GAV), and Anti-Spyware are processed on the standby Security Appliance, while other services, such as firewall, NAT, and other types of traffic are processed on the Active Security Appliance concurrently.

  NOTE: Active/Active DPI is:
  - Included on the SM 9000 series platforms.
  - Supported on the NSA 5600 and above platforms with a SonicOS Expanded License or a High Availability (Stateful) License.

  For licensing information, see *SonicOS 6.5 Updates*.

- **Active/Active Clustering**—In this mode, multiple Security Appliances are grouped together as cluster nodes, with multiple Active units processing traffic (as multiple gateways), doing DPI and sharing the

network load. Each cluster node consists of two units acting as a Stateful HA pair. Active/Active Clustering provides Stateful Failover support in addition to load-sharing. Optionally, each cluster node can also consist of a single unit, in which case Stateful Failover and Active/Active DPI are not available.

NOTE: Active/Active Clustering is:

- Included on the SM 9000 series platforms.
- Supported on NSA 3600 and above platforms only with the purchase of a SonicOS Expanded License.

For licensing information, see *SonicOS 6.5 Updates*.

- **Active/Active DPI Clustering**—This mode allows for the configuration of up to four HA cluster nodes for failover and load sharing, where the nodes load balance the application of DPI security services to network traffic. This mode can be enabled for additional performance gain, utilizing the standby units in each cluster node.

  NOTE: Active/Active DPI Clustering is:

  - Included on the SM 9000 series platforms.
  - Supported on NSA 3600 and above platforms only with the purchase of a SonicOS Expanded License.

  For licensing information, see *SonicOS 6.5 Updates*.

High Availability Encryption

High Availability encryption adds security to the communication between appliances in a HA pair. HA control messages between active and standby firewalls, such as heartbeats, configuration sync and HA state information, are encrypted to ensure security for inter-node communication.

This option is available in Active-Standby HA mode only and does not apply to messages exchanged for stateful synchronization even in Active-Standby mode. Discovery messages (find-peer and found-peer) are transmitted without encryption. After the discovery stage, however, all control messages are encrypted between the firewalls:

- Heartbeats
- Messages used for incremental config updates
- `prefSync` messages
- Various messages for sending HA commands between the firewall pair
- Firmware sync messages

Crash Detection

The HA feature has a thorough self-diagnostic mechanism for both the Active and Standby Security Appliances. The failover to the standby unit occurs when critical services are affected, physical (or logical) link failure is detected on monitored interfaces, or when the Security Appliance loses power.

The self-checking mechanism is managed by software diagnostics, which check the complete system integrity of the Security Appliance. The diagnostics check internal system status, system process status, and network connectivity. There is a weighting mechanism on both sides to decide which side has better connectivity to avoid potential failover looping.

Critical internal system processes such as NAT, VPN, and DHCP (among others) are checked in real time. The failing service is isolated as early as possible, and the failover mechanism repairs it automatically.

Virtual MAC Address

The Virtual MAC address allows the High Availability pair to share the same MAC address, which dramatically reduces convergence time following a failover. Convergence time is the amount of time it takes for the devices in a network to adapt their routing tables to the changes introduced by high availability.

Without Virtual MAC enabled, the Active and Standby Security Appliances each have their own MAC addresses. Because the Security Appliances are using the same IP address, when a failover occurs, it breaks the mapping between the IP address and MAC address in the ARP cache of all clients and network resources. The Secondary Security Appliance must issue an ARP request, announcing the new MAC address/IP address pair. Until this ARP request propagates through the network, traffic intended for the Primary Security Appliance's MAC address can be lost.

The Virtual MAC address greatly simplifies this process by using the same MAC address for both the Primary and Secondary Security Appliances. When a failover occurs, all routes to and from the Primary Security Appliance are still valid for the Secondary Security Appliance. All clients and remote sites continue to use the same Virtual MAC address and IP address without interruption.

By default, this Virtual MAC address is provided by the SonicWall firmware and is different from the physical MAC address of either the Primary or Secondary Security Appliances. This eliminates the possibility of configuration errors and ensures the uniqueness of the Virtual MAC address, which prevents possible conflicts. Optionally, you can manually configure the Virtual MAC address on **High Availability > Monitoring Settings**.

The Virtual MAC setting is available even if Stateful High Availability is not licensed. When Virtual MAC is enabled, it is always used even if Stateful Synchronization is not enabled.

Dynamic WAN Interfaces with PPPoE HA

NOTE: Dynamic WAN interfaces with PPPoE HA is not supported on the SuperMassive 9800. Only the DHCP Server dynamic WAN mode is supported.

PPPoE can be enabled on interfaces in non-stateful mode, HA Active/Standby mode. PPPoE HA provides HA where a Secondary Security Appliance assumes connection to the PPPoE server when the Active Security Appliance fails.

NOTE: One WAN interface must be configured as PPPoE; see Configuring a WAN Interface on page 292.

After the Active unit connects to the PPPoE server, the Security Appliance synchronizes the PPPoE session ID and server name to the Secondary unit.

When the Active Security Appliance fails, it terminates the PPPoE HA connection on the client side by timing out. The Secondary Security Appliance connects to the PPPoE server, terminates the original connection on the server side, and starts a new PPPoE connection. All pre-existing network connections are rebuilt, the PPPoE sessions are re-established, and the PPP process is renegotiated.

Stateful Synchronization with DHCP

DHCP can be enabled on interfaces in both Active/Standby (non-stateful) and Stateful Synchronization modes.

Only the Active Security Appliance can get a DHCP lease. The Active Security Appliance synchronizes the DHCP IP address along with the DNS and gateway addresses to the Secondary Security Appliance. The DHCP client ID is also synchronized, allowing this feature to work even without enabling Virtual MAC.

During a failover, the Active Security Appliance releases the DHCP lease and, as it becomes the Active unit, the Secondary Security Appliance renews the DHCP lease using the existing DHCP IP address and client ID. The IP address does not change, and network traffic, including VPN tunnel traffic, continues to pass.

If the Active Security Appliance does not have an IP address when failover occurs, the Secondary Security Appliance starts a new DHCP discovery.

Stateful Synchronization with DNS Proxy

DNS Proxy supports stateful synchronization of DNS cache. When the DNS cache is added, deleted, or updated dynamically, it synchronizes to the idle Security Appliance.

About HA Monitoring

On **MANAGE | System Setup > High Availability > Monitoring Settings**, you can configure both physical and logical interface monitoring:

- By enabling physical interface monitoring, you enable link detection for the designated HA interfaces. The link is sensed at the physical layer to determine link viability.
- Logical monitoring involves configuring the SonicWall to monitor a reliable device on one or more of the connected networks.

Failure to periodically communicate with the device by the Active unit in the HA Pair triggers a failover to the Standby unit. If neither unit in the HA Pair can connect to the device, no action is taken.

The Primary and Secondary IP addresses configured on **High Availability > Monitoring Settings** can be configured on LAN or WAN interfaces, and are used for multiple purposes:

- As independent management addresses for each unit (supported on all physical interfaces)
- To allow synchronization of licenses between the Standby unit and the SonicWall licensing server
- As the source IP addresses for the probe pings sent out during logical monitoring

Configuring unique management IP addresses for both units in the HA Pair allows you to log in to each unit independently for management purposes. Note that non-management traffic is ignored if it is sent to one of these IP addresses. The Primary and Secondary Security Appliances' unique LAN IP addresses cannot act as an active gateway; all systems connected to the internal LAN needs to use the virtual LAN IP address as their gateway.

If WAN monitoring IP addresses are configured, then X0 monitoring IP addresses are not required. If WAN monitoring IP addresses are not configured, then X0 monitoring IP addresses are required, because in such a scenario the Standby unit uses the X0 monitoring IP address to connect to the licensing server with all traffic routed through the Active unit.

The management IP address of the Secondary/Standby unit is used to allow license synchronization with the SonicWall licensing server, which handles licensing on a per-Security Appliance basis (not per-HA Pair). Even if the Secondary unit was already registered on MySonicWall before creating the HA association, you must use the link on **MANAGE | Updates > Licenses** to connect to the SonicWall server while accessing the Secondary Security Appliance through its management IP address (for more information, see *SonicOS Updates*).

When using logical monitoring, the HA Pair pings the specified Logical Probe IP address target from the Primary as well as from the Secondary unit. The IP address set in the Primary IP Address or Secondary IP Address field is used as the source IP address for the ping. If both units can successfully ping the target, no failover occurs. If both cannot successfully ping the target, no failover occurs, as SonicOS assumes that the problem is with the target, and not the Security Appliances. If one Security Appliance can ping the target but the other cannot, however, the HA Pair failovers to the unit that can ping the target.

The configuration tasks on **High Availability > Monitoring Settings** are performed on the Primary unit and then are automatically synchronized to the Secondary.

About Active/Standby HA

HA allows two identical Security Appliances running SonicOS to be configured to provide a reliable, continuous connection to the public Internet. One Security Appliance is configured as the Primary unit, and an identical

Security Appliance is configured as the Secondary unit. In the event of the failure of the Primary Security Appliance, the Secondary Security Appliance takes over to secure a reliable connection between the protected network and the Internet. Two Security Appliances configured in this way are also known as a High Availability Pair (HA Pair).

Active/Standby HA provides standard, high availability, and hardware failover functionality with the option of enabling stateful HA and Active/Active DPI.

HA provides a way to share licenses between two Security Appliances when one is acting as a high availability system for the other. To use this feature, you must register the Security Appliances on MySonicWall as Associated Products. Both Security Appliances must be the same SonicWall model.

Topics:

- Benefits of Active/Standby HA on page 680
- How Active/Standby HA Works on page 680

Benefits of Active/Standby HA

- **Increased network reliability** – In a High Availability configuration, the Secondary Security Appliance assumes all network responsibilities when the Primary unit fails, ensuring a reliable connection between the protected network and the Internet.
- **Cost-effectiveness** – High Availability is a cost-effective option for deployments that provide high availability by using redundant Security Appliances. You do not need to purchase a second set of licenses for the Secondary unit in a High Availability Pair.
- **Virtual MAC for reduced convergence time after failover** – The Virtual MAC address setting allows the HA Pair to share the same MAC address, which dramatically reduces convergence time following a failover. Convergence time is the amount of time it takes for the devices in a network to adapt their routing tables to the changes introduced by high availability. By default, the Virtual MAC address is provided by the SonicWall firmware and is different from the physical MAC address of either the Primary or Secondary Security Appliances.

How Active/Standby HA Works

NOTE: The TZ300 series and TZ400 series Security Appliances can be operated in Active/Standby HA mode without Stateful Synchronization. The SOHO W does not support High Availability with or without Stateful Synchronization.

HA requires one SonicWall Security Appliance configured as the Primary SonicWall, and an identical Security Appliance configured as the Secondary SonicWall. During normal operation, the Primary SonicWall is in an Active state and the Secondary SonicWall in an Standby state. If the Primary device loses connectivity, the Secondary SonicWall transitions to Active mode and assumes the configuration and role of Primary, including the interface IP addresses of the configured interfaces.

Basic Active/Standby HA provides stateless high availability. After a failover to the Secondary Security Appliance, all the pre-existing network connections must be re-established, including the VPN tunnels that must be re-negotiated. Stateful Synchronization can be licensed and enabled separately. For more information, see About Stateful Synchronization on page 681.

The failover applies to loss of functionality or network-layer connectivity on the Primary SonicWall. The failover to the Secondary SonicWall occurs when critical services are affected, physical (or logical) link failure is detected on monitored interfaces, or when the Primary SonicWall loses power. The Primary and Secondary SonicWall devices are currently only capable of performing Active/Standby High Availability or Active/Active DPI – complete Active/Active high availability is not supported at present.

There are two types of synchronization for all configuration settings:

- **Incremental** – If the timestamps are in sync and a change is made on the Active unit, an incremental synchronization is pushed to the Standby unit.
- **Complete** –If the timestamps are out of sync and the Standby unit is available, a complete synchronization is pushed to the Standby unit. When incremental synchronization fails, a complete synchronization is automatically attempted.

About Stateful Synchronization

Stateful Synchronization provides dramatically improved failover performance. When enabled, the network connections and VPN tunnel information is continuously synchronized between the two units so that the Secondary can seamlessly assume all network responsibilities if the Primary Security Appliance fails, with no interruptions to existing network connections.

NOTE: Stateful HA is included on NSA 4600 and higher NSA platforms and on all SuperMassive platforms. Stateful HA is supported on the TZ500 and higher TZ platforms and NSA 2600 and NSA 3600 platforms with an Extended or Stateful HA upgrade license. For licensing information, see *SonicOS Updates*.

Topics:

- Benefits of Stateful Synchronization on page 681
- How Does Stateful Synchronization Work? on page 681

Benefits of Stateful Synchronization

- **Improved reliability** - By synchronizing most critical network connection information, Stateful Synchronization prevents down time and dropped connections in case of Security Appliance failure.
- **Faster failover performance** - By maintaining continuous synchronization between the Primary and Secondary Security Appliances, Stateful Synchronization enables the Secondary Security Appliance to take over in case of a failure with virtually no down time or loss of network connections.
- **Minimal impact on CPU performance** - Typically less than 1% usage.
- **Minimal impact on bandwidth** - Transmission of synchronization data is throttled so as not interfere with other data.

How Does Stateful Synchronization Work?

Stateful Synchronization is not load-balancing. It is an active-standby configuration where the Primary Security Appliance handles all traffic. When Stateful Synchronization is enabled, the Primary Security Appliance actively communicates with the Secondary to update most network connection information. As the Primary Security Appliance creates and updates network connection information (such as VPN tunnels, active users, connection cache entries), it immediately informs the Secondary Security Appliance. This ensures that the Secondary Security Appliance is always ready to transition to the Active state without dropping any connections.

The synchronization traffic is throttled to ensure that it does not interfere with regular network traffic. All configuration changes are performed on the Primary Security Appliance and automatically propagated to the Secondary Security Appliance. The High Availability pair uses the same LAN and WAN IP addresses—regardless of which Security Appliance is currently Active.

When using SonicWall Global Management System (GMS) to manage the Security Appliances, GMS logs into the shared WAN IP address. In case of a failover, GMS administration continues seamlessly, and GMS administrators currently logged into the Security Appliance are not logged out; however, **Get** and **Post** commands may result in a time out with no reply returned.

Synchronized and non-synchronized information lists the information that is synchronized and information that is not currently synchronized by Stateful Synchronization.

Synchronized and non-synchronized information

Information that is Synchronized	Information that is not Synchronized
VPN information	Dynamic WAN clients (L2TP, PPPoE, and PPTP)
Basic connection cache	Deep Packet Inspection (GAV, IPS, and Anti Spyware)
FTP	IPHelper bindings (such as NetBIOS and DHCP)
Oracle SQL*NET	SYNFlood protection information
Real Audio	Content Filtering Service information
RTSP	VoIP protocols
GVC information	Dynamic ARP entries and ARP cache time outs
Dynamic Address Objects	Active wireless client information
DHCP server information	Wireless client packet statistics
Multicast and IGMP	Rogue AP list
Active users	
ARP	
SonicPoint status	
Wireless guest status	
License information	
Weighted Load Balancing information	
RIP and OSPF information	

Stateful Synchronization Example

In case of a failover, the following sequence of events occurs:

1. A PC user connects to the network, and the Primary Security Appliance creates a session for the user.
2. The Primary Security Appliance synchronizes with the Secondary Security Appliance. The Secondary now has all of the user's session information.
3. The administrator restarts the Primary unit.
4. The Secondary unit detects the restart of the Primary unit and switches from Standby to Active.
5. The Secondary Security Appliance begins to send gratuitous ARP messages to the LAN and WAN switches using the same Virtual MAC address and IP address as the Primary Security Appliance. No routing updates are necessary for downstream or upstream network devices.
6. When the PC user attempts to access a Web page, the Secondary Security Appliance has all of the user's session information and is able to continue the user's session without interruption.

About Active/Active DPI HA

IMPORTANT: Capture functionality is not supported in Active/Active DPI mode.

With Active/Active DPI enabled on a Stateful HA pair, the Deep Packet Inspection services are processed on the standby Security Appliance of an HA pair concurrently with the processing of Security Appliance, NAT, and other modules on the active Security Appliance. The following DPI services are affected:

- Intrusion Prevention Service (IPS)
- Gateway Anti-Virus (GAV)
- Gateway Anti-Spyware
- Application Control

To use the Active/Active DPI feature, you must configure an additional interface as the **Active/Active DPI Interface**. For example, if you choose to make X5 the Active/Active DPI Interface, you must physically connect X5 on the active unit to X5 on the standby unit in the HA pair. Certain packet flows on the active unit are selected and offloaded to the standby unit on the Active/Active DPI Interface. DPI is performed on the standby unit and then the results are returned to the active unit over the same interface. The remaining processing is performed on the active unit.

NOTE: Active/Active DPI is included on SuperMassive 9200, 9400, and 9600 platforms and is supported on the NSA 5600 and NSA 6600 only with extended licenses. For licensing information, see *SonicOS Updates*.

Benefits of Active/Active DPI HA

Active/Active DPI taps into the unused CPU cycles available in the standby unit, but the traffic still arrives and leaves through the active unit. The standby unit only sees the network traffic offloaded by the active unit, and processing of all modules other than DPI services is restricted to the active unit.

Active/Standby and Active/Active DPI Prerequisites

This section lists the supported platforms, provides recommendations and requirements for physically connecting the units, and describes how to register, associate, and license the units for High Availability.

Topics:

- Supported Platforms and Licensing for HA on page 683
- Physically Connecting Your Security Appliances on page 685
- Connecting the Active/Active DPI Interfaces for Active/Active DPI on page 686

Supported Platforms and Licensing for HA

Licenses included with the purchase of a SonicWall Security Appliance are shown in HA licenses available with SonicWall Security Appliances. Some platforms require additional licensing to use the HA features.

NOTE: HA licenses must be activated on each Security Appliance, either by registering the unit on MySonicWall from the SonicOS management interface, or by applying the license keyset to each unit if Internet access is not available.

The HA licenses included with the purchase of the SonicWall Security Appliance is shown in HA licenses available with SonicWall Security Appliances. Some platforms require additional licensing to use the Stateful Synchronization or Active/Active DPI features. SonicOS Expanded licenses or High Availability licenses can be purchased on MySonicWall or from a SonicWall reseller.

NOTE: Stateful High Availability licenses must be activated on each Security Appliance, either by registering the unit on MySonicWall from the SonicOS management interface, or by applying the license keyset to each unit if Internet access is not available.

HA licenses available with SonicWall Security Appliances

Platform	Active/Standby HA [a]	Stateful HA	A/A Clustering	A/A DPI
SM 9600	Included	Included	Included	Included
SM 9400	Included	Included	Included	Included
SM 9200	Included	Included	Included	Included
NS*a* 9650	Included	Included	Included	Included
NS*a* 9450	Included	Included	Included	Included
NS*a* 9250	Included	Included	Included	Included
NS*a* 6650	Included	Included	Expanded License	Expanded License
NSA 6600	Included	Included	Expanded License	Expanded License
NS*a* 5650	Included	Included	Expanded License	Expanded License
NSA 5600	Included	Included	Expanded License	Expanded License
NS*a* 4650	Included	Included	Expanded License	N/A
NSA 4600	Included	Included	Expanded License	N/A
NS*a* 3650	Included	Stateful HA License Expanded License	Expanded License	N/A
NSA 3600	Included	Stateful HA License Expanded License	Expanded License	N/A
NS*a* 2650	Included	Stateful HA License Expanded License	N/A	N/A
NSA 2600	Included	Stateful HA License Expanded License	N/A	N/A
TZ600/TZ600 P	Included	Expanded License Stateful HA License	N/A	N/A
TZ500/TZ500 W	Included	Expanded License Stateful HA License	N/A	N/A
TZ400/TZ400 W	Included	N/A	N/A	N/A
TZ350/TZ350 W	Included	N/A	N/A	N/A
TZ300/TZ300 W/ TZ300 P	Included	N/A	N/A	N/A
SOHO 250/ SOHO 250 W	N/A	N/A	N/A	N/A
SOHO W	N/A	N/A	N/A	N/A

a. NA = Feature not available

You can view system licenses on **MANAGE | Updates > Licenses**. This page also provides a way to log into MySonicWall and to apply licenses to a Security Appliance. For further information, see *SonicOS 6.5 Updates*.

There is also a way to synchronize licenses for an HA pair whose Security Appliances do not have Internet access. When live communication with SonicWall's licensing server is not permitted due to network policy, you can use license keysets to manually apply security services licenses to your Security Appliances. When you register a Security Appliance on MySonicWall, a license keyset is generated for the Security Appliance. If you add a new security service license, the keyset is updated. However, until you apply the licenses to the Security Appliance, it cannot perform the licensed services.

IMPORTANT: In a High Availability deployment without Internet connectivity, you must apply the license keyset to both of the Security Appliances in the HA pair.

You can view system licenses on **MANAGE | Updates > Licenses**. This page also provides a way to log into MySonicWall. For information about licensing, see *SonicOS 6.5 Updates*.

IMPORTANT: Even if you first register your Security Appliances on MySonicWall, you must individually register both the Primary and the Secondary Security Appliances from the SonicOS management interface while logged into the individual management IP address of each Security Appliance. This allows the Secondary unit to synchronize with the SonicWall license server and share licenses with the associated Primary Security Appliance. When Internet access is restricted, you can manually apply the shared licenses to both Security Appliances.

Physically Connecting Your Security Appliances

NOTE: For complete procedures for connecting your Security Appliances, see the *Quick Start Guide* for your Security Appliance. For procedures for connecting Active/Active Cluster Security Appliances, see Connecting the HA Ports for Active/Active Clustering on page 701 and Connecting Redundant Port Interfaces on page 701.

If you are connecting the Primary and Secondary Security Appliances to an Ethernet switch that uses the spanning tree protocol, be aware that it may be necessary to adjust the link activation time on the switch port to which the SonicWall interfaces connect. For example, on a Cisco Catalyst-series switch, it is necessary to activate **spanning tree port fast** for each port connecting to the SonicWall Security Appliance's interfaces.

High Availability requires additional physical connections among the affected SonicWall Security Appliances. For all modes, you need connections for HA Control and HA Data. Active/Active DPI requires an additional connection.

In any High Availability deployment, you must physically connect the LAN and WAN ports of all units to the appropriate switches.

It is important that the X0 interfaces from all units be connected to the same broadcast domain. Otherwise, traffic failover does not work. Also, X0 is the default redundant HA port; if the normal HA Control link fails, X0 is used to communicate heartbeats between units. Without X0 in the same broadcast domain, both units would become active if the HA Control link fails.

TIP: SonicOS Security Appliances now allow heartbeats to be exchanged between an HA pair across the MGMT interface in addition to the HA control interface.

A WAN connection to the Internet is useful for registering your Security Appliances on MySonicWall and for synchronizing licensing information. Unless live communication with SonicWall's licensing server is not permitted due to network policy, the WAN (X1) interface should be connected before registration and licensing are performed.

Connecting the Active/Active DPI Interfaces for Active/Active DPI

For Active/Active DPI, you must physically connect at least one additional interface, called the **Active/Active DPI Interface**, between the two Security Appliances in each HA pair, or Cluster Node. The connected interfaces must be the same number on both Security Appliances, and must initially appear as unused, unassigned interfaces in **MANAGE | System Setup | Network > Interfaces**. For example, you could connect X5 on the Primary unit to X5 on the Secondary if X5 is an unassigned interface. After enabling Active/Active DPI, the connected interface has a Zone assignment of **HA Data-Link**.

Certain packet flows on the active unit are selected and offloaded to the standby unit on the Active/Active DPI Interface. DPI is performed on the standby unit and then the results are returned to the active unit over the same interface.

Optionally, for port redundancy with Active/Active DPI, you can physically connect a second Active/Active DPI Interface between the two Security Appliances in each HA pair. This interface takes over transferring data between the two units during Active/Active DPI processing if the first Active/Active DPI Interface has a fault.

To connect the Active/Active DPI interfaces for Active/Active DPI:

1. Decide which interface to use for the additional connection between the Security Appliances in the HA pair. The same interface must be selected on each Security Appliance.
2. In the SonicOS Management Interface, navigate to **MANAGE | System Setup | Network > Interfaces** and ensure that the **Zone** is **Unassigned** for the intended Active/Active DPI Interface.
3. Using a standard Ethernet cable, connect the two interfaces directly to each other.
4. Optionally, for port redundancy with Active/Active DPI, physically connect a second Active/Active DPI Interface between the two Security Appliances in each HA pair.

Maintenance

Topics:

- Removing an HA Association on page 686
- Replacing a SonicWall Security Appliance on page 687

Removing an HA Association

You can remove the association between two SonicWall Security Appliances on MySonicWall at any time. You might need to remove an existing HA association if you replace an Security Appliance or reconfigure your network. For example, if one of your SonicWall Security Appliances fails, you need to replace it. Or, you might need to switch the HA Primary Security Appliance with the Secondary, or HA Secondary, unit after a network reconfiguration. In either case, you must first remove the existing HA association, and then create a new association that uses a new Security Appliance or changes the parent-child relationship of the two units (see Replacing a SonicWall Security Appliance on page 687).

To remove the association between two registered SonicWall Security Appliances:

1. Login to MySonicWall.
2. In the left navigation bar, click **My Products**.
3. On the **My Products** page, under **Registered Products**, scroll down to find the secondary Security Appliance from which you want to remove associations. Click the product **name** or **serial number.**

4 On the **Service Management - Associated Products** page, scroll down to the **Parent Product** section, just above the **Associated Products** section.

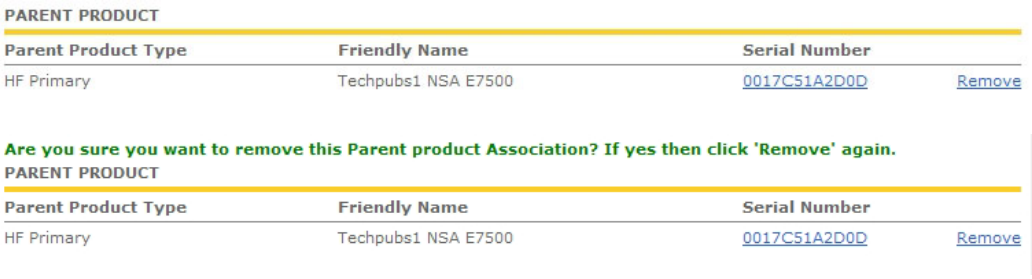
5 Under **Parent Product**, to remove the association for this Security Appliance:

   a Click **Remove**.

   b Wait for the page to reload.

   c Scroll down.

   d Click **Remove** again.



Replacing a SonicWall Security Appliance

If your SonicWall Security Appliance has a hardware failure while still under warranty, SonicWall replaces it. In this case, you need to remove the HA association containing the failed Security Appliance in MySonicWall, and add a new HA association that includes the replacement. If you contact SonicWall Technical Support to arrange the replacement (known as an RMA), Support often takes care of this for you.

After replacing the failed Security Appliance in your equipment rack with the new unit, you can update MySonicWall and your SonicOS configuration.

Replacing a failed HA Primary unit is slightly different than replacing an HA Secondary unit. Both procedures are provided in these sections:

- Replacing an HA Primary Unit on page 687
- Replacing an HA Secondary Unit on page 688

Replacing an HA Primary Unit

To replace an HA Primary unit:

1 In the SonicOS management interface of the remaining SonicWall Security Appliance (the Secondary unit), on the High Availability page, uncheck **Enable High Availability** to disable it.

2 Check **Enable High Availability**.

   The old Secondary unit now becomes the Primary unit. Its serial number is automatically displayed in the Primary SonicWall Serial Number field.

3 Type the serial number for the replacement unit into the **Secondary SonicWall Serial Number** field.

4 Click **Synchronize Settings**.

5 On MySonicWall, remove the old HA association. See Removing an HA Association on page 686.

6 On MySonicWall, register the replacement SonicWall Security Appliance and create an HA association with the new Primary (original Secondary) unit as the HA Primary, and the replacement unit as the HA Secondary. See *SonicOS Updates*.

7 Contact SonicWall Technical Support to transfer the security services licenses from the former HA Pair to the new HA Pair.

This step is required when the HA Primary unit has failed because the licenses are linked to the Primary unit in an HA Pair.

Replacing an HA Secondary Unit

To replace an HA Secondary unit:

1 On MySonicWall, remove the old HA association as described in Removing an HA Association on page 686.

2 On MySonicWall, register the replacement SonicWall Security Appliance.

3 Create an HA association with the original HA Primary, using the replacement unit as the HA Secondary as described in Replacing an HA Primary Unit on page 687.

Active/Active Clustering

NOTE: Active/Active Clustering is supported on NSA 3600 and above Security Appliances. See HA licenses available with SonicWall Security Appliances and Licensing requirements for A/A Clustering

About Active/Active Clustering

An Active/Active Cluster is formed by a grouping up to four Cluster Nodes, with multiple Active units processing traffic (as multiple gateways), doing DPI, and sharing the network load. A Cluster Node can consist of a Stateful HA pair, a Stateless HA pair with standard failover, or a single standalone unit, in which case Stateful Failover and Active/Active DPI are not available. Dynamic state synchronization is only available in a Cluster Node if it is a Stateful HA pair. The traditional SonicWall High Availability protocol or Stateful HA protocol is used for communication within the Cluster Node, between the units in the HA pair.

When a Cluster Node is a Stateful HA pair, Active/Active DPI can be enabled within the Cluster Node for higher performance.

With Active/Active Clustering, you can assign certain traffic flows to each node in the cluster, providing load sharing in addition to redundancy, and supporting a much higher throughput without a single point of failure.

A typical recommended setup includes four Security Appliances of the same SonicWall model configured as two Cluster Nodes, where each node consists of one Stateful HA pair. For larger deployments, the cluster can include eight Security Appliances, configured as four Cluster Nodes (or HA pairs). Within each Cluster Node, Stateful HA keeps the dynamic state synchronized for seamless failover with zero loss of data on a single point of failure. Stateful HA is not required, but is highly recommended for best performance during failover.

Load sharing is accomplished by configuring different Cluster Nodes as different gateways in your network. Typically this is handled by another device downstream (closer to the LAN devices) from the Active/Active Cluster, such as a DHCP server or a router.

A Cluster Node can also be a single Security Appliance, allowing an Active/Active cluster setup to be built using two Security Appliances. In case of a fault condition on one of the Security Appliances in this deployment, the failover is not stateful because neither Security Appliance in the Cluster Node has an HA Secondary.

Redundancy is achieved at several levels with Active/Active Clustering:

- The cluster provides redundant Cluster Nodes, each of which can handle the traffic flows of any other Cluster Node, if a failure occurs.

- The Cluster Node consists of a Stateful HA pair, in which the Secondary Security Appliance can assume the duties of the Primary unit in case of failure.
- Port redundancy, in which an unused port is assigned as a secondary to another port, provides protection at the interface level without requiring failover to another Security Appliance or node.
- Active/Active DPI can be enabled, providing increased throughput within each Cluster Node.

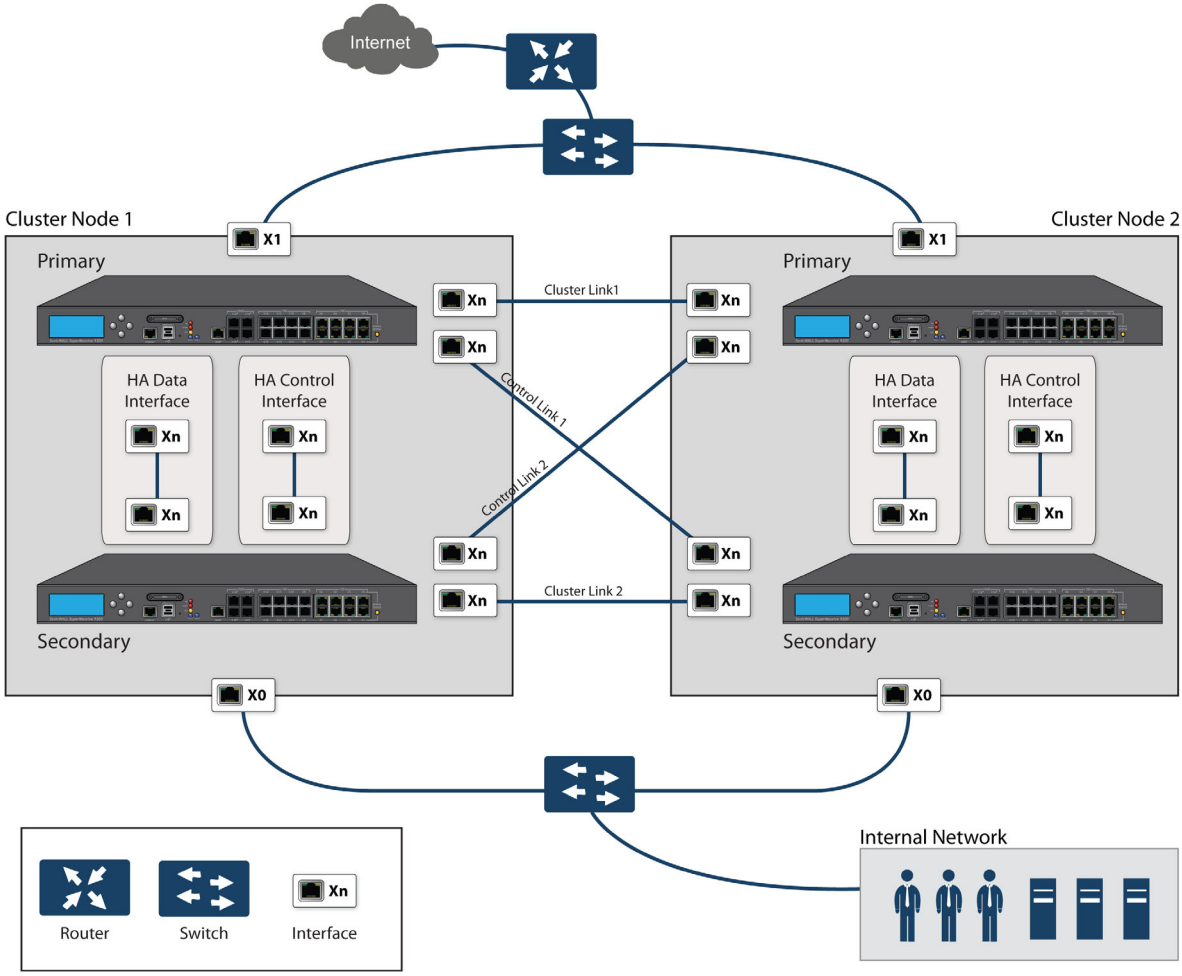
Topics:

- Example: Active/Active Clustering – Four-Unit Deployment on page 689
- Example: Active/Active Clustering – Two-Unit Deployment on page 691
- Benefits of Active/Active Clustering on page 691
- How Does Active/Active Clustering Work? on page 692
- Features Supported with Active/Active Clustering on page 697

Example: Active/Active Clustering – Four-Unit Deployment

Active/Active four-unit cluster shows a four-unit cluster. Each Cluster Node contains one HA pair. The designated HA ports of all four Security Appliances are connected to a Layer 2 switch. These ports are used for Cluster Node management and monitoring state messages sent over SVRRP, and for configuration synchronization. The two units in each HA pair are also connected to each other using another interface (shown as the Xn interface). This is the Active/Active DPI Interface necessary for Active/Active DPI. With Active/Active DPI enabled, certain packets are offloaded to the standby unit of the HA pair for DPI processing.

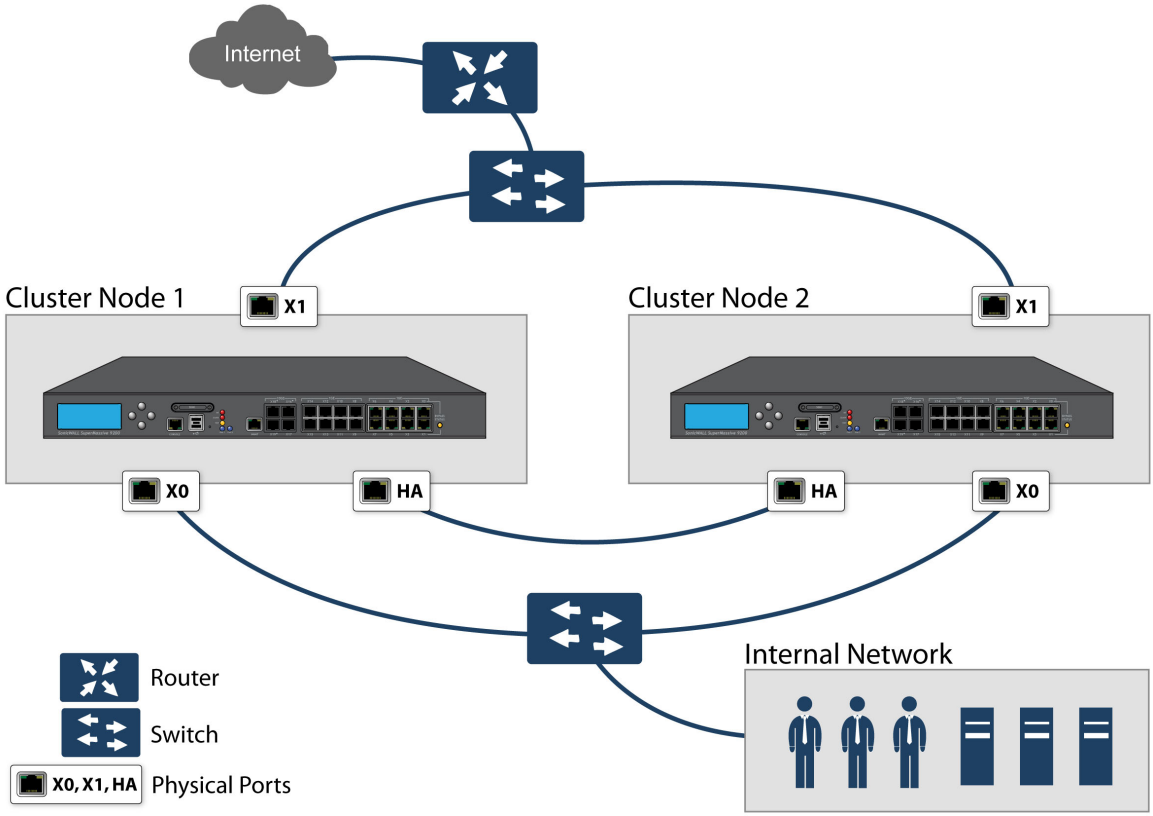
Active/Active four-unit cluster



Example: Active/Active Clustering – Two-Unit Deployment

Active/Active two-unit cluster shows a two-unit cluster. In a two-unit cluster, HA pairs are not used. Instead, each Cluster Node contains a single Security Appliance. The designated HA ports on the two Security Appliances are connected directly to each other using a cross-over cable. The SonicWall Virtual Router Redundancy Protocol (SVRRP) uses this HA port connection to send Cluster Node management and monitoring state messages. SVRRP management messages are initiated on the Master Node, and monitoring information is communicated from every Security Appliance in the cluster. The HA port connection is also used for configuration synchronization between Cluster Nodes.

Active/Active two-unit cluster



Benefits of Active/Active Clustering

The benefits of Active/Active Clustering include the following:

- All the Security Appliances in the cluster are utilized to derive maximum throughput
- Can run in conjunction with Active/Active DPI to perform concurrent processing of IPS, GAV, Anti-Spyware, and App Rules services, which are the most processor intensive, on the standby Security Appliance in each HA pair while the active Security Appliance performs other processing
- Load sharing is supported by allowing the assignment of particular traffic flows to each node in the cluster
- All nodes in the cluster provide redundancy for the other nodes, handling traffic as needed if other nodes go down
- Interface redundancy provides secondary for traffic flow without requiring failover

- Both Full Mesh and non-Full Mesh deployments are supported

How Does Active/Active Clustering Work?

There are several important concepts that are introduced for Active/Active Clustering.

Topics:

- About Cluster Nodes on page 692
- About the Cluster on page 692
- About Virtual Groups on page 694
- About SVRRP on page 695
- About Failover on page 695
- About DPI with Active/Active Clustering on page 696
- About High Availability Monitoring with Active/Clustering on page 696

About Cluster Nodes

An Active/Active Cluster is formed by a collection of Cluster Nodes. A Cluster Node can consist of a Stateful HA pair, a Stateless HA pair or a single standalone unit. Dynamic state synchronization is only available in a Cluster Node if it is a Stateful HA pair. The traditional SonicWall High Availability protocol or Stateful HA protocol is used for communication within the Cluster Node, between the units in the HA pair.

When a Cluster Node is a Stateful HA pair, Active/Active DPI can be enabled within the Cluster Node for higher performance.

About the Cluster

All Security Appliances in the Cluster must be of same product model and be running the same firmware version.

Within the cluster, all Security Appliances are connected and communicating with each other; see Active/Active two-node cluster. For communication between Cluster Nodes, a new protocol, called SonicWall Virtual Router Redundancy Protocol (SVRRP), is used. Cluster Node management and monitoring state messages are sent using SVRRP.

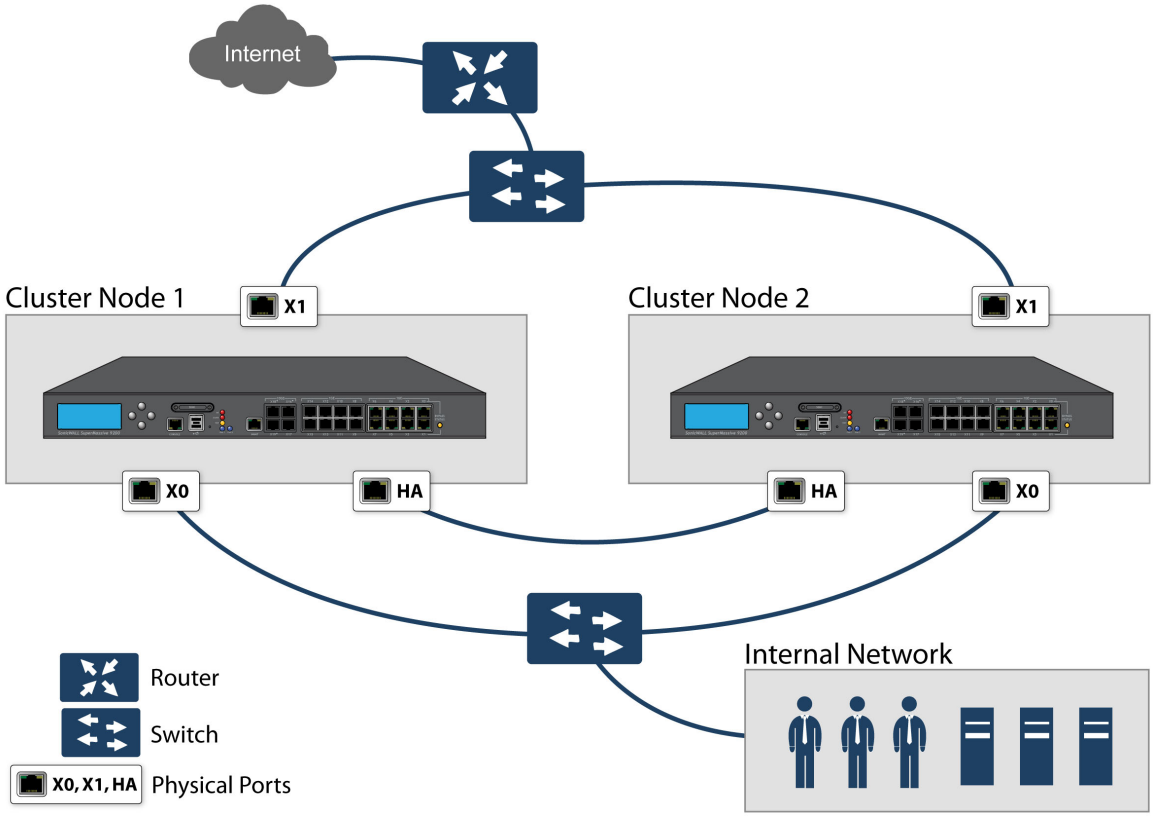
All Cluster Nodes share the same configuration, which is synchronized by the Master Node. The Master Node is also responsible for synchronizing firmware to the other nodes in the cluster. The HA port connection is used to synchronize configuration and firmware updates.

Dynamic state is not synchronized across Cluster Nodes, but only within a Cluster Node. When a Cluster Node contains an HA pair, Stateful HA can be enabled within that Cluster Node, with the advantages of dynamic state synchronization and stateful failover as needed. In the event of the failure of an entire Cluster Node, the failover is stateless. This means that preexisting network connections must be rebuilt. For example, Telnet and FTP sessions must be re-established and VPN tunnels must be renegotiated.

About Failover on page 695 provides more information about how failover works.

The maximum number of Cluster Nodes in a cluster is currently limited to four. If each Cluster Node is an HA pair, the cluster includes eight Security Appliances.

Active/Active two-node cluster



Actions Allowed Within the Cluster

The types of administrative actions that are allowed differ based on the state of the Security Appliance in the cluster. All actions are allowed for admin users with appropriate privileges on the active Security Appliance of the Master Node, including all configuration actions. A subset of actions are allowed on the active Security Appliance of Non-Master nodes, and even fewer actions are allowed on Security Appliances in the standby state. Administrative Actions Allowed lists the allowed actions for active Security Appliances of Non-Master nodes and standby Security Appliances in the cluster.

Administrative Actions Allowed

Administrative Action	Active Non-Master	Standby
Read-only actions	Allowed	Allowed
Registration on MySonicWall	Allowed	Allowed
License Synchronization with SonicWall License Manager	Allowed	Allowed
Diagnostic tools in **INVESTIGATE | Tools | System Diagnostics** (for information about these tools, see *SonicOS Investigation*).	Allowed	Allowed
Packet capture	Allowed	Allowed
HA Synchronize Settings (syncs settings to the HA peer within the node)	Not Allowed	Not allowed
HA Synchronize Firmware (syncs firmware to the HA peer within the node)	Allowed	Not allowed
Administrative logout of users	Allowed	Not allowed
Authentication tests (such as test LDAP, test RADIUS, test Authentication Agent)	Allowed	Not allowed

About Virtual Groups

Active/Active Clustering also supports the concept of Virtual Groups. Currently, a maximum of four Virtual Groups are supported.

A Virtual Group is a collection of virtual IP addresses for all the configured interfaces in the cluster configuration (unused/unassigned interfaces do not have virtual IP addresses). When Active/Active Clustering is enabled for the first time, the configured IP addresses for the interfaces on that Security Appliance are converted to virtual IP addresses for Virtual Group 1. Thus, Virtual Group 1 includes virtual IP addresses for X0, X1, and any other interfaces that are configured and assigned to a zone.

A Virtual Group can also be thought of as a logical group of traffic flows within a failover context, in that the logical group of traffic flows can failover from one node to another depending upon the fault conditions encountered. Each Virtual Group has one Cluster Node acting as the owner and one or more Cluster Nodes acting as standby. A Virtual Group is only owned by one Cluster Node at a time, and that node becomes the owner of all the virtual IP addresses associated with that Virtual Group. The owner of Virtual Group 1 is designated as the Master Node, and is responsible for synchronizing configuration and firmware to the other nodes in the cluster. If the owner node for a Virtual Group encounters a fault condition, one of the standby nodes becomes the owner.

As part of the configuration for Active/Active Clustering, the serial numbers of other Security Appliances in the cluster are entered into the SonicOS management interface, and a ranking number for the standby order is assigned to each. When the Active/Active Clustering configuration is applied, up to three additional Virtual Groups are created, corresponding to the additional Cluster Nodes added, but virtual IP addresses are not created for these Virtual Groups. You need to configure these virtual IP addresses on **MANAGE | System Setup | Network > Interfaces**.

There are two factors in determining Virtual Group ownership (which Cluster Node owns which Virtual Group):

- **Rank of the Cluster Node** – The rank is configured in the SonicOS management interface to specify the priority of each node for taking over the ownership of a Virtual Group.
- **Virtual Group Link Weight of the Cluster Nodes** – This is the number of interfaces in the Virtual Group that are up and have a configured virtual IP address.

When more than two Cluster Nodes are configured in a cluster, these factors determine the Cluster Node that is best able to take ownership of the Virtual Group. In a cluster with two Cluster Nodes, one of which has a fault, naturally the other takes ownership.

SVRRP is used to communicate Virtual Group link status and ownership status to all Cluster Nodes in the cluster.

The owner of Virtual Group 1 is designated as the Master Node. Configuration changes and firmware updates are only allowed on the Master Node, which uses SVRRP to synchronize the configuration and firmware to all the nodes in the cluster. On a particular interface, virtual IP addresses for Virtual Group 1 must be configured before other Virtual Groups can be configured.

Load Sharing and Multiple Gateway Support

The traffic for the Virtual Group is processed only by the owner node. A packet arriving on a Virtual Group leaves the Security Appliance on the same Virtual Group. In a typical configuration, each Cluster Node owns a Virtual Group, and therefore processes traffic corresponding to one Virtual Group.

This Virtual Group functionality supports a multiple gateway model with redundancy. In a deployment with two Cluster Nodes, the X0 Virtual Group 1 IP address can be one gateway and the X0 Virtual Group 2 IP address can be another gateway. It is up to the network administrator to determine how the traffic is allocated to each gateway. For example, you could use a smart DHCP server which distributes the gateway allocation to the PCs on the directly connected client network, or you could use policy-based routes on a downstream router.

When Active/Active Clustering is enabled, the SonicOS internal DHCP server is turned off and cannot be enabled. Networks needing a DHCP server can use an external DHCP server which is aware of the multiple gateways, so that the gateway allocation can be distributed.

NOTE: When Active/Active Clustering is enabled, the SonicOS internal DHCP server is turned off.

Effect on Related Configuration Pages

When Active/Active Clustering is initially enabled, the existing IP addresses for all configured interfaces are automatically converted to virtual IP addresses for Virtual Group 1. When Virtual Group 1 or any Virtual Group is created, default interface objects are created for virtual IP addresses with appropriate names, such as `Virtual Group 1` or `Virtual Group 2`. The same interface can have multiple virtual IP addresses, one for each Virtual Group that is configured. You can view these virtual IP addresses in **MANAGE | System Setup | Network > Interfaces**.

NOTE: All Cluster Nodes in the Active/Active cluster share the same configuration

A virtual MAC address is associated with each virtual IP address on an interface and is generated automatically by Sonic OS. The virtual MAC address is created in the format `00-17-c5-6a-XX-YY`, where `XX` is the interface number such as `03` for port X3, and `YY` is the internal group number such as `00` for Virtual Group 1, or `01` for Virtual Group 2.

NOTE: The Active/Active virtual MAC address is different from the High Availability virtual MAC address. The High Availability virtual MAC address functionality is not supported when Active/Active Clustering is enabled.

NAT policies are automatically created for the affected interface objects of each Virtual Group. These NAT policies extend existing NAT policies for particular interfaces to the corresponding virtual interfaces. You can view these NAT policies in **MANAGE | Policies | Rules > NAT Policies**. Additional NAT policies can be configured as needed and can be made specific to a Virtual Group if desired. For information about NAT policies, see *SonicOS Policies*.

After Active/Active Clustering is enabled, you must select the Virtual Group number during configuration when adding a VPN policy.

About SVRRP

For communication between Cluster Nodes in an Active/Active cluster, a new protocol called SonicWall Virtual Router Redundancy Protocol (SVRRP) is used. Cluster Node management and monitoring state messages are sent using SVRRP over the Active/Active Cluster links.

SVRRP is also used to synchronize configuration changes, firmware updates, and signature updates from the Master Node to all nodes in the cluster. In each Cluster Node, only the active unit processes the SVRRP messages.

If the Active/Active Cluster links fail, SVRRP heartbeat messages are sent on the X0 interface. However, while the Active/Active Cluster links are down, configuration is not synchronized. Firmware or signature updates, changes to policies, and other configuration changes cannot be synchronized to other Cluster Nodes until the Active/Active Cluster links are fixed.

About Failover

There are two types of failover that can occur when Active/Active Clustering is enabled:

High Availability failover	Within an HA pair, the Secondary unit takes over for the Primary. If Stateful HA is enabled for the pair, the failover occurs without interruption to network connections.

Active/Active failover	If all the units in the owner node for a Virtual Group encounter a fault condition, then the standby node for the Virtual Group takes over the Virtual Group ownership. Active/Active failover transfers ownership of a Virtual Group from one Cluster Node to another. The Cluster Node that becomes the Virtual Group owner also becomes the owner of all the virtual IP addresses associated with the Virtual Group and starts using the corresponding virtual MAC addresses. Active/Active failover is stateless, meaning that network connections are reset and VPN tunnels must be renegotiated. Layer 2 broadcasts inform the network devices of the change in topology as the Cluster Node which is the new owner of a Virtual Group generates ARP requests with the virtual MACs for the newly owned virtual IP addresses. This greatly simplifies the failover process as only the connected switches need to update their learning tables. All other network devices continue to use the same virtual MAC addresses and do not need to update their ARP tables, because the mapping between the virtual IP addresses and virtual MAC addresses is not broken.

When both High Availability failover and Active/Active failover are possible, HA failover is given precedence over Active/Active failover:

- HA failover can be stateful, whereas Active/Active failover is stateless.
- The standby Security Appliance in an HA pair is lightly loaded and has resources available for taking over the necessary processing, although it may already be handling DPI traffic if Active/Active DPI is enabled. The alternative Cluster Node might already be processing traffic comparable in amount to the failed unit, and could become overloaded after failover.

Active/Active failover always operates in Active/Active preempt mode. Preempt mode means that, after failover between two Cluster Nodes, the original owner node for the Virtual Group seizes the active role from the standby node after the owner node has been restored to a verified operational state. The original owner has a higher priority for a Virtual Group due to its higher ranking if all virtual IP interfaces are up and the link weight is the same between the two Cluster Nodes.

About DPI with Active/Active Clustering

Active/Active DPI can be used along with Active/Active Clustering. When Active/Active DPI is enabled, it utilizes the standby Security Appliance in the HA pair for DPI processing.

For increased performance in an Active/Active cluster, enabling Active/Active DPI is recommended, as it utilizes the standby Security Appliance in the HA pair for Deep Packet Inspection (DPI) processing.

About High Availability Monitoring with Active/Clustering

When Active/Active Clustering is enabled, HA monitoring configuration is supported for the HA pair in each Cluster Node. The HA monitoring features are consistent with previous versions. HA monitoring can be configured for both physical/link monitoring and logical/probe monitoring. After logging into the Master Node, monitoring configuration needs to be added on a per Node basis from **MANAGE | System Setup | High Availability > Monitoring Settings**.

NOTE: High Availability > Monitoring Settings apply only to the HA pair that you are logged into, not to the entire cluster.

Physical interface monitoring enables link detection for the monitored interfaces. The link is sensed at the physical layer to determine link viability.

When physical interface monitoring is enabled, with or without logical monitoring enabled, HA failover takes precedence over Active/Active failover. If a link fails or a port is disconnected on the active unit, the standby unit in the HA pair becomes active.

NOTE: For interfaces with configured virtual IP addresses, Active/Active physical monitoring is implicit and is used to calculate the Virtual Group Link Weight. Physical monitoring cannot be disabled for these interfaces. This is different from HA monitoring.

Logical monitoring involves configuring SonicOS to monitor a reliable device on one or more of the connected networks. Failure to periodically communicate with the device by the active unit in the HA pair triggers a failover to the standby unit. If neither unit in the HA pair can connect to the device, the problem is assumed to be with the device and no failover occurs.

If both physical monitoring and logical monitoring are disabled, Active/Active failover occurs on link failure or port disconnect.

The Primary and Secondary IP addresses configured on **MANAGE | System Setup | High Availability > Monitoring Settings** can be configured on LAN or WAN interfaces, and are used for multiple purposes:

- As independent management addresses for each unit, regardless of the Active or Standby status of the unit (supported on all physical interfaces)
- To allow synchronization of licenses between the standby unit and the SonicWall licensing server
- As the source IP addresses for the probe pings sent out during logical monitoring

Configuring monitoring IP addresses for both units in the HA pair allows you to log in to each unit independently for management purposes. Note that non-management traffic is ignored if it is sent to one of the monitoring IP addresses. The Primary and Secondary Security Appliance's unique LAN IP addresses cannot act as an active gateway; all systems connected to the internal LAN need to use a virtual LAN IP address as their gateway.

NOTE: When HA Monitoring/Management IP addresses are configured only on WAN interfaces, they need to be configured on all the WAN interfaces for which a Virtual IP address has been configured.

The management IP address of the Secondary unit is used to allow license synchronization with the SonicWall licensing server, which handles licensing on a per-Security Appliance basis (not per-HA pair). Even if the standby unit was already registered on MySonicWall before creating the HA association, you must use the link on **MANAGE | Updates | Licenses** to connect to the SonicWall server while accessing the Secondary Security Appliance through its management IP address. This allows synchronization of licenses (such as the Active/Active Clustering or the Stateful HA license) between the standby unit and the SonicWall licensing server.

When using logical monitoring, the HA pair pings the specified Logical Probe IP address target from the Primary as well as from the Secondary SonicWall. The IP address set in the Primary IP Address or Secondary IP Address field is used as the source IP address for the ping. If both units can successfully ping the target, no failover occurs. If both cannot successfully ping the target, no failover occurs, as the SonicWalls assumes that the problem is with the target, and not the SonicWalls. But, if one SonicWall can ping the target but the other SonicWall cannot, the HA pair failsover to the SonicWall that can ping the target.

The configuration tasks on **MANAGE | System Setup | High Availability > Monitoring Settings** are performed on the Primary unit and then are automatically synchronized to the Secondary.

Features Supported with Active/Active Clustering

Topics:

- Caveats on page 698
- Backward Compatibility on page 698
- SonicPoint Compatibility on page 698
- WAN Load Balancing Compatibility on page 698

- Routing Topology and Protocol Compatibility on page 699

Caveats

When Active/Active Clustering is enabled, only static IP addresses can be used on the WAN.

These features are not supported when Active/Active Clustering is enabled:

- DHCP Server
- L3 Transparent Mode
- L2 Bridging / L2 Transparent Mode
- Dynamic DNS
- Wire Mode

These features are only supported on Virtual Group 1:

- SonicWall GVC
- SonicOS SSL VPN
- IP Helper

Backward Compatibility

The Active/Active Clustering feature is not backward compatible. When upgrading to SonicOS from a previous release that did not support Active/Active Clustering, it is highly recommended that you disable High Availability before exporting the preferences from an HA pair running a previous version of SonicOS. The preferences can then be imported without potential conflicts after upgrading.

SonicPoint Compatibility

There are two points to consider when using SonicWall SonicPoints or SonicWaves together with Active/Active Clustering:

- SonicPoints and SonicWaves only communicate with the Master node for downloading firmware and other aspects of operation.
- SonicPoints and SonicWaves need access to an independent DCHP server. SonicPoints and SonicWaves require a DHCP server to provide IP addresses to wireless clients, but the embedded SonicOS DHCP server is automatically disabled when Active/Active Clustering is enabled.

WAN Load Balancing Compatibility

When WAN Load Balancing (WLB) is enabled in an Active/Active Cluster, the same WLB interface configuration is used for all nodes in the cluster.

A WAN interface failure can trigger either a WLB failover, an HA pair failover, or an Active/Active failover to another Cluster Node, depending on:

- WAN goes down logically due to WLB probe failure – WLB failover
- Physical WAN goes down while Physical Monitoring is enabled – HA pair failover
- Physical WAN goes down while Physical Monitoring is not enabled – Active/Active failover

Routing Topology and Protocol Compatibility

This section describes the current limitations and special requirements for Active/Active Clustering configurations with regard to routing topology and routing protocols.

Topics:

- Layer-2 Bridge Support on page 699
- OSPF Support on page 699
- RIP Support on page 699
- BGP Support on page 699
- Asymmetric Routing In Cluster Configurations on page 699

Layer-2 Bridge Support

Layer-2 Bridged interfaces are not supported in a cluster configuration.

OSPF Support

OSPF is supported with Active/Active Clustering. When enabled, OSPF runs on the OSPF-enabled interfaces of each active Cluster Node. From a routing perspective, all Cluster Nodes appear as parallel routers, each with the virtual IP address of the Cluster Node's interface. In general, any network advertised by one node is advertised by all other nodes.

The OSPF router-ID of each Cluster Node must be unique and is derived from the router-ID configured on the Master node as follows:

- If the user enters **0** or **0.0.0.0** for the router-ID in the OSPF configuration, each node's router-ID is assigned the node's X0 virtual IP address.
- If the user enters any value other than **0** or **0.0.0.0** for the router-ID, each node is assigned a router-ID with consecutive values incremented by one for each node. For example, in a 4-node cluster, if the router-ID `10.0.0.1` was configured on the Master node, the router-IDs assigned would be:
  - Node 1: `10.0.0.1`
  - Node 2: `10.0.0.2`
  - Node 3: `10.0.0.3`
  - Node 4: `10.0.0.4`

RIP Support

RIP is supported, and like OSPF, runs on the RIP-enabled interfaces of each Cluster Node. From a routing perspective, all Cluster Nodes appear as parallel routers with the virtual IP address of the Cluster Node's interface. In general, any network advertised by one node is advertised by all other nodes.

BGP Support

BGP is supported in clusters, and also appears as parallel BGP routers using the virtual IP address of the Cluster Node's interface. As with OSPF and RIP, configuration changes made on the Master node are applied to all other Cluster Nodes. In the case of BGP, where configuration may only be applied through the CLI, the configuration is distributed when the running configuration is saved with the **`write file`** CLI command.

Asymmetric Routing In Cluster Configurations

SonicOS supports asymmetric routing for traffic flows across different layer 2 bridged pair interfaces on the Security Appliance or when it flows across different Security Appliances in a high availability cluster.

Active/Active Clustering Prerequisites

NOTE: In addition to the requirements described in this section, ensure that you have completed the prerequisites described in Active/Standby and Active/Active DPI Prerequisites on page 683.

For Active/Active Clustering, additional physical connections are required:

- **Active/Active Cluster Link**—Each Active/Active cluster link must be at least a 100MB interface, but a 1GB interface is preferred.

Active/Active Clustering configuration can include configuring Virtual Group IDs and redundant ports. Procedures are provided in this section for both of these tasks within **High Availability > Base Setup** on page 703.

Topics:

- Licensing Requirements for Active/Active Clustering on page 700
- Connecting the HA Ports for Active/Active Clustering on page 701
- Connecting Redundant Port Interfaces on page 701

Licensing Requirements for Active/Active Clustering

Active/Active Clustering licenses included with the purchase of a SonicWall Security Appliance are shown in Licensing requirements for A/A Clustering. Some platforms require additional licensing to use the Active/Active Clustering features. SonicOS Expanded licenses can be purchased on MySonicWall or from a SonicWall reseller.

NOTE: Active/Active Clustering licenses must be activated on each Security Appliance, either by registering the unit on MySonicWall from the SonicOS management interface, or by applying the license keyset to each unit if Internet access is not available.

Licensing requirements for A/A Clustering

Platform	License requirement[a]	Platform	License requirement
NSA 6600	Expanded license	SM 9600	Included
NS*a* 5650	Expanded license	SM 9400	Included
NSA 5600	Expanded license	SM 9200	Included
NS*a* 4650	Expanded license		
NSA 4600	Expanded license	TZ600/600P	N/A
NS*a* 3650	Expanded license	TZ500/TZ500 W	N/A
NSA 3600	Expanded license	TZ400/TZ400 W	N/A
NS*a* 2650	N/A	TZ350/TZ350 W	N/A
NSA 2600	N/A	TZ300/TZ300P/TZ300 W	N/A
		SOHO 250/SOHO 250 W	N/A
		SOHO W	N/A

a. N/A = not applicable; Included = included with base license

You can view system licenses on **MANAGE | Updates > Licenses**. This page also provides a way to log into MySonicWall. For information about licensing, see *SonicOS Updates*.

When the Security Appliances in the Active/Active cluster have Internet access, each Security Appliance in the cluster must be individually registered from the SonicOS management interface while you are logged into the individual management IP address of each Security Appliance. This allows the Secondary units to synchronize

with the SonicWall licensing server and share licenses with the associated Primary Security Appliances in each HA pair.

Connecting the HA Ports for Active/Active Clustering

For Active/Active Clustering, you must physically connect the designated HA ports of all units in the Active/Active cluster to the same Layer 2 network.

SonicWall recommends connecting all designated HA ports to the same Layer 2 switch. You can use a dedicated switch or simply use some ports on an existing switch in your internal network. All of these switch ports must be configured to allow Layer 2 traffic to flow freely amongst them.

In the case of a two-unit Active/Active cluster deployment, where the two Cluster Nodes each have only a single Security Appliance, you can connect the HA ports directly to each other using a cross-over cable. No switch is necessary in this case.

The SonicWall Virtual Router Redundancy Protocol (SVRRP) uses this HA port connection to send Cluster Node management and monitoring state messages. SVRRP management messages are initiated on the Master Node, and monitoring information is communicated from every Security Appliance in the cluster.

The HA port connection is also used to synchronize configuration from the Master Node to the other Cluster Nodes in the deployment. This includes firmware or signature upgrades, policies for VPN and NAT, and other configuration.

Connecting Redundant Port Interfaces

You can assign an unused physical interface as a redundant port to a configured physical interface called the Primary interface. On each Cluster Node, each primary and redundant port pair must be physically connected to the same switch, or preferably, to redundant switches in the network.

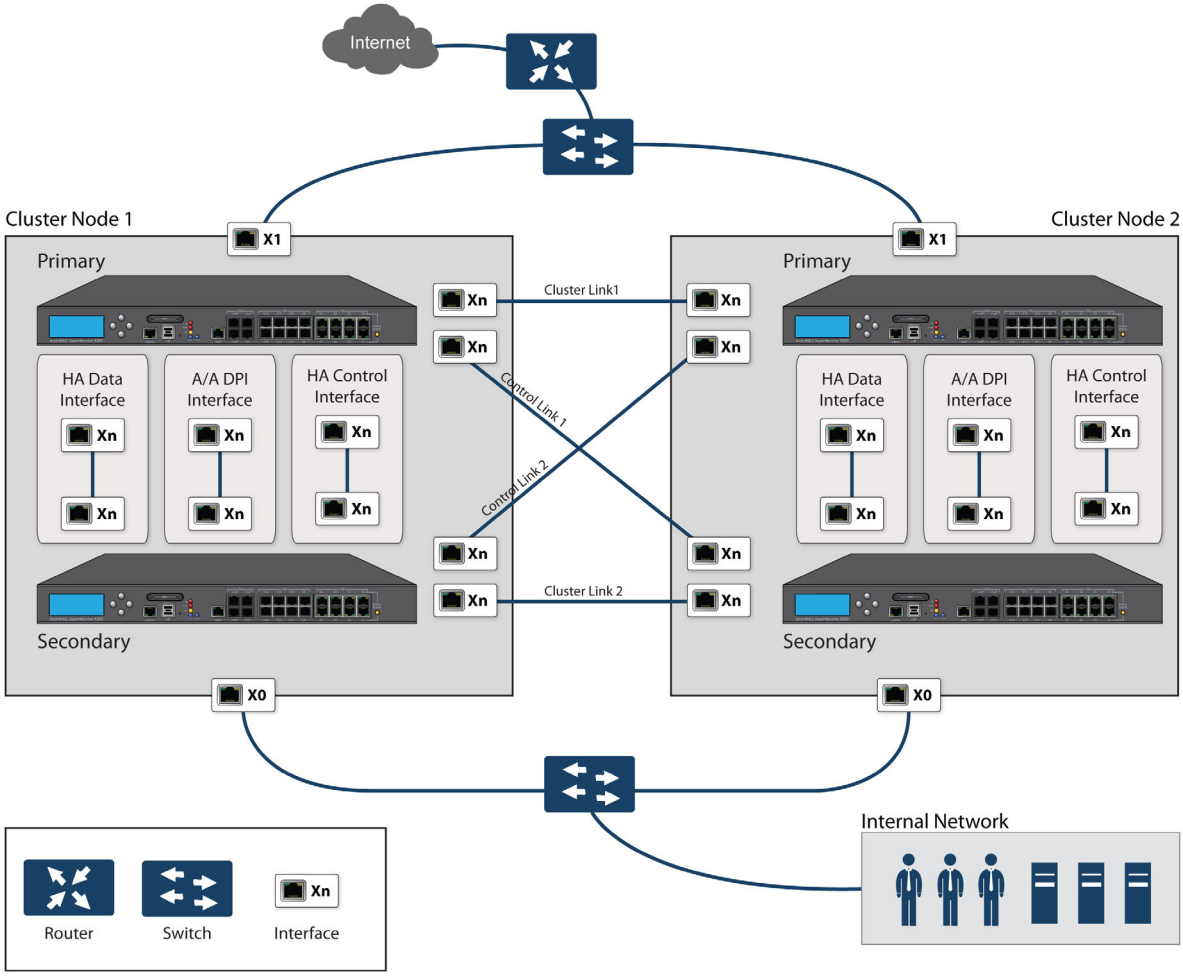
NOTE: Because all Cluster Nodes share the same configuration, each node must have the same redundant ports configured and connected to the same switch(es).

To use Active/Active Clustering, you must register all SonicWall Security Appliances in the cluster on MySonicWall. The two Security Appliances in *each* HA pair must also be associated as HA Primary and HA Secondary on MySonicWall. That is, associate the two Security Appliances in the HA pair for Cluster Node 1, then associate the Security Appliances in the HA pair for Cluster Node 2, and so on for any other Cluster Nodes.

Active/Active DPI Clustering High Availability

Active/Active DPI Clustering High Availability allows for the configuration of up to four HA cluster nodes for failover and load sharing, where the nodes load balance the application of Deep Packet Inspection (DPI) security services to network traffic. See Active/Active DPI clustering high availability.

Active/Active DPI clustering high availability



For the Cluster Links and the Control Links, each unit in Cluster Node 1 is connected to each unit in the peer node (Cluster Node 2). For best practice, use the same set of interfaces on each unit in each node. (For example, connect X8 in one unit to X8 in the peer unit, and do the same if you are using X9 or X10.) However, there is no restriction on which ports you use.

47

Configuring High Availability

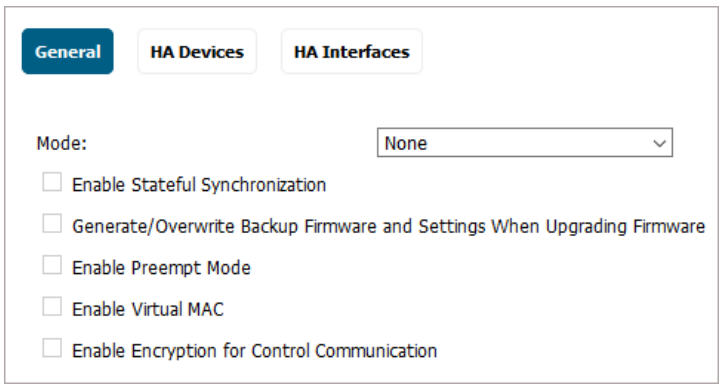
IMPORTANT: High Availability cannot be used along with PortShield except with the SonicWall X-Series/N-Series Solution. Before configuring HA, remove any existing PortShield configuration from **MANAGE | System Setup > Network > PortShield Groups**. For more information, go to https://www.sonicwall.com/support/technical-documentation/ and search for the ***SonicWall SonicOS 6.5 X-Series/N-Series Solution Deployment Guide*** by selecting **NS*a* Series**, **NSA Series**, **SuperMassive 9000 Series**, and **TZ Series** in the **Select A Product** field.

TIP: For a description of High Availability in SonicOS, see About High Availability and Active/Active Clustering on page 674.

Topics:

- **High Availability > Base Setup** on page 703
  - Configuring Active/Standby High Availability Settings on page 704
  - Configuring HA with Dynamic WAN Interfaces on page 706
  - Configuring Active/Active DPI High Availability Settings on page 707
  - Configuring Active/Active Clustering on page 709
  - Verifying Active/Active Clustering Configuration on page 714
  - Configuring Network DHCP and Interface Settings on page 716
  - Active/Active Clustering Full-Mesh on page 718

High Availability > Base Setup



You configure High Availability (HA) on **MANAGE | System Setup > High Availability > Base Setup**:

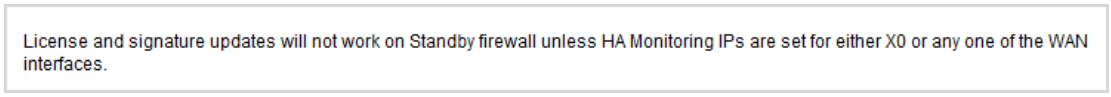
- Configuring Active/Standby High Availability Settings on page 704
- Configuring HA with Dynamic WAN Interfaces on page 706

- Configuring Active/Active DPI High Availability Settings on page 707

NOTE: For more information on High Availability, see About High Availability on page 675 and Active/Standby and Active/Active DPI Prerequisites on page 683. If your Active/Active Clustering environment uses VPN or NAT, see Configuring VPN and NAT with Active/Active Clustering on page 714 after you have finished the Active/Active configuration.

- Configuring Active/Active Clustering on page 709
- Verifying Active/Active Clustering Configuration on page 714
- Configuring Network DHCP and Interface Settings on page 716
- Active/Active Clustering Full-Mesh on page 718

License and signature updates do not work on Standby Security Appliances unless HA Monitoring IP addresses are set for either X0 or any WAN interface. If these interfaces have not been set, a message displays:

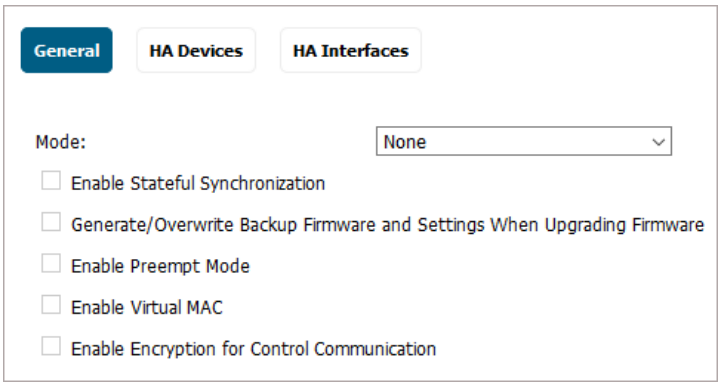


Configuring Active/Standby High Availability Settings

The configuration tasks on **High Availability > Base Setup** are performed on the Primary firewall and then are automatically synchronized to the Secondary firewall.

To configure Active/Standby:

1 Navigate to **System Setup | High Availability > Base Setup**.

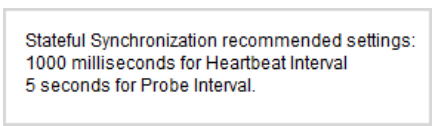


2 From **Mode**, select **Active/Standby**. The options become available.

3 Select **Enable Stateful Synchronization**. This option is not selected by default.

When Stateful High Availability is not enabled, session state is not synchronized between the Primary and Secondary firewalls. If a failover occurs, any session that had been active at the time of failover needs to be renegotiated.

A recommendation message displays.



4 Click **OK**.

5 To back up the settings when you upgrade the firmware version, select **Generate/Overwrite Backup Firmware and Settings When Upgrading Firmware**. This option is not selected by default.

6 To configure the High Availability Pair so that the Primary firewall takes back the Primary role when it restarts after a failure, select **Enable Preempt Mode**. This option is not selected by default.

> **TIP:** It is recommended that preempt mode be disabled when enabling Stateful High Availability because preempt mode can be over-aggressive about failing over to the Secondary firewall.

7 Select **Enable Virtual MAC** to allow the Primary and Secondary firewalls to share a single MAC address. This greatly simplifies the process of updating network ARP tables and caches when a failover occurs. This option is not selected by default.

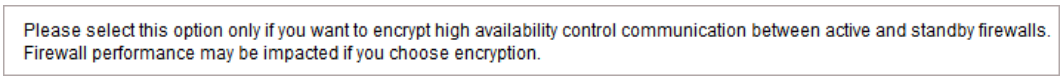
> **IMPORTANT:** If PPPoE Unnumbered is configured, you must select **Enable Virtual MAC**.

Only the switch to which the two firewalls are connected needs to be notified. All outside devices continue to route to the single shared MAC address.

8 To encrypt HA control communication between the active and standby firewalls, select **Enable Encryption for Control Communication**. This option is not selected by default.

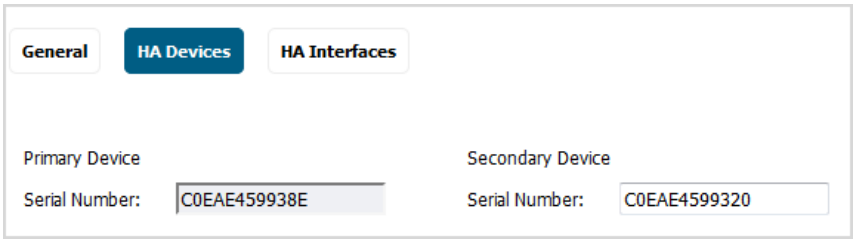
> **IMPORTANT:** Firewall performance may be affected if you choose encryption.

A confirmation message displays:



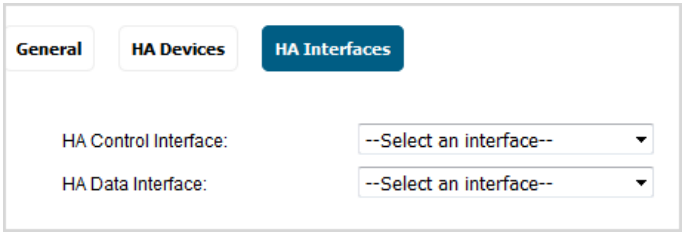
Click **OK**.

9 Click **HA Devices** to configure the Secondary firewall serial number. The serial number for the Primary Device is displayed, but the field is dimmed and cannot be edited.



10 Enter the **Serial Number** of the **Secondary Device**.

11 Click **HA Interfaces**.



12 Select the interface for the **HA Control Interface**. This option is dimmed and the interface displayed if the firewall detects that the interface is already configured.

13 Select the interface for the **Active/Active DPI Interface**. This option is dimmed and the interface displayed out if the firewall detects that the interface is already configured.

14 When finished with all High Availability configuration, click **Accept**. All settings are synchronized to the Secondary firewall, and the Secondary firewall reboots.

Configuring HA with Dynamic WAN Interfaces

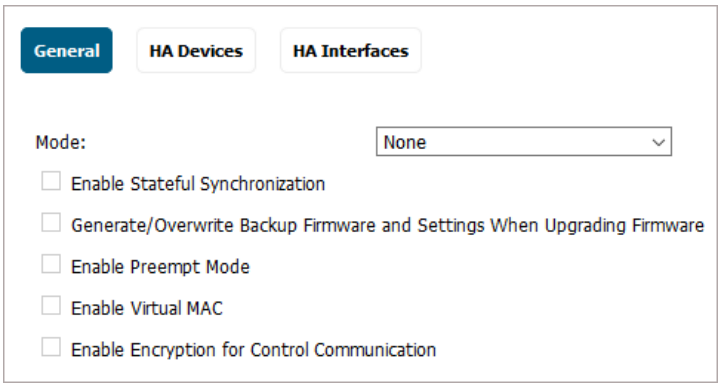
The configuration tasks on **High Availability > Base Setup** are performed on the Primary firewall and then are automatically synchronized to the Secondary.

To configure HA with a dynamic WAN interface:

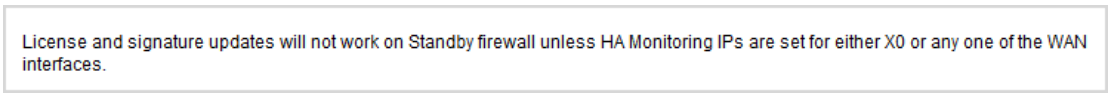
1 Navigate to **MANAGE | System Setup > Network > Interfaces**.

2 Configure a WAN interface as PPPoE, as described in Configuring a WAN Interface on page 292.

3 Navigate to **High Availability > Base Setup**.



4 Select HA mode from **Mode**. If you chose **Active/Active DPI** or **Active/Active Clustering**, a message about license and signature updates displays.



5 Click **OK**.

6 Ensure **Enable Stateful Synchronization** is not selected. This option is not selected by default.

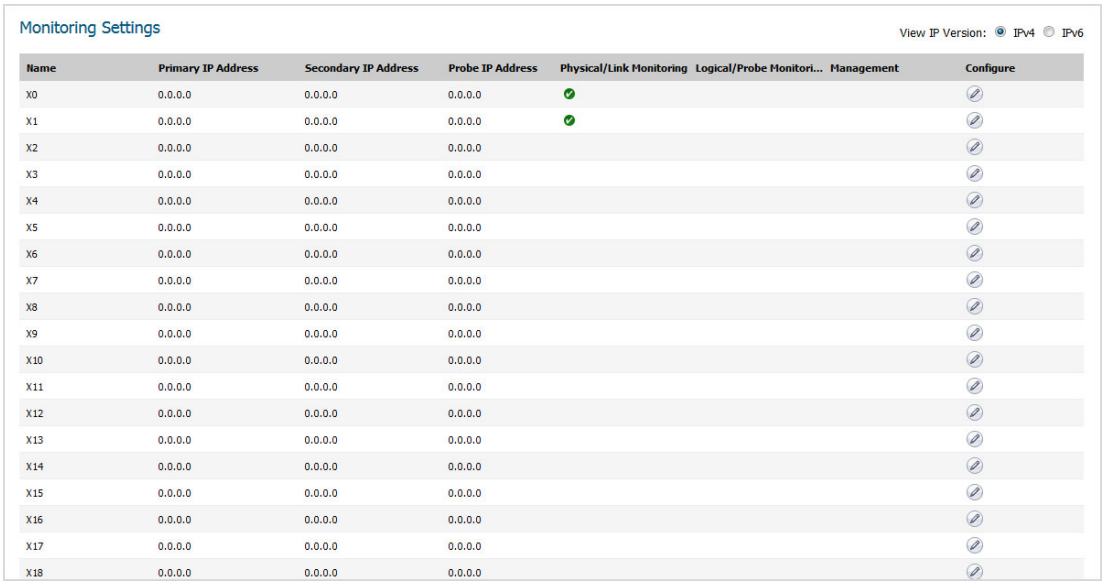
7 Ensure **Enable Preempt Mode** is not selected. This option is not selected by default.

8 Select **Enable Virtual MAC**. This option is not selected by default.

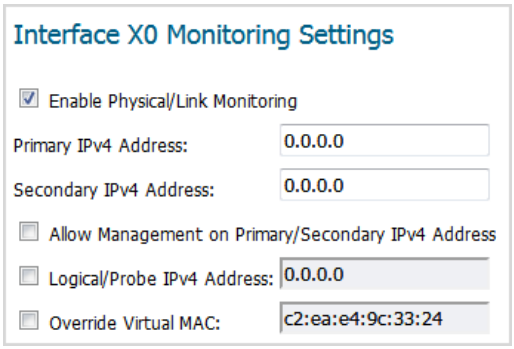
9 Configure **HA Devices** and **HA Interfaces** options as described in Configuring Active/Standby High Availability Settings on page 704.

10 Click **APPLY**.

11 Navigate to **MANAGE | System Setup > High Availability > Monitoring Settings**.



12 Click the **Configure** icon for the PPPoE interface. The **Edit HA Monitoring** dialog displays.



13 Select the **Enable Physical/Link Monitoring** checkbox. This option is not selected by default.

14 Ensure the **Primary IPv4 Address** and **Secondary IPv4 Address** fields are set to `0.0.0.0`.

15 Ensure none of the other options are selected.

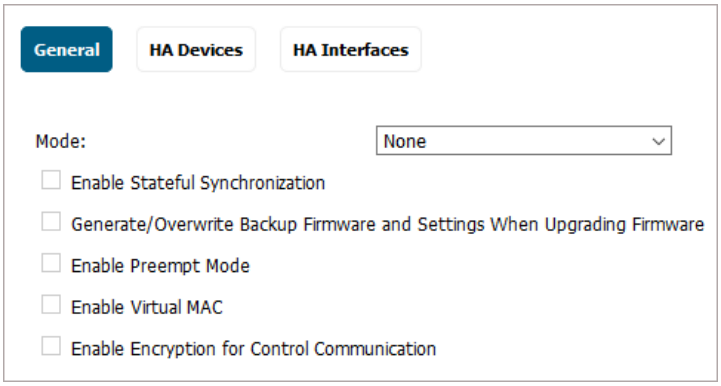
16 Click **OK**.

Configuring Active/Active DPI High Availability Settings

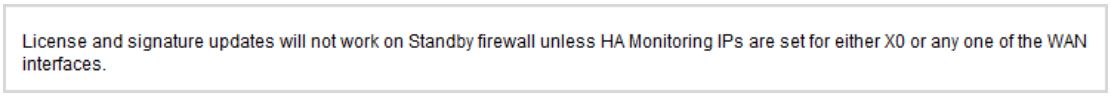
The configuration tasks on **MANAGE | System Setup > High Availability > Base Setup** are performed on the Primary firewall and then are automatically synchronized to the Secondary.

To configure Active/Active DPI:

1 Navigate to **High Availability > Base Setup**.



2 In the **Mode** drop-down menu, select **Active/Active DPI**. A message about license and signature updates displays.



3 Click **OK**.

The **Enable Stateful Synchronization** option is automatically enabled for Active/Active DPI, and the option is dimmed.

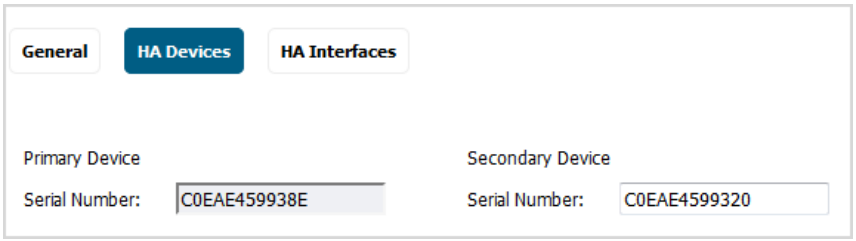
4 To back up the settings when you upgrade the firmware version, select **Generate/Overwrite Backup Firmware and Settings When Upgrading Firmware**. This option is not selected by default.

5 Under normal conditions, Preempt Mode should be disabled for Active/Active DPI. Ensure **Enable Preempt Mode** is not selected. This option is not selected by default.

NOTE: This option instructs the Primary firewall to take back the Primary role when it restarts after a failure; thus, this option only applies to Active/Standby configurations.

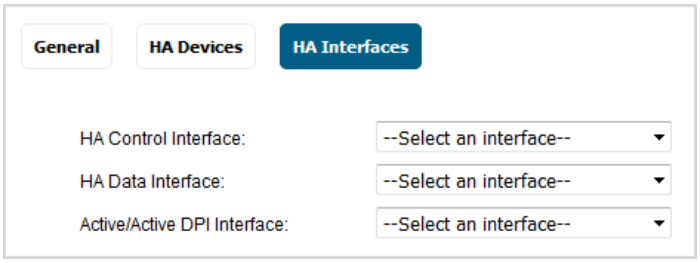
6 To allow both Security Appliances in the HA pair to share a single MAC address, select **Enable Virtual MAC**. This option greatly simplifies the process of updating network ARP tables and caches when a failover occurs. Only the switch to which the two Security Appliances are connected needs to be notified. All outside devices continue to route to the single shared MAC address. This option is not selected by default.

7 Click the **HA Devices** tab. The Serial Number for the Primary Device is displayed, and the field is dimmed and cannot be edited.



8 Enter the **Serial Number** of the **Secondary Device**.

9 Click **HA Interfaces**.



10 Select the HA control interface form **HA Control Interface**. This option is dimmed and the interface displayed if the Security Appliance detects that the interface is already configured.

11 Select the interface number for the **HA Data Interface**. This option is dimmed and the interface displayed if the Security Appliance detects that the interface is already configured.

12 Select the interface number for the **Active/Active DPI Interface**. This option is dimmed and the interface displayed if the Security Appliance detects that the interface is already configured.

This interface is used for transferring data between the two Security Appliances during Active/Active DPI processing. Only unassigned, available interfaces appear in the drop-down menu. The connected interfaces must be the same number on both Security Appliances, and must initially appear as unused, unassigned interfaces in **MANAGE | Network > Interfaces**. For example, you could connect X5 on the Primary unit to X5 on the Secondary if X5 is an unassigned interface. After enabling Active/Active DPI, the connected interface has a **Zone** assignment of **HA Data-Link**.

13 When finished with all High Availability configuration, click **Accept**. All settings are synchronized to the Standby Security Appliance, and the Standby Security Appliance reboots.

Configuring Active/Active Clustering

Topics:

- Configuring Active/Active Clustering High Availability on page 709
- Configuring Active/Active DPI Clustering High Availability on page 711
- Configuring VPN and NAT with Active/Active Clustering on page 714

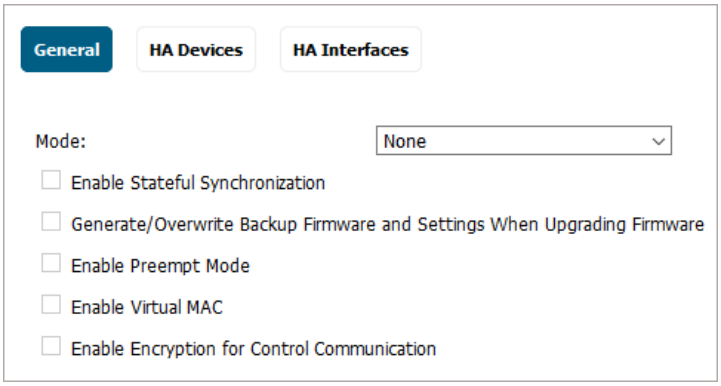
Configuring Active/Active Clustering High Availability

Active/Active Clustering High Availability allows for the configuration of up to four HA cluster nodes for failover and load sharing. Each node can contain either a single Security Appliance or an HA pair.

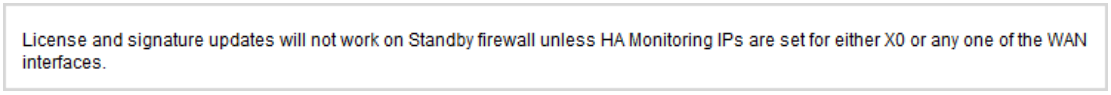
To configure Active/Active Clustering High Availability:

1 Login to the Primary unit of the Master Cluster Node.

2 Navigate to **MANAGE | System Setup > High Availability > Base Setup**.



3 In the **Mode** drop-down menu, select **Active/Active Clustering.** A message about license and signature updates displays.

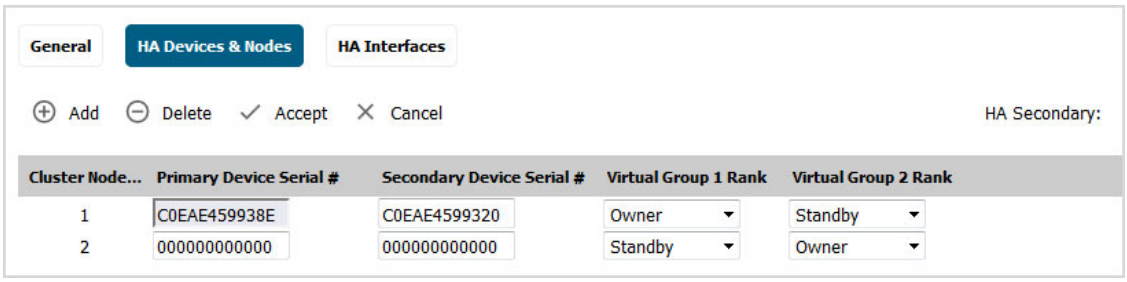


4 Click **OK**. **HA Devices changes to HA Devices & Nodes**.

5 Select **Enable Stateful Synchronization**.

6 To automatically backup the firmware and configuration settings when you upload new firmware to the Security Appliance, select **Generate/Overwrite Backup Firmware and Settings When Upgrading Firmware**. As the Master Node synchronizes new Security Appliance to other Security Appliances in the cluster, secondary units are created on those Security Appliances.

7 To configure the Active/Active cluster information, click **HA Devices & Nodes**.



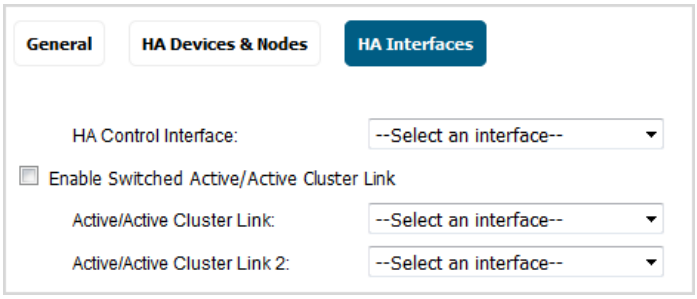
8 In the **Cluster Node** table, enter the serial numbers of the Security Appliances in each Cluster Node in the appropriate **Primary Device Serial #** and **Secondary Device Serial #** fields.

9 Select the rank that Cluster Node 1 holds for each Virtual Group from the **Virtual Group *n* Rank** drop-down menus. By default, Cluster Node 1 is the **Owner** of Group 1 and typically is ranked as **Standby** for other Groups.

To exclude a Security Appliance from a cluster, select **None** for the **Virtual Group *n* Rank**.

10 In the second row, select the rank that Cluster Node 2 holds for each Virtual Group in the **Virtual Group *n* Rank** drop-down menus.

11 Click **HA Interfaces**.



12 Select the HA control interface form **HA Control Interface**. This option is dimmed and the interface displayed if the Security Appliance detects that the interface is already configured.

13 Select **Enable Switched Active/Active Cluster Link**. The options change.



14 Select the interface to be used for transferring data between the two units during Active/Active processing from **Active/Active Cluster Link**. Only unassigned, available interfaces are listed.

15 If you selected **Enable Switched Active/Active Cluster Link**, go to Step 17.

16 Select the interface to be used for transferring data between the two units on the second link during Active/Active processing from **Active/Active Cluster Link 2**. Only unassigned, available interfaces are listed.

17 Click **APPLY**. All settings are synchronized to the Standby unit, and the Standby unit reboots.

18 Go to **MANAGE | System Setup > High Availability > Advanced Settings** and follow the steps in **High Availability > Advanced Settings** on page 725.

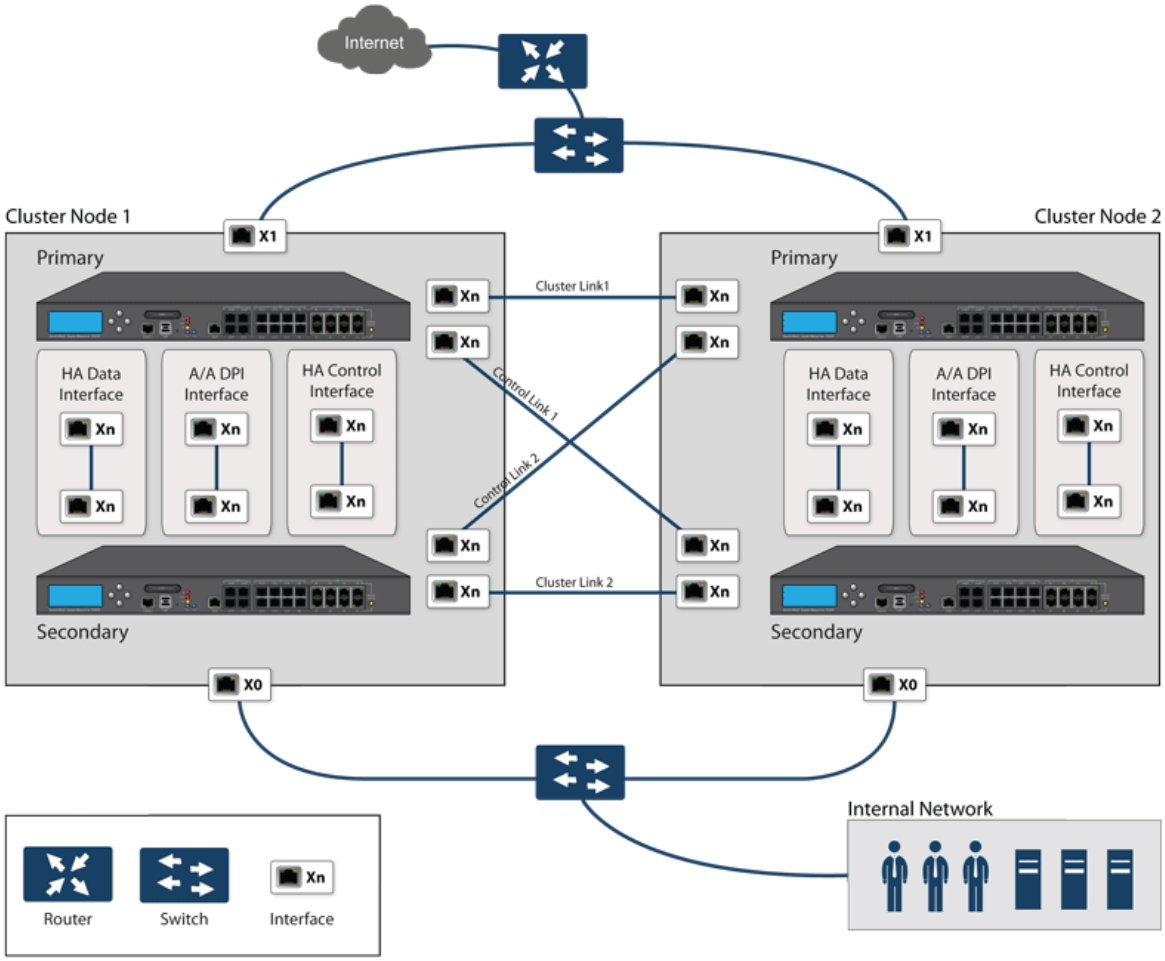
19 Go to **MANAGE | System Setup > Network > Interfaces** page to verify that you have successfully configured the Active/Active interfaces that you want.

20 Go to **MANAGE | System Setup > High Availability > Monitoring Settings** to verify your settings for Active/Active Clustering.

Configuring Active/Active DPI Clustering High Availability

Active/Active DPI Clustering High Availability allows for the configuration of up to four HA cluster nodes for failover and load sharing, where the nodes load balance the application of Deep Packet Inspection (DPI) security services to network traffic. See Active/Active DPI clustering high availability.

Active/Active DPI clustering high availability



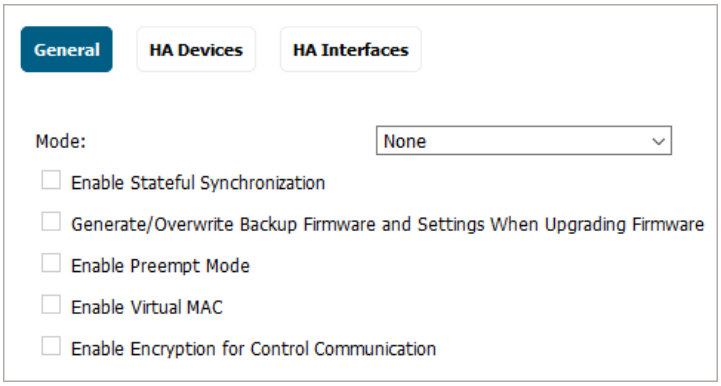
For the Cluster Links and the Control Links, each unit in Cluster Node 1 is connected to each unit in the peer node (Cluster Node 2). For best practice, use the same set of interfaces on each unit in each node. (For example, connect X8 in one unit to X8 in the peer unit, and do the same if you are using X9 or X10.) However, there is no restriction on which ports you use.

To configure Active/Active DPI Clustering High Availability:

NOTE: If you have physically connected the Active/Active DPI Interface as described in Physically Connecting Your Security Appliances on page 685, you are ready to configure Active/Active DPI in the SonicOS management interface.

1 Login to the Primary unit of the Master Cluster Node.

2 Navigate to **MANAGE | System Setup > High Availability > Base Setup**.



3 From **Mode**, select **Active/Active DPI Clustering**.

4 The **Enable Stateful Synchronization** option is automatically enabled for Active/Active DPI Clustering.

5 Select **Generate/Overwrite Secondary Firmware and Settings When Upgrading Firmware** to automatically create a secondary of the firmware and configuration settings when you upload new firmware to the Security Appliance. As the Master Node synchronizes new firmware to other Security Appliances in the cluster, secondarys are created on those Security Appliances.

6 Click **HA Devices** to configure the Active/Active cluster information.

7 For the **HA Secondary** option at the top of the tab, select

- **Internal** if the configured secondary Security Appliance is part of the cluster node for this Security Appliance.
- **External** if the configured secondary Security Appliance is part of a different cluster node.

8 In the table, enter the serial numbers of the Security Appliances in each Cluster Node.

TIP: The serial number for the Primary Device may be populated and dimmed.

9 Enter the rank that Cluster Node 1 holds for each Virtual Group in the **Virtual Group *X* Rank** fields to the right of the serial numbers. By default, Cluster Node 1 is the **Owner** of Group 1, and typically is ranked as **Standby** for Group 2. To exclude an firewall from a cluster, select **None** for the **Virtual Group *X* Rank**.

10 In the second row, enter the rank that Cluster Node 2 holds for each Virtual Group in the **Virtual Group *X* Rank** fields to the right of the serial numbers.

11 Click the **HA Interfaces** tab. Select the interface for the **HA Control Interface**. This option is grayed out if the Security Appliance detects that the interface is already configured.

12 Select the interface for the **Active/Active DPI Interface**. This option is grayed out if the Security Appliance detects that the interface is already configured.

13 Select the **Active/Active DPI Interface**. This interface is used for transferring data between the two units during Active/Active DPI processing. Only unassigned, available interfaces appear in the drop-down menu.

14 Select the **Active/Active Cluster Link** interface.

15 When finished with all High Availability configuration, click **ACCEPT**. All settings are synchronized to the Standby unit, and the Standby unit reboots.

16 Go to **MANAGE | System Setup > High Availability > Advanced Settings** and follow the steps in Fine Tuning High Availability on page 725.

17 Go to **MANAGE | System Setup > Network > Interfaces** to verify that you have successfully configured the Active/Active interfaces that you want.

18 Go to **MONITOR | Current Status > High Availability Status** to verify your settings for Active/Active Clustering. For information about **High Availability Status**, see *SonicOS 6.5 Monitoring*.

Configuring VPN and NAT with Active/Active Clustering

Extra considerations must be taken when configuring these features in an Active/Active Clustering environment:

- Configuring VPN with Active/Active Clustering on page 714
- Configuring a NAT Policy with Active/Active Clustering on page 714

Configuring VPN with Active/Active Clustering

VPN policy configuration requires association with a Virtual Group when running in Active/Active Clustering mode. You configure the options for creating this association on **MANAGE | Connectivity > VPN | Base Settings**. For information about configuring VPN policies, see *SonicOS 6.5 Connectivity*.

Virtual Group address objects are available for local networks. These Virtual Group address objects are created by SonicOS when virtual IP addresses are added and are deleted when the virtual IP is deleted. If creating a VPN Policy for a remote network, Virtual Group address objects may also be available. For example, a custom name **Active-Active-Lan-Host-1**.

Configuring a NAT Policy with Active/Active Clustering

When running in Active/Active Clustering mode, NAT policy configuration includes Virtual Group settings. Default NAT policies are created by SonicOS when virtual IP addresses are added and are deleted when the virtual IP is deleted. You can specify a Virtual Group when creating custom NAT policies; for example, a NAT policy automatically created for Virtual Group 2 on interface X1. For information about creating NAT policies, see *SonicOS 6.5 Policies*.

Verifying Active/Active Clustering Configuration

This section describes several methods of verifying the correct configuration of Active/Active Clustering and Active/Active DPI. See the following:

- Comparing CPU Activity on Firewalls in a Cluster on page 714
- Verifying Settings in MONITOR | Current Status > High Availability Status on page 715
- Additional Parameters in TSR on page 715
- Responses to DPI Matches on page 715
- Logging on page 716

Comparing CPU Activity on Firewalls in a Cluster

When Active/Active DPI is enabled on a Stateful HA pair, you can observe a change in CPU utilization on Security Appliances in the HA pair. CPU activity goes down on the active unit, and goes up on the standby unit.

You can view the CPU utilization on the Multi-Core Monitor. On the active Security Appliance of the Master node, go to **MONITOR | Appliance Health > Live Monitor** and scroll to Multi-Core Monitor to show the activity of all Security Appliances in the Active/Active cluster. For information about the Multi-Core Monitor, see *SonicOS 6.5 Monitoring*.

When viewing the Multi-Core Monitor on an active unit in the cluster, all fSecurity Appliances in the cluster are displayed. However, if you log into the individual IP address of an standby unit in the cluster, the Multi-Core Monitor page only displays the core usage for the two Security Appliances in that particular HA pair.

NOTE: To see the core usage for all Security Appliances in the cluster, SonicWall recommends viewing the Multi-Core Monitor on the active unit of the Master node.

Verifying Settings in MONITOR | Current Status > High Availability Status

In the **Active/Active Clustering Node Status** table, **MONITOR | Current Status > High Availability Status** provides status for the entire Active/Active cluster and for each Cluster Node in the deployment. For information about viewing HA status, see *SonicOS 6.5 Monitoring*.

Additional Parameters in TSR

You can tell that Active/Active DPI is correctly configured on your Stateful HA pair by generating a Tech Support Report on I**NVESTIGATE | Tools > System Diagnostics**. These configuration parameters should appear with their correct values in the Tech Support Report:

- Enable Active/Active DPI
- Active/Active DPI Interface configuration

For information about generating a TSR, see *SonicOS 6.5 Investigate*.

To generate a TSR for this purpose:

1. Log into the Stateful HA pair using the shared IP address.
2. Navigate to I**NVESTIGATE | Tools > System Diagnostics**.
3. Under **Tech Support Report**, click **Download Report**.

Responses to DPI Matches

Responses, or actions, are always sent out from the active unit of the Stateful HA pair running Active/Active DPI when DPI matches are found in network traffic.

NOTE: This does not indicate that all the processing was performed on the active unit.

Deep Packet Inspection discovers network traffic that matches IPS signatures, virus attachments, App Rules policies, and other malware. When a match is made, SonicOS performs an action such as dropping the packet or resetting the TCP connection.

Some DPI match actions inject additional TCP packets into the existing stream. For example, when an SMTP session carries a virus attachment, SonicOS sends the SMTP client a `552` error response code, with a message saying `the email attachment contains a virus`. A TCP reset follows the error response code and the connection is terminated.

These additional TCP packets are generated as a result of the DPI processing on the standby Security Appliance. The generated packets are sent to the active Security Appliance over the Active/Active DPI Interface, and are sent out from the active Security Appliance as if the processing occurred on the active Security Appliance. This ensures seamless operation and it appears as if the DPI processing was done on the active Security Appliance.

Logging

If Active/Active DPI is enabled and DPI processing on the standby Security Appliance results in a DPI match action as described above, then the action is logged on the active unit of the Stateful HA pair, rather than on the standby unit where the match action was detected. This does not indicate that all the processing was performed on the active unit.

High Availability related log events can be viewed in **INVESTIGATE | Logs > Event Logs**. For information about logs, see *SonicOS 6.5 Investigate*.

Configuring Network DHCP and Interface Settings

When Active/Active Clustering is enabled, the SonicOS internal DHCP server is turned off and cannot be enabled. Networks needing a DHCP server can use an external DHCP server. The SonicOS DHCP server should be disabled in the management interface before enabling Active/Active Clustering, and all DHCP server lease scopes deleted.

On **MANAGE | System Setup > Network > Interfaces**, you can configure additional virtual IP addresses for interfaces in a Virtual Group, and redundant ports for interfaces.

For information about performing these tasks, see:

- Disabling the SonicOS DHCP Server on page 716
- Configuring Virtual IP Addresses on page 716
- Configuring Redundant Ports on page 717

Disabling the SonicOS DHCP Server

To disable the SonicOS DHCP server and delete all DHCP server lease scopes:

1. Login to the Primary unit of the Cluster Node.
2. Navigate to the **MANAGE | System Setup > Network > DHCP Server**.
3. Choose IP version: **IPv4** or **IPv6**.
4. Clear **Enable DHCPv4/6 Server**.
5. Under **DHCPv4/6 Server Lease Scopes**, select **All** for the **View Style** to select all lease scopes in the table.
6. Select the checkbox in the header of the table.
7. Click **DELETE**.
8. Click **OK** in the confirmation dialog.
9. Click **ACCEPT**.

Configuring Virtual IP Addresses

When Active/Active Clustering is enabled for the first time, the configured IP addresses for the interfaces on that Security Appliance are automatically converted to virtual IP addresses for Virtual Group 1. Thus, Virtual Group 1 includes virtual IP addresses for X0, X1, and any other interfaces which are configured and assigned to a zone.

Active/Active Clustering requires additional configuration of virtual IP addresses for additional Virtual Groups. You can assign multiple virtual IP addresses to each interface, one per Virtual Group. Each additional virtual IP

address is associated with one of the other Virtual Groups in the cluster. Each interface can have up to a maximum of four virtual IP addresses. VLAN interfaces can also have up to four virtual IP addresses.

NOTE: A packet cannot be forwarded on an interface if a virtual IP address is not configured on it for the Virtual Group handling that traffic flow.

To configure a virtual IP address on an interface:

1 Login to the Primary unit of the Cluster Node.

2 Navigate to **MANAGE | System Setup > Network > Interfaces**.

3 In the **Interface Settings** table, click the **Configure** icon for the interface you want to configure.

4 In the **Edit Interface** dialog, type the virtual IP address into the **IP Address (Virtual Group X)** field, where *X* is the virtual group number.

NOTE: The new virtual IP address must be in the same subnet as any existing virtual IP address for that interface.

5 Click **OK**. The configured virtual IP address appears in the **Interface Settings** table.

Configuring Redundant Ports

Redundant ports can be used along with Active/Active Clustering. You can assign an unused physical interface as a redundant port to a configured physical interface called the “primary interface”. If there is a physical link failure on the primary interface, the redundant interface can continue processing traffic without any interruption. One advantage of this feature is that in case of a physical link failure, there is no need to do a device failover.

You can configure a redundant port on **MANAGE | System Setup > Network > Interfaces > Edit Interface > Advanced** dialog. The **Redundant Port** field is only available when Active/Active Clustering is enabled.

NOTE: Because all Cluster Nodes share the same configuration, each node must have the same redundant ports configured and connected to the same switch(es).

To configure a redundant port for an interface:

1 Login to the Primary unit of the Cluster Node.

2 Navigate to **MANAGE | System Setup > Network > Interfaces**.

3 In the **Interface Settings** table, click the **Configure** icon for the primary interface for which you want to create a redundant port. For example, click the **Configure** icon for **X2**. The **Edit Interface** dialog displays.

4 Click **Advanced**.

5 From **Redundant/Aggregate Ports**, select **Port Redundancy**. The options on the dialog change.

6 From **Redundant Port**, select the redundant port. Only unused interfaces are available for selection. For example, select **X4** for the redundant port.

7 Click **3**.

The selected interface is dimmed in the **Interface Settings** table. A note indicates that it is a redundant Port and lists the primary interface. The interface also appears in the **Redundant Port** field in the **Edit Interface** dialog of the primary port.

NOTE: The primary and redundant ports must be physically connected to the same switch, or preferably, to redundant switches in the network.

8 On each Cluster Node, replicate the redundant physical connections using the same interface numbers for primary and redundant ports. All Cluster Nodes share the same configuration as the Master node.

Active/Active Clustering Full-Mesh

Topics:

- Active/Active Clustering Full-Mesh Overview on page 718
- Configuring Active/Active Clustering Full Mesh on page 720
- Cabling for Active/Active Full Mesh on page 721
- Configuring Active/Active Cluster Security Appliances on page 722
- Configuring Active/Active Cluster Full-Mesh 2-Unit Deployment on page 724

Active/Active Clustering Full-Mesh Overview

Active/Active Clustering Full-Mesh configuration is an enhancement to the Active/Active Clustering configuration option and prevents any single point of failure in the network. All firewall and other network devices are partnered for complete redundancy. Full-Mesh ensures that there is no single point of failure in your deployment, whether it is a device (Security Appliance/switch/router) or a link. Every device is wired twice to the connected devices. Active/Active Clustering with Full-Mesh provides the highest level of availability possible with high performance.

NOTE: The routers in the Security Appliance's upstream network should be pre-configured for Virtual Router Redundancy Protocol (VRRP).

Topics:

- About Full Mesh Deployments on page 718
- Benefits of Active/Active Clustering Full Mesh on page 718
- Redundant Ports and Redundant Switches on page 719

About Full Mesh Deployments

Active/Active Clustering Full Mesh configuration is an enhancement to the Active/Active Clustering configuration option and provides the highest level of availability possible with high performance. Full Mesh deployments provide a very high level of availability for the network, because all devices have one or more redundant partners, including routers, switches, and Security Appliances. Every device is wired twice to the connected devices, so that no single point of failure exists in the entire network. For example, every SonicWall firewall uses redundant ports to connect twice to each networking device.

NOTE: Full Mesh deployments require that Port Redundancy is enabled and implemented.

Benefits of Active/Active Clustering Full Mesh

- **No Single Point of Failure in the Core Network**: In an Active/Active Clustering Full-Mesh deployment, there is no single point of failure in the entire core network, not just for the Security Appliances. An alternative path for a traffic flow is always available in case there are simultaneous failures of switch, router, Security Appliance on a path, thus providing the highest levels of availability.
- **Port Redundancy:** Active/Active Clustering Full-Mesh utilizes port redundancy in addition to HA redundancy within each Cluster Node, and node level redundancy within the cluster. With port redundancy, a backup link takes over in a transparent manner if the primary port fails. This prevents the need for device level failover.

Redundant Ports and Redundant Switches

Redundant ports can be used along with Active/Active Clustering. If one port should have a fault, the traffic is seamlessly handled through the redundant port without causing an HA or Active/Active failover. A **Redundant Port** field in **MANAGE | System Setup > Network > Interfaces > Edit Interface** dialog becomes available when Active/Active Clustering is enabled.

When configuring a redundant port, the interface must be unused; that is, not assigned to any zone. The two ports must be physically connected to the same switch, or preferably, to redundant switches in the network.

NOTE: Because all Cluster Nodes shares the same configuration, each node must have the same redundant ports configured and connected to the same switch(es).

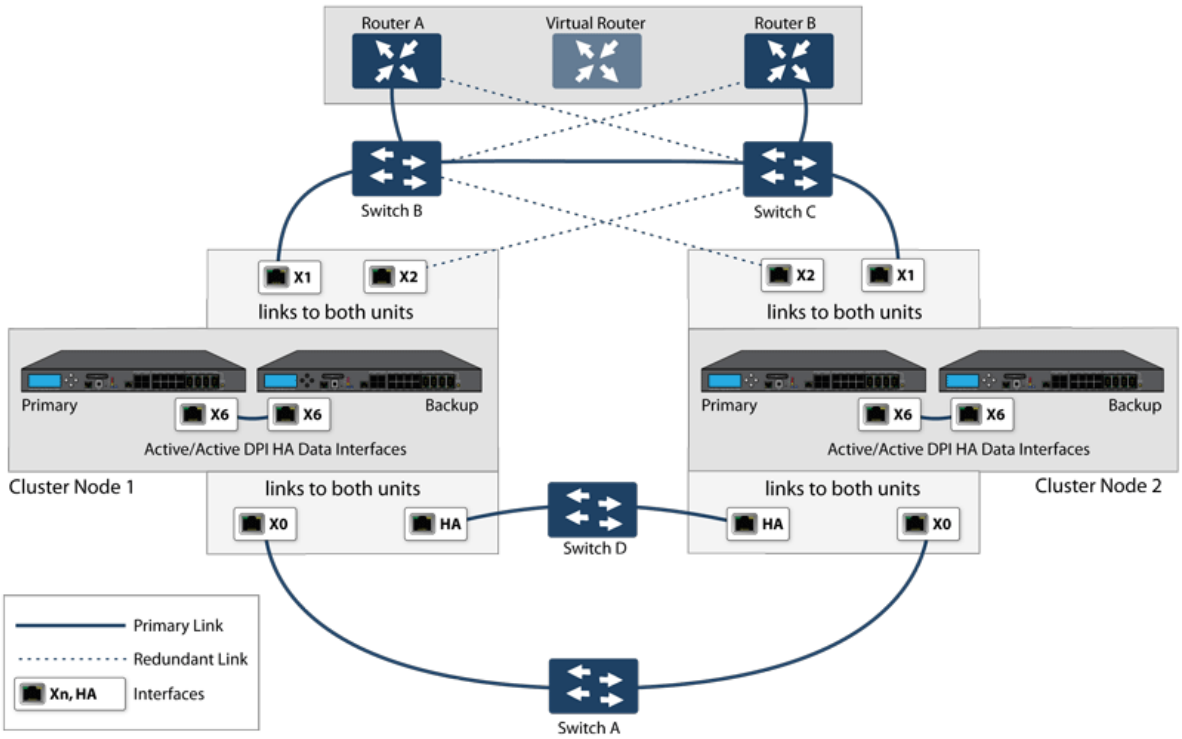
While all Cluster Nodes are up and processing traffic normally, redundant ports remain standby and are ready for use if the partner port goes down for any reason. If one Cluster Node goes down, causing an Active/Active failover, the redundant port on the remaining Cluster Node is put to use immediately to handle the traffic for the Virtual Group that was owned by the failed node. This provides load sharing.

For example, a deployment in which Virtual Group 1 is owned by Cluster Node 1 and Virtual Group 2 is owned by Cluster Node 2. The Cluster Nodes are configured with redundant ports, X3 and X4. No traffic is sent on X4 while all nodes are functioning properly. If Cluster Node 2 goes down, Virtual Group 2 is now also owned by Cluster Node 1. At this point, the redundant port X4 begins to be used for load sharing. Virtual Group 1 traffic is sent on X3, while Virtual Group 2 traffic is sent on X4. In a larger deployment, if Cluster Node 1 owns three or four Virtual Groups, traffic is distributed among the redundant ports – traffic for Virtual Groups 1 & 3 is sent on X3, while traffic for Virtual Groups 2 & 4 is sent on X4.

When a redundant switch is configured, SonicWall recommends using a redundant port to connect to it. While it is possible to connect a redundant switch without using a redundant port, this involves complex configuration using probes. A redundant switch can be deployed anywhere in the network depending on the need for high availability. For example, a redundant switch might be deployed on the WAN side if traffic passing through it is business-critical.

WAN-side redundancy shows a deployment that includes redundant routers, switches, and ports on the WAN side, but is not a Full Mesh deployment because the LAN side does not use redundancy.

WAN-side redundancy



Full Mesh is not required when deploying redundant ports or switches, but a Full Mesh deployment includes them. A Full Mesh deployment uses redundant ports on each of the main traffic ports (LAN, WAN, etc.), and uses redundant upstream routers in addition to redundant switches.

Configuring Active/Active Clustering Full Mesh

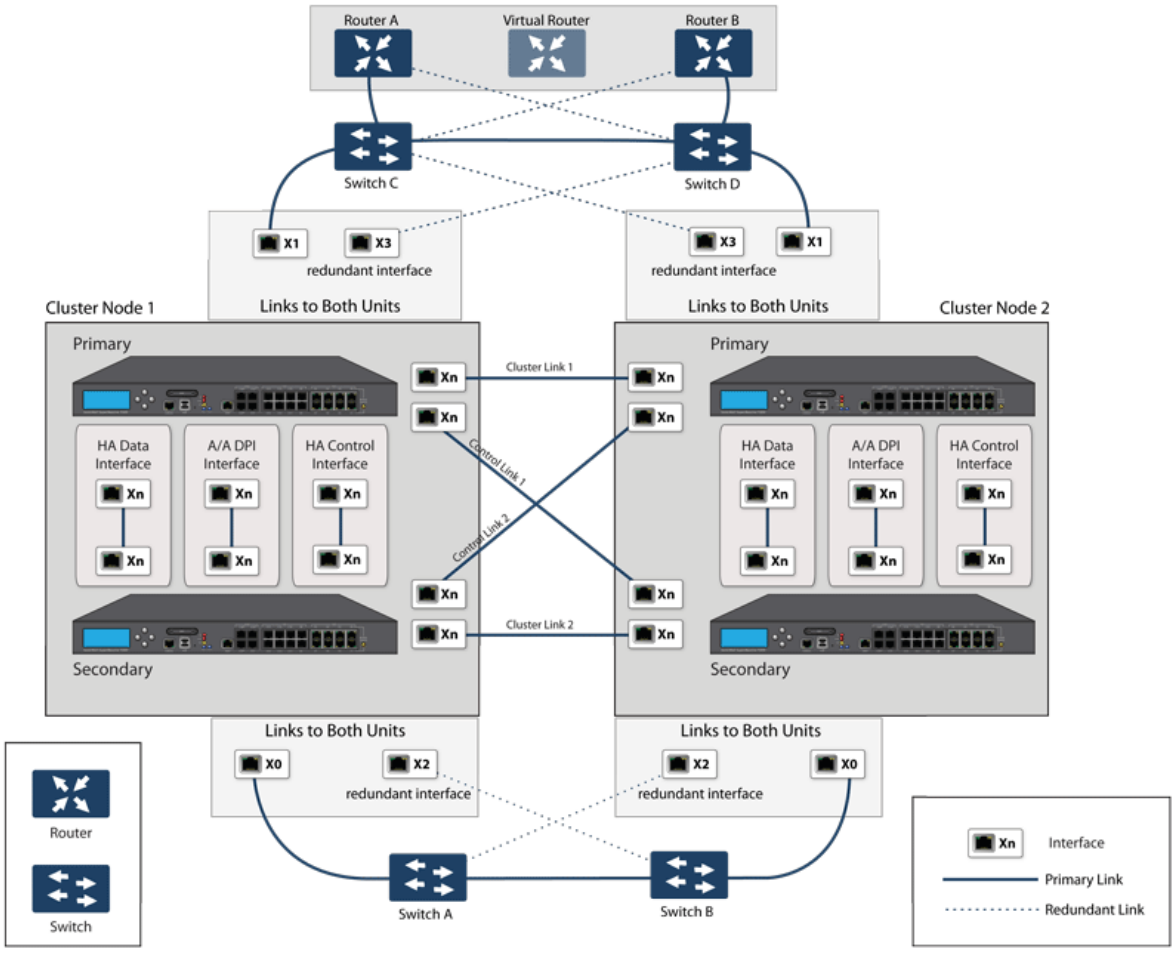
This section describes the procedure for setting up a 4-unit Active/Active Cluster Full-Mesh deployment (see Active/Active four-unit cluster full mesh):

- Cabling for Active/Active Full Mesh on page 721
- Configuring Active/Active Cluster Security Appliances on page 722
- Configuring Active/Active Cluster Full-Mesh 2-Unit Deployment on page 724

The deployments described are examples. Your actual deployment might differ based on the following factors:

- Topology/design of your network and the types of network devices you use (such as switches, routers, load balancers)
- Level of availability desired
- Resource constraints

Active/Active four-unit cluster full mesh



Cabling for Active/Active Full Mesh

This procedure describes the cabling for the deployment illustrated in Active/Active four-unit cluster full mesh.

To physically connect your network devices for a full-mesh deployment:

1 Connect all the HA links of all the firewalls into a port-based VLAN on Switch E.

2 In this setup, X2 is the redundant port of X0. Connect the cables as follows for the X0, X2 ports:

   a Connect CN2-Primary Firewall's X0 to Switch A and X2 to Switch B.

   b Connect CN2-Backup Firewall's X0 to Switch A and X2 to Switch B.

   c Connect CN2-Primary Firewall's X0 to Switch B and X2 to Switch A.

   d Connect CN2-Backup Firewall's X0 to Switch B and X2 to Switch A.

3 On Switch A and Switch B:

   a Configure all the Switch ports connected to the X0,X2 interfaces to be in the same port-based VLAN.

   b Enable Spanning Tree, but also enable Port Fast (or equivalent command) on the ports connected to the firewalls.

4 X3 is the redundant port of X1. Connect the cables as follows for the X1, X3 ports:

a Connect CN2-Primary Firewall's X1 to Switch C and X3 to Switch D.

b Connect CN2-Backup Firewall's X1 to Switch C and X3 to Switch D.

c Connect CN2-Primary Firewall's X1 to Switch D and X3 to Switch C.

d Connect CN2-Backup Firewall's X1 to Switch D and X3 to Switch C.

5 On Switch C and Switch D:

a Configure all the Switch ports connected to the X1,X3 interfaces to be in the same port-based VLAN.

b Enable Spanning Tree, but also enable Port Fast (or equivalent command) on the ports connected to the firewalls.

6 Cable Switch A and Switch B together.

7 Cable Switch C and Switch D together.

8 If the Router A and Router B have redundant port support, then connect the Routers to Switches in the same way as we connected the Firewall ports to Switches. That is, connect the primary port on Router A to Switch C and the backup port on Router A to Switch D. Connect the ports in the same way for Router B.

9 If the Routers do not have redundant port support, but have switching support, then you create two ports in the same VLAN on Router A and assign an IP address to the VLAN instead of the port. Then connect one port to Switch C and the other port to Switch D. Do a similar configuration for Router B. (This is the setup shown in Active/Active four-unit cluster full mesh.)

10 Active/Active DPI is used along with Active/Active Clustering. Ports X6 and X7 are the two HA data ports for redundancy and load-sharing of offloaded traffic from Active to Standby Security Appliance. Perform the following cabling (X6,X7 ports and cabling have not been shown in Active/Active four-unit cluster full mesh for brevity):

a Connect X6 of CN1-Primary to X6 of CN1-Backup with a Cross-over cable.

b Connect X7 of CN1-Primary to X7 of CN1-Backup with a Cross-over cable.

c Connect X6 of CN2-Primary to X6 of CN2-Backup with a Cross-over cable.

d Connect X7 of CN2-Primary to X7 of CN2-Backup with a Cross-over cable.

Configuring Active/Active Cluster Security Appliances

Topics:

- Configuration Procedure on page 722
- Testing for No Point of Failure on page 723

Configuration Procedure

To configure the Active/Active Cluster Security Appliances:

1 Shut down all firewalls except the CN1-Primary unit.

2 On **MANAGE | System Setup > High Availability > Base Setup** page:

a Choose **Active/Active Clustering** from **Mode**.

b Select **Enable Stateful Synchronization**.

c Click **HA Devices & Nodes**.

d Enter the serial numbers of the Cluster Node Primary and Secondary devices in the appropriate **Primary Device Serial #** and **Secondary Device Serial #** fields.

e For CN1, select **Owner** from **Virtual Group 1 Rank** and **Standby** for **Virtual Group 2 Rank**.

f For CN2, select **Owner** from **Virtual Group 1 Rank** and **Standby** for **Virtual Group 2 Rank**.

g Enable Active/Active DPI with X6 and X7 as the two HA data ports.

h Click **Apply**.

3 On **MANAGE | System Setup > Network > Interfaces**:

a Add the Virtual Group (VG) IP addresses for both the X0 and X1 interfaces.

b Add the redundant port configuration (X2 as redundant port of X0, X3 as redundant port of X1).

4 On **MANAGE | System Setup > High Availability > Monitoring Settings**, add the monitoring/management IP addresses either on X0 or X1 for each unit in the cluster.

5 Turn on all the other Security Appliances. A complete synchronization of the configuration is made from the CN1-Primary to all other Security Appliances.

6 Login to each Security Appliance using the dedicated monitoring/management address and do the following:

a Register the Security Appliance on MySonicWall.

b Synchronize the licenses with MySonicWall.

Testing for No Point of Failure

After the above deployment is connected and configured, CN1 owns Virtual Group1 (VG1), and CN2 owns Virtual Group 2 (VG2).

Configure the VG1 IP address on X0 as the gateway for a certain set of traffic flows and the VG2 IP address on X0 as the gateway for other sets of traffic flows. You can use different methods to accomplish this:

- Use a smart DHCP server that distributes the gateway allocation to the PCs on the directly connected client network.
- Using policy-based routes on a downstream router.

When the traffic setup is done, both Cluster Nodes actively process network traffic.

To test for no single point of failure on all devices and links:

1 **Device Failures**: Traffic should continue to flow through both Cluster Nodes in each of these device failures:

a Power down Switch A while Switch B is up and ready.

b Power down Switch B while Switch A is up and ready.

c Restart the Active unit in CN1 from the SonicOS Management Interface while the Standby unit in CN1 is up and ready (this scenario is similar to a software failure on the CN1-Active unit).

NOTE: There is a Stateful HA failover in this case.

d Shut down the CN1-Active unit while the CN1-Standby unit is up and ready (this scenario is similar to a hardware failure on the CN1-Active unit).

NOTE: There is a Stateful HA failover in this case.

e Repeat Step c and Step d for CN2.

f Shut down Router A while Router B is up and ready.

g Shut down Router B while Router A is up and ready.

2 **Link Failures**: Traffic should continue to flow in each of these link failures:

a On each of the Active Security Appliances in the Cluster Node, disconnect the X0 cable while X2 is connected.

b On each of the Active Security Appliances in the Cluster Node, disconnect the X1 cable while X3 is connected.

c Disconnect the primary link from upstream switches to the router which is the Active virtual router.

d Disconnect X6, the Active-Active DPI HA data interface.

Configuring Active/Active Cluster Full-Mesh 2-Unit Deployment

You can deploy Active/Active Cluster Full-Mesh with two Security Appliances, where each CN consists of only one Security Appliance (no HA backup). However, such a setup has these limitations:

- Failover is not stateful and existing connections need to be re-built.
- If the traffic on each unit is greater than 50% of the capacity of the single Security Appliance at the time of failover, then after the failover, the traffic in excess of 50% is dropped.

The procedure for the 2-unit Full-Mesh is similar to the procedure for the 4-unit Full-Mesh, with these exceptions:

- The steps involving the Backup unit in each node do not apply.
- The steps for configuring Stateful Sync and Active-Active DPI do not apply.
- There is no Switch required for connecting the HA ports (as there are only two, they can be directly connected with a cross-over cable).

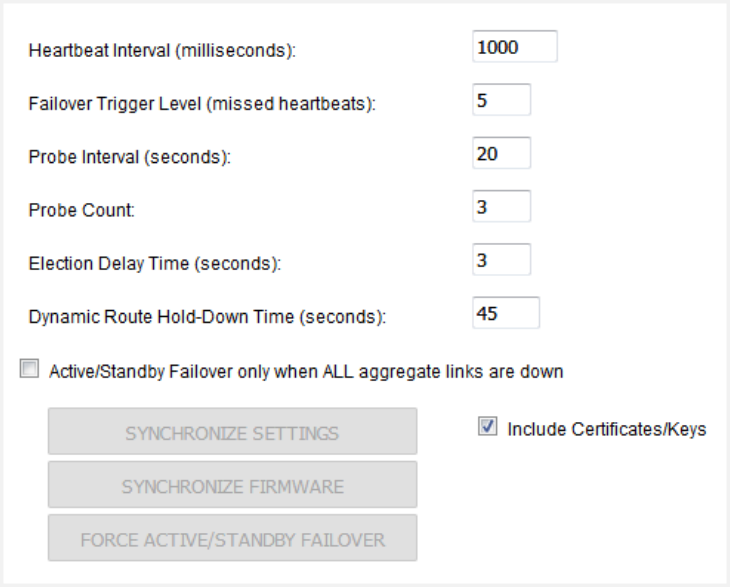
48

Fine Tuning High Availability

Topics:

- **High Availability > Advanced Settings** on page 725
  - Configuring Advanced High Availability on page 726

High Availability > Advanced Settings



MANAGE | System Setup > High Availability > Advanced Settings provides the ability to fine-tune the High Availability configuration as well as synchronize setting and firmware among the High Availability Security Appliances. **High Availability > Advanced Settings** is identical for both Active/Standby and Active/Active configurations.

The **Heartbeat Interval** and **Failover Trigger Level (missed heartbeats)** settings apply to both the SVRRP heartbeats (Active/Active Clustering heartbeat) and HA heartbeats. Other settings on **High Availability > Advanced Settings** apply only to the HA pairs within the Cluster Nodes.

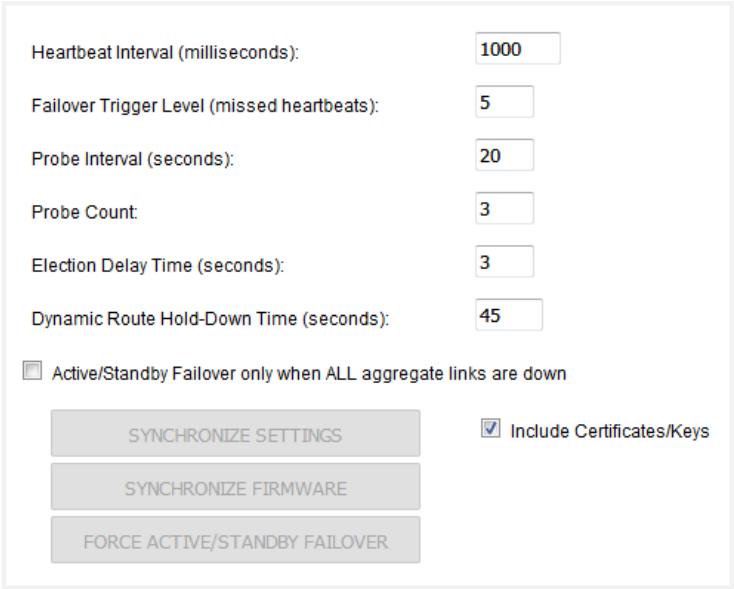
NOTE: For more information on High Availability, see About High Availability on page 675 and Active/Standby and Active/Active DPI Prerequisites on page 683.

Configuring Advanced High Availability

To configure advanced settings:

1 Login as an administrator to the SonicOS Management Interface on the Master Node, that is, on the Virtual Group1 IP address (on X0 or another interface with HTTP management enabled).

2 Navigate to **MANAGE | System Setup > High Availability > Advanced Settings**.



3 Optionally adjust the **Heartbeat Interval** to control how often the Security Appliances in the Active/Active cluster communicate. This setting applies to all units in the Active/Active cluster. The default is **1,000** milliseconds (1 second), the minimum value is 1,000 milliseconds, and the maximum is 300000.

NOTE: SonicWall recommends that you set the Heartbeat Interval to at least 1000.

You can use higher values if your deployment handles a lot of network traffic. Lower values may cause unnecessary failovers, especially when the Security Appliance is under a heavy load.

This timer is linked to the **Failover Trigger Level (missed heartbeats)** timer.

4 Set the **Failover Trigger Level** to the number of heartbeats that can be missed before failing over. This setting applies to all units in the Active/Active cluster. The default is **5**, the minimum is 4, and the maximum is 99.

This timer is linked to the Heartbeat Interval timer. If the **Failover Trigger Level** is set to 5 and the **Heartbeat Interval** is set to 10000 milliseconds (10 seconds), it takes 50 seconds without a heartbeat before a failover is triggered.

5 Set the **Probe Interval** to the interval, in seconds, between probes sent to specified IP addresses to monitor that the network critical path is still reachable. This interval is used in logical monitoring for the local HA pair. The default is **20** seconds, and the allowed range is 5 to 255 seconds.

TIP: SonicWall recommends that you set the interval for at least 5 seconds.

You can set the Probe IP Address(es) on **MANAGE | System Setup > High Availability > Advanced Settings**. See **High Availability > Monitoring Settings** on page 728.

6 Set the **Probe Count** to the number of consecutive probes before SonicOS concludes that the network critical path is unavailable or the probe target is unreachable. This count is used in logical monitoring for the local HA pair. The default is **3**, and the allowed range is 3 to 10.

7 Set the **Election Delay Time** to the number of seconds the Primary Security Appliance waits to consider an interface up and stable. The default is **3** seconds, the minimum is 3 seconds, and the maximum is 255 seconds.

TIP: This timer is useful with switch ports that have a spanning-tree delay set.

8 Set the **Dynamic Route Hold-Down Time** to the number of seconds the newly-active Security Appliance keeps the dynamic routes it had previously learned in its route table. The default value is **45** seconds, the minimum is 0 seconds, and the maximum is 1200 seconds (20 minutes).

NOTE: The **Dynamic Route Hold-Down Time** setting is displayed only when the **Advanced Routing** option is selected on **MANAGE | System Setup > Network > Routing**.

TIP: In large or complex networks, a larger value may improve network stability during a failover

This setting is used when a failover occurs on a High Availability pair that is using either RIP or OSPF dynamic routing. During this time, the newly-active appliance relearns the dynamic routes in the network. When the **Dynamic Route Hold-Down Time** duration expires, SonicOS deletes the old routes and implements the new routes it has learned from RIP or OSPF.

9 If you want Failover to occur only when ALL aggregate links are down, select **Active/Standby Failover only when ALL aggregate links are down**. This option is not selected by default.

10 To have the appliances synchronize all certificates and keys within the HA pair. select **Include Certificates/Keys**. This option is selected by default.

11 (Optional) To synchronize the SonicOS preference settings between your primary and secondary HA firewalls, click **SYNCHRONIZE SETTINGS**.

12 (Optional) To synchronize the firmware version between your primary and secondary HA firewalls, click **SYNCHRONIZE FIRMWARE**.

13 (Optional) To test the HA failover functionality is working properly by attempting an Active/Standby HA failover to the secondary Security Appliance, click **FORCE ACTIVE/STANDBY FAILOVER**.

14 When finished with all High Availability configuration, click **ACCEPT**. All settings are synchronized to the Secondary Security Appliance or to other units in the cluster.

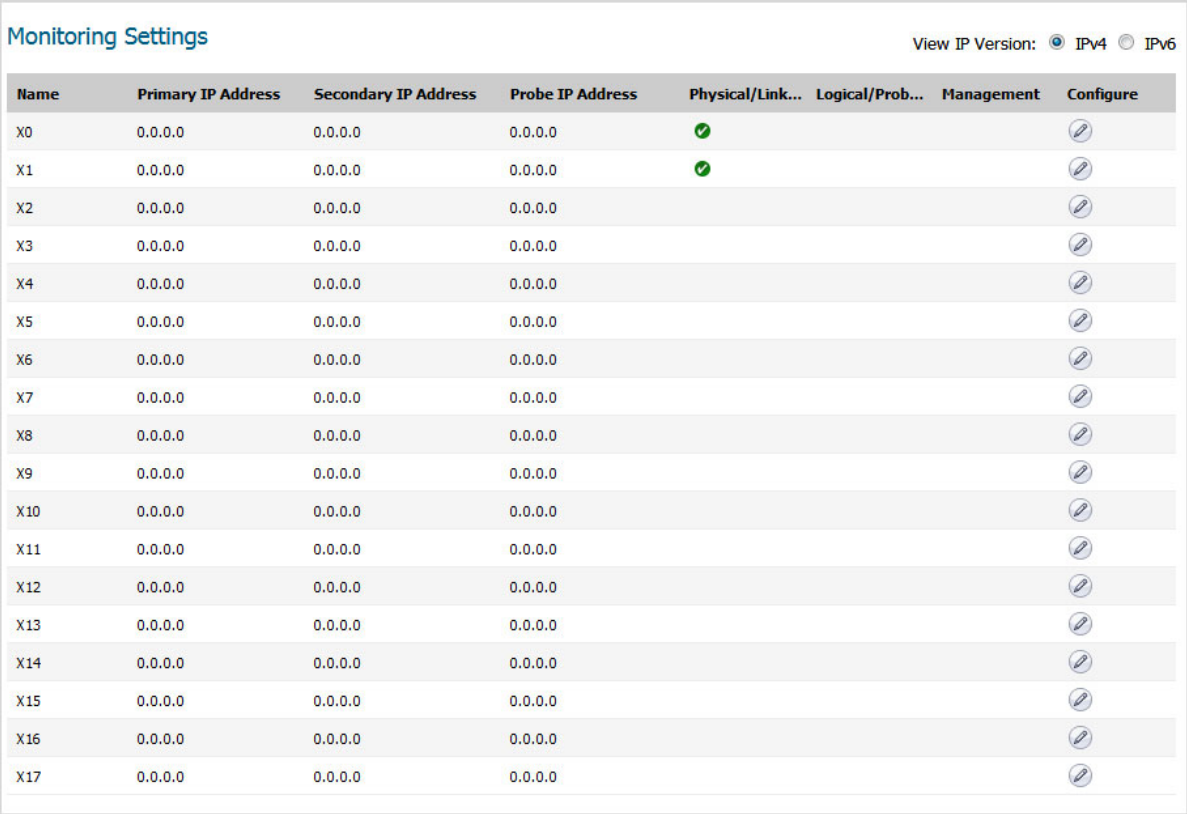
49

Monitoring High Availability

Topics:

- **High Availability > Monitoring Settings** on page 728
  - Configuring Active/Standby High Availability Monitoring on page 729
  - IPv6 High Availability Monitoring on page 730

High Availability > Monitoring Settings



Name	Primary IP Address	Secondary IP Address	Probe IP Address	Physical/Link...	Logical/Prob...	Management	Configure
X0	0.0.0.0	0.0.0.0	0.0.0.0	✔			
X1	0.0.0.0	0.0.0.0	0.0.0.0	✔			
X2	0.0.0.0	0.0.0.0	0.0.0.0				
X3	0.0.0.0	0.0.0.0	0.0.0.0				
X4	0.0.0.0	0.0.0.0	0.0.0.0				
X5	0.0.0.0	0.0.0.0	0.0.0.0				
X6	0.0.0.0	0.0.0.0	0.0.0.0				
X7	0.0.0.0	0.0.0.0	0.0.0.0				
X8	0.0.0.0	0.0.0.0	0.0.0.0				
X9	0.0.0.0	0.0.0.0	0.0.0.0				
X10	0.0.0.0	0.0.0.0	0.0.0.0				
X11	0.0.0.0	0.0.0.0	0.0.0.0				
X12	0.0.0.0	0.0.0.0	0.0.0.0				
X13	0.0.0.0	0.0.0.0	0.0.0.0				
X14	0.0.0.0	0.0.0.0	0.0.0.0				
X15	0.0.0.0	0.0.0.0	0.0.0.0				
X16	0.0.0.0	0.0.0.0	0.0.0.0				
X17	0.0.0.0	0.0.0.0	0.0.0.0				

On **MANAGE | System Setup > High Availability > Monitoring Settings**, you can configure independent management IP addresses for each unit in the HA Pair, using either LAN or WAN interfaces. You can also configure physical/link monitoring and logical/probe monitoring. For more information about the HA Monitoring settings, see About High Availability and Active/Active Clustering on page 674.

Topics:

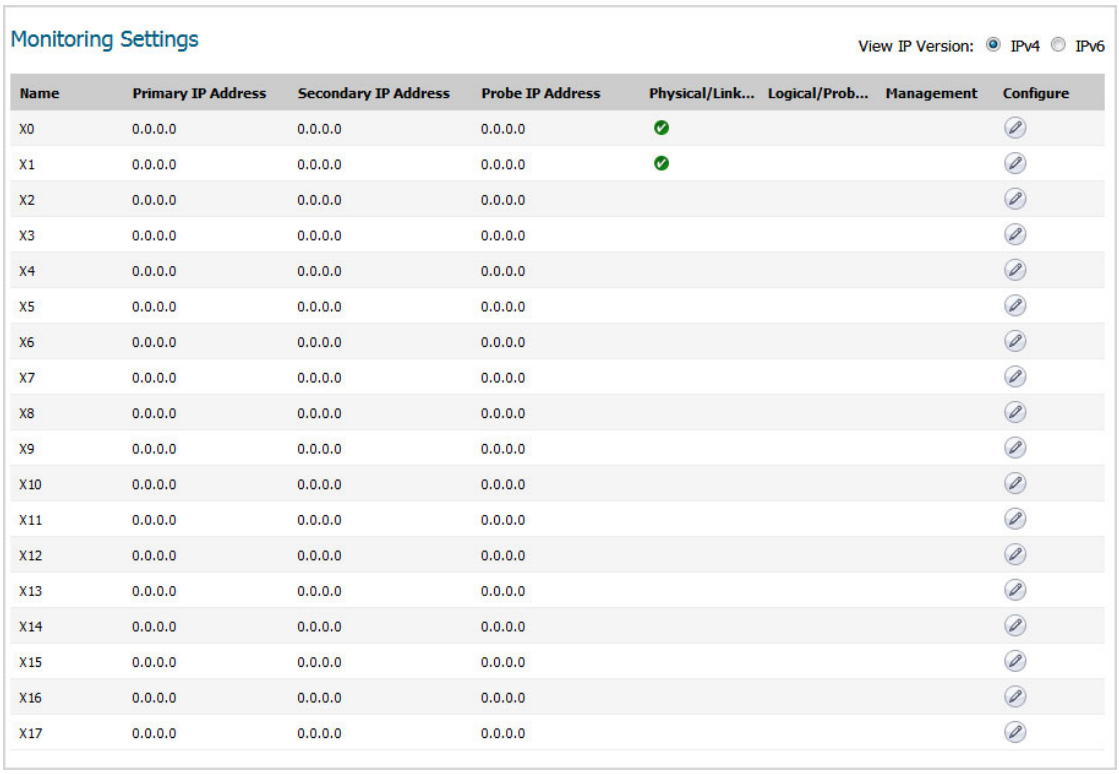
- Configuring Active/Standby High Availability Monitoring on page 729
- IPv6 High Availability Monitoring on page 730

Configuring Active/Standby High Availability Monitoring

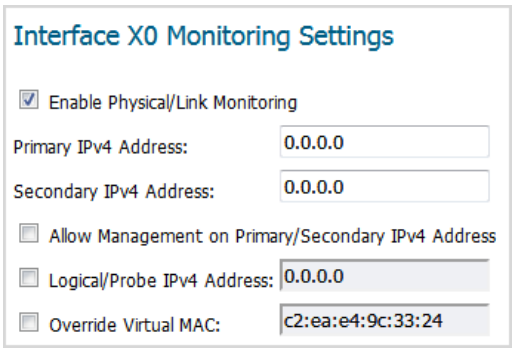
To set the independent LAN management IP addresses and configure physical and/or logical interface monitoring:

1 Login as an administrator to the SonicOS Management Interface on the Primary SonicWall Security Appliance.

2 Navigate to **MANAGE | System Setup > High Availability > Monitoring Settings**.



3 Click the **Configure** icon for an interface on the LAN, such as **X0**. The **Edit HA Monitoring** dialog displays.



4 To enable link detection between the designated HA interfaces on the Primary and Secondary units, leave **Enable Physical Interface Monitoring** selected. This option is selected by default.

5 In the **Primary IPv4/v6 Address** field, enter the unique LAN management IP address of the Primary unit. The default is `0.0.0.0`.

6 In the **Secondary IPv4/v6 Address** field, enter the unique LAN management IP address of the Secondary unit. The default is `0.0.0.0`.

7 Select **Allow Management on Primary/Secondary IP Address**. When this option is enabled for an interface, a green icon appears in the interface's **Management** column in the **Monitoring Settings** table. Management is only allowed on an interface when this option is enabled. This option is not selected by default.

8 In the **Logical Probe IPv4/v6 Address** field, enter the IP address of a downstream device on the LAN network that should be monitored for connectivity. Typically, this should be a downstream router or server. (If probing is desired on the WAN side, an upstream device should be used.) This option is not selected by default.

The Primary and Secondary Security Appliances regularly ping this probe IP address. If both successfully ping the target, no failover occurs. If neither successfully ping the target, no failover occurs, because it is assumed that the problem is with the target, and not the Security Appliances. But, if one Security Appliance can ping the target but the other cannot, failover occurs to the Security Appliance that can ping the target.

The **Primary IPv4/v6 Address** and **Secondary IPv4/v6 Address** fields must be configured with independent IP addresses on a LAN interface, such as X0, (or a WAN interface, such as X1, for probing on the WAN) to allow logical probing to function correctly.

9 Optionally, to manually specify the virtual MAC address for the interface, select **Override Virtual MAC** and enter the MAC address in the field. The format for the MAC address is six pairs of hexadecimal numbers separated by colons, such as `A1:B2:C3:d4:e5:f6`. This option is not selected by default.

IMPORTANT: Care must be taken when choosing the Virtual MAC address to prevent configuration errors.

When **Enable Virtual MAC** is selected on **MANAGE | System Setup | High Availability > Advanced Settings**, the SonicOS firmware automatically generates a Virtual MAC address for all interfaces. Allowing the SonicOS firmware to generate the Virtual MAC address eliminates the possibility of configuration errors and ensures the uniqueness of the Virtual MAC address, which prevents possible conflicts.

10 Click **OK**.

11 To configure monitoring on any of the other interfaces, repeat Step 3 through Step 10 for each interface.

12 When finished with all High Availability configuration, click **ACCEPT.** All settings are synchronized to the Secondary unit automatically.

IPv6 High Availability Monitoring

For complete information on the SonicOS implementation of IPv6, see IPv6 on page 932.

IPv6 High Availability (HA) Monitoring is implemented as an extension of HA Monitoring in IPv4. After configuring HA Monitoring for IPv6, both the primary and backup Security Appliances can be managed from the IPv6 monitoring address, and IPv6 Probing is capable of detecting the network status of HA pairs.

For easy configuration of both IP versions, toggle between IPv6 and IPv4 displays in **MANAGE | System Setup > High Availability > Monitoring Settings**.

The IPv6 HA Monitoring configuration page is inherited from IPv4, so the configuration procedures are almost identical. Just select IPv6 and refer to About High Availability on page 675 and IPv6 HA Monitoring Considerations on page 730 for configuration details.

IPv6 HA Monitoring Considerations

Consider the following when configuring IPv6 HA Monitoring:

- In the **Edit HA Monitoring** dialog, **Enable Physical/Link Monitoring** and **Override Virtual MAC** are dimmed because they are layer 2 properties. That is, the properties are used by both IPv4 and IPv6, so you configure them in the IPv4 monitoring page.
- The primary/backup IPv6 address must be in the same subnet of the interface, and it can not be same as the global IP and Link-Local-IP of the primary/backup Security Appliance.
- If the primary/backup monitoring IP is set to (not **::**), then they cannot be the same.
- If **Allow Management on Primary/Secondary IPv6 Address** is enabled, then primary/backup monitoring IPv6 addresses cannot be unspecified (that is, **::**).
- If **Logical/Probe IPv6 Address** is enabled, then the probe IP cannot be unspecified.

Part 8

System Setup | WAN Acceleration

- Using WAN Acceleration

50

Using WAN Acceleration

Topics:

- About WAN Acceleration on page 733
  - Supported Platforms on page 734
  - Transmission Control Protocol Acceleration on page 734
  - Windows File Sharing Acceleration on page 734
  - Web Cache on page 735
  - Deployment Prerequisites for the WAN Acceleration Service on page 735
  - About WXA Clustering on page 736
  - How Does WXA Clustering Work? on page 738
  - Permitting Acceleration on a Route Policy on page 739
- System Setup > WAN Acceleration on page 739
  - Enabling WAN Acceleration on page 740
  - Managing Groups on page 740
  - Managing WXAs through the WXAs Table on page 744
  - Configuring WXA for VPN Polices on page 760
  - Configuring Acceleration of SSL VPN Traffic on page 760

About WAN Acceleration

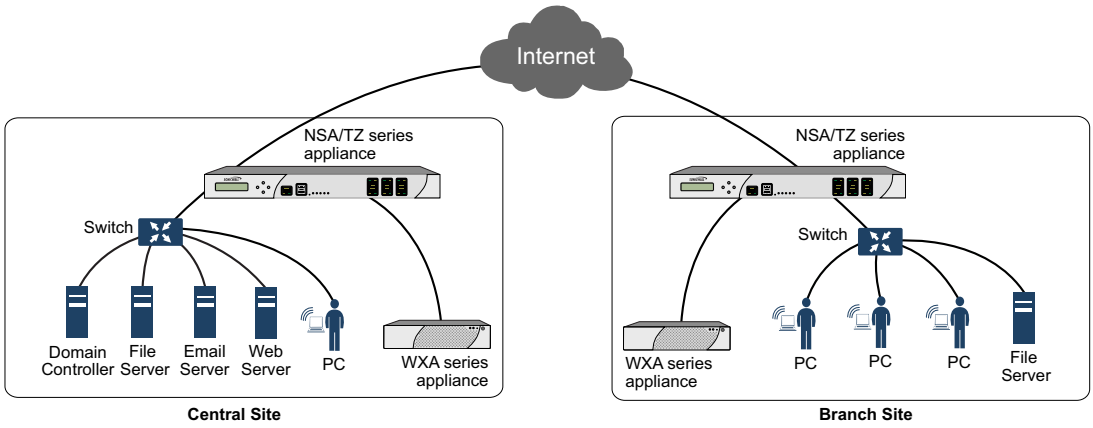
The WAN Acceleration (WXA) service allows you to accelerate WAN traffic between a central site and a branch site by using Transmission Control Protocol (TCP), and Windows File Sharing (WFS). The SonicWall WXA series appliance is deployed in conjunction with a SonicWall NSA series appliance. In this type of deployment, the NSA series appliance provides dynamic security services, such as attack prevention, Virtual Private Network (VPN), routing, and Web Content Filtering. The WAN Acceleration service can increase NSA series appliance performance.

IMPORTANT: If **Wireless-Controller-Only** mode has been selected for **Wireless LAN Controller**, WXA is disabled and a message displays:



SonicWall WXA series appliance topology displays the basic network topology for the SonicWall WXA series appliance and the SonicWall Security Appliances.

SonicWall WXA series appliance topology



Topics:

- Supported Platforms on page 734
- Transmission Control Protocol Acceleration on page 734
- Windows File Sharing Acceleration on page 734
- Web Cache on page 735

Supported Platforms

WAN Acceleration is available for SonicOS on these platforms:

- SuperMassive 9200, 9400, and 9600
- TZ600, TZ500/500 W, TZ400/400 W, TZ350/350 W, TZ300/300 W
- NSA 6600/5600/4600/3600/2650/2600
- SOHO 250/250 W, SOHO W

WXA Clustering is currently only available for the NSA and SuperMassive series of Security Appliances.

Transmission Control Protocol Acceleration

The TCP Acceleration service is a process that decreases the amount of data passing over the WAN by using compression. This accelerates selected traffic passing between a central site and a branch site. The selected traffic is stored in the SonicWall WXA series appliances' shared databases as blocks of data and tagged with reference indexes. This allows the WXA series appliances to only send the reference indexes (which are smaller in size) over the WAN instead of the actual data.

Windows File Sharing Acceleration

WAN Acceleration refers to a wide range of technologies that are aimed at accelerating applications, improving throughput, and reducing latency. Windows File Sharing (WFS) Acceleration is a subset of WAN Acceleration.

The use of WFS Acceleration within your network reduces the impact of high-latency and low-bandwidth links by approximating streaming behavior through the use of read-ahead and write-behind functionality and

differential file transfer to avoid re-transferring parts of files that have not changed. WFS Acceleration allows branch users to access and share commonly used files at near-LAN speeds over the WAN.

Distributed enterprises that deploy WFS Acceleration solutions may be able to consolidate storage to corporate central sites, eliminating the need to back up and manage data that previously resided in their branch sites.

Where storage is not consolidated, the costs and delays of accessing local and branch storage data from other sites is also reduced.

The WXA series appliance offers WFS Acceleration for:

- Unsigned SMB Traffic – In networks that support Unsigned SMB traffic, configuring WFS Acceleration is greatly simplified, because Unsigned SMB traffic does not have a security layer. So, the WXA series appliances can intercept traffic without joining a domain, which eliminates the need to configure custom DNS zones, reverse lookup, and file shares.
- Signed SMB Traffic – In networks that require SMB signing, the WXA series appliances must join a domain, as Signed SMB traffic has a security layer. Signed SMB configuration is more complex than Unsigned SMB configuration and offers more granularity. Signed SMB configuration also has an Advanced configuration mode with more options.

Extended Support for Signed SMB

Extended Support for Signed SMB traffic is handled by a single WXA and is configured outside the groups settings used elsewhere in WXA Clustering. By configuring it on the Windows Domain, users can fully benefit from the extra functionality of the WFS Acceleration module on networks that support signed SMB. After the WXA series appliance has joined the domain, you can configure the shares on the remote servers that you would like to be included in the WFS Acceleration process.

IMPORTANT: It is strongly recommended that you configure the WXA series appliances at the sites where the file servers are located before configuring the WXA series appliances at the branch sites requiring remote access to the shares.

Web Cache

The Web Cache feature stores copies of Web pages and Youtube videos passing through the network that are frequently and recently requested. So, when a user requests one of these Web pages, it is retrieved from the local web cache instead of the Internet, saving bandwidth and response time. Minimal, Moderate, and Aggressive caching strategies are available. These determine which objects are placed into the web cache and how long they stay there.

Deployment Prerequisites for the WAN Acceleration Service

A SonicWall Security Appliance is required to deploy the SonicWall WXA series appliance.

Traffic passing through the SonicWall WXA series appliance requires Internet Protocol version 4 (IPv4). The WAN Acceleration service is not compatible with IPv6.

Deployment Considerations

Consider the following when deploying the SonicWall WXA series appliance:

- WXA Clustering is supported by NSA and SuperMassive series Security Appliances, where multiple WXAs can be plugged into the Security Appliance.
- For WXA Clustering, the WXA series appliances are supported to work with SonicWall NSA 2600 or above and SuperMassive series Security Appliances running SonicOS 6.2.2 or higher firmware.
- WXA 500 can run in memory mode by booting the PC with the CD inserted. Alternatively, it can be installed onto the hard disk. In the latter case, more features are available.
- Typically the WXA series appliances are deployed in a site-to-site VPN configuration through their respective SonicWall Security Appliances. However, you can also use routing or L2 Bridge Mode.
- If a WXA series appliance is used in a high availability configuration, a switched connection to both high availability pairs is required.
- The initial configuration of the WXA series appliance should be performed by using the WXA Setup Wizard, which is available by clicking **Quick Configuration** in SonicWall Security Appliance's Management Interface. For more information on the WXA Setup Wizard, see *SonicOS Quick Configuration*.
- Encrypted traffic is highly randomized and does not materially benefit from the WXA series appliance's WAN Acceleration service. Therefore, SSL and TLS traffic types are not accelerated.
- WFS Acceleration using Signed SMB supports Windows file services using Active Directory, Kerberos, and NTLM for authentication and authorization.
- WFS Acceleration using Signed SMB with NTLM clients provides credentials to the SonicWall WXA series appliance that are valid in the domain. The SonicWall WXA series appliance obtains the Kerberos credentials through the Domain Controller. This permits client devices that have not joined the domain to be used by users who have valid domain credentials.
- Create a DHCP scope on the managing SonicWall Security Appliance before the WXA series appliance is physically connected.
- If the branch offices have Domain Controllers and DNS Servers, it is recommended that you use those DNS Server addresses and the Domain DNS name in the DHCP scope. Configure the only Domain Name and Domain DNS server IP addresses in the configured DHCP scope. The WXA series appliance discovers Kerberos, LDAP, and NTP servers automatically based on this type of information to assist in joining the appliance to the domain.
- Review the LDAP, Kerberos, and NTP services. In a multi-site domain where sites and services are not explicitly configured, the WXA series appliance might not choose the closest servers.
- SonicWall recommends that the WXA series appliance retrieves NTP updates from the Domain Controller. If an NTP server is not configured, this is done automatically.
- SonicWall recommends that the Active Directory DNS zones that hold the WXA name or IP address be configured to accept secure updates only.
- Configure the zone properties of an interface to which the WXA series appliance is connected as a LAN zone.

About WXA Clustering

NOTE: WXA Clustering is supported on NSA 2600 and higher appliances.

SonicOS supports WXA Clustering with two or more NSA series appliances or SuperMassive appliances. The largest SonicWall WXA series appliances can support up to 1200 connections, which roughly translates into support for as many as 240 concurrent users. The number of supported users is now increased with WXA Clustering:

- SonicOS can monitor or probe multiple WXA series appliances simultaneously and store a friendly name for each WXA.

- SonicOS implements three forms of load sharing: TCP Acceleration, Unsigned SMB acceleration, and Web Cache.
- A VPN Policy can be assigned to always use the same WXA group.
- The number of connections is spread equally across all WXA series appliances in a group.
- When one of the WXAs reaches the connection capacity, the next WXA in the group is used.

Topics:

- Supported Platforms for Clustering on page 737
- What is WXA Clustering? on page 737
- How Does WXA Clustering Work? on page 738

Supported Platforms for Clustering

WXA Clustering is supported:

- By WXA firmware versions 1.3.2 and above.
- On these SonicWall Security Appliances:

SM 9600	NSA 6600
SM 9400	NSA 5600
SM 9200	NSA 4600
	NSA 2600
	NSA 3600

What is WXA Clustering?

WXA Clustering is defined as two or more WXA series appliance working in conjunction to provide increased throughput and resilience.

Benefits

Clustering WXA series appliance significantly increases the number of connections that can be accelerated simultaneously. By simply adding more WXA devices, you can increase that capacity several times over. Maximum users and connections by WXA model shows the maximum number of users and connections that are available for each WXA platform.

Maximum users and connections by WXA model

	WXA series appliance				
	WXA 6000	WXA 4000	WXA 2000	WXA 5000	WXA 500 Live
Platform	Software	Hardware			
Appliance	Hardware				
Appliance	Virtual				
Appliance	Software				
Maximum Users	2000	240	120	360	20
Maximum Connections	10,000	1,200	600	1,800	100

Clustering WXA appliances provides these benefits:

- Increases scalability of the acceleration solution for both users and the WAN infrastructure
- Is an elastic solution that can be scaled to meet enterprise and application requirements
- Is a flexible solution that allows one or more WXAs to be dedicated to particular tasks or network segments
- Is a resilient infrastructure for WAN Acceleration

How Does WXA Clustering Work?

WXA Clustering is accomplished by connecting multiple WXA series appliances together and using load balancing and connection balancing to increase the number of simultaneous connections that are possible. It is not necessary to have WXA Clustering implemented at both the remote and local locations, but there must be at least one WXA at each location.

When multiple WXAs are connected and are working together, the amount of data that can be accelerated over the WAN is increased significantly.

In a WXA Clustering configuration, the WXAs are members of groups, and there can be multiple groups. The WXAs in each group have the same configuration, but different groups of WXAs can have different configurations.

The WXA configurations are pushed from SonicOS on the SonicWall Security Appliance.

Topics:

- Restrictions on page 738
- Licensing for WXA Clustering on page 738

Restrictions

WXA Clustering does not support WFS Acceleration for Signed SMB. WFS Acceleration for Signed SMB is supported if you use a single WXA that is dedicated to accelerating Signed SMB. That WXA may or may not be part of a group. However, by keeping it outside of the clustering groups, it can be engaged solely in accelerating Signed SMB traffic.

Licensing for WXA Clustering

Licensing for WXA Clustering is based on the maximum number of concurrent accelerated connections that you wish to support. Customers can purchase a WXA Clustering Licenses for a specified number of connections that are accelerated.

On the WXA 500, WXA 5000 and WXA 6000, WXA Clustering licenses are purchased based on the number of connections desired. Each license represents the maximum number of connections that are allowed. Only the maximum number of connections that are licensed are accelerated. If more than the maximum number of connections passes through the SonicWall Security Appliance, the excess connections are still established, but are not accelerated.

On the WXA 2000 and WXA 4000, no extra licenses are needed. For these models, the maximum number of connections that is built into the appliance is also the maximum number of connections that are accelerated.

If a WXA 2000 or a WXA 4000 is added to a cluster that has a WXA 500, WXA 5000 or WXA 6000 in it, the number of connections is increased accordingly. For example, adding a WXA 2000 to a cluster adds 600 concurrent connections to the allowed limit.

You can add any number of virtual WXA 500s, WXA 5000s and WXA 6000s to a SonicWall Security Appliance, but the number of accelerated connections is dependent on the license that has been purchased.

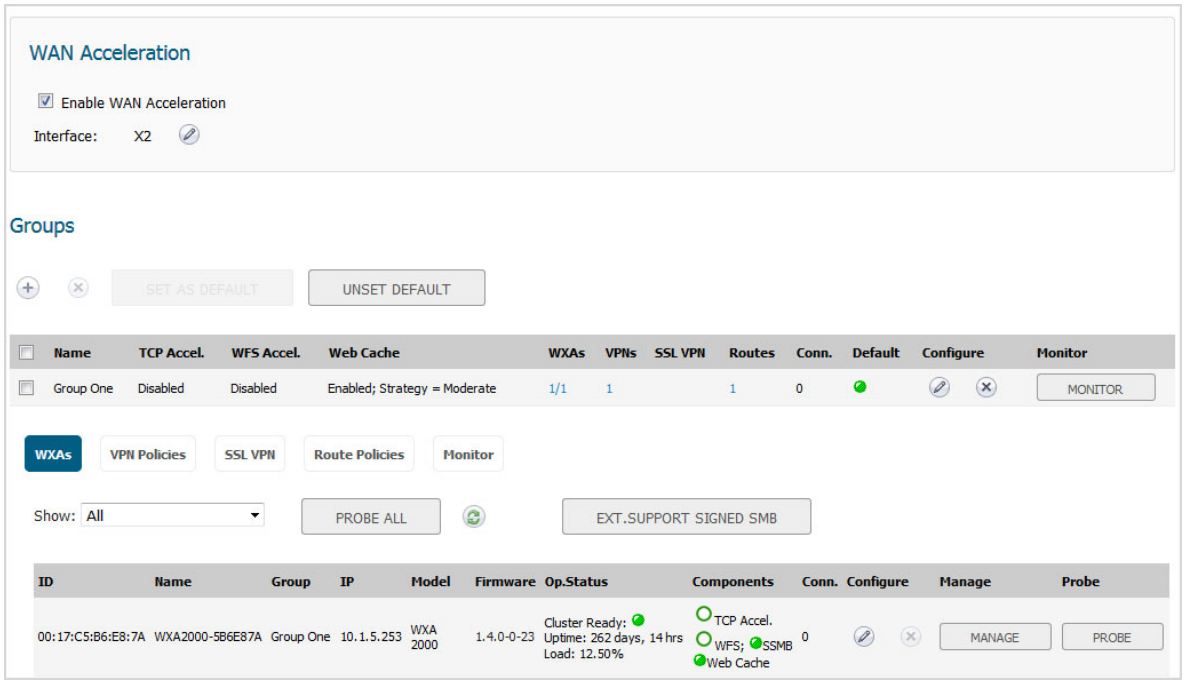
If the allowed number of accelerated connections is exceeded, all excess connections bypass the cluster, no matter how many WXAs are in the cluster. It is the Administrator's responsibility to ensure that a sufficient number of WXAs are attached to the Security Appliance to handle the number of licensed connections that they wish to support.

Permitting Acceleration on a Route Policy

After you have configured the WXA, you can permit acceleration for a Route Policy. You can either permit acceleration for a Route Policy on the **System Setup > WAN Acceleration** page or on the **Network > Routing** page (see Configuring Route Advertisements and Route Policies on page 450).

If you do not have a VPN configured on your network and you are using a custom Route Policy, you need to add two route Policies on each site: One for outgoing traffic, and one for incoming traffic.

System Setup > WAN Acceleration



Topics:

- Enabling WAN Acceleration on page 740
- Managing Groups on page 740
- Managing WXAs through the WXAs Table on page 744
- Configuring WXA for VPN Polices on page 760
- Configuring Acceleration of SSL VPN Traffic on page 760

Enabling WAN Acceleration

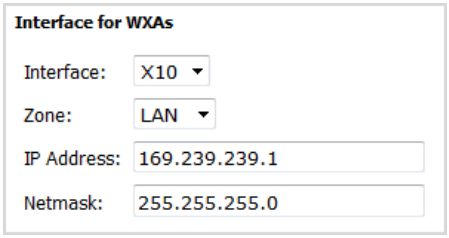
To enable WAN Acceleration:

1 Navigate to **System Setup > WAN Acceleration**.



2 In the **WAN Acceleration** section, select **Enable WAN Acceleration**.

3 Select the interface to which the WXAs are connected by clicking the **Interface Edit** icon. The **Interface for WXAs** popup displays.



4 From **Interface**, select the interface to which the WXAs are connected.

5 From **Zone**, select the zone for the interface.

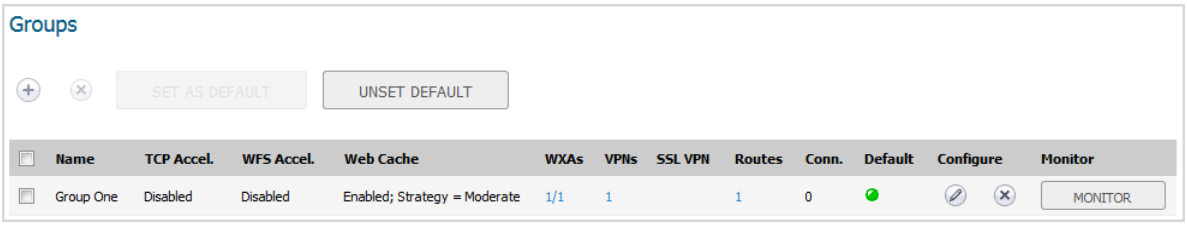
6 In the **IP Address** field, enter the IP address of the selected interface.

> **IMPORTANT:** The DHCP range, used to allocate addresses to the WXAs, is determined by the interface's IP address and the netmask.

7 In the **Netmask** field, enter the netmask of the selected interface.

8 Click **OK**.

Managing Groups



This Column	Indicates or displays
Name	Name of the group
TCP Accel.	Whether TCP Acceleration is **Enabled** or **Disabled**.
WFS Accel.	Whether WFS Acceleration is **Enabled** or **Disabled**.

This Column	Indicates or displays
Web Cache	• Whether Web Cache is **Enabled** or **Disabled**. • The caching strategy: **Minimal**, **Moderate**, or **Aggressive**.
WXAs	Number of available WXAs (online and cluster ready) found versus the number configured for the group: *found/configured*.
VPNs	Number of VPNs whose acceleration is governed by the group.
SSL VPN	Number of SSL VPNs whose acceleration is governed by the group.
Routes	Number of routes whose acceleration is governed by the group.
Conn.	Number of connections currently passing through the WXAs in the group. If you mouse over the number, a popup displays showing the: • Sum of supported connections for each WXA model in the group. • Overall licensed number of concurrent connections that can be accelerated.
Default	A green icon to indicate the default group.
Configure	**Edit** and **Delete** icons for the group.
Monitor	**Monitor** button to display the **WXA Connections** monitor.

Topics:

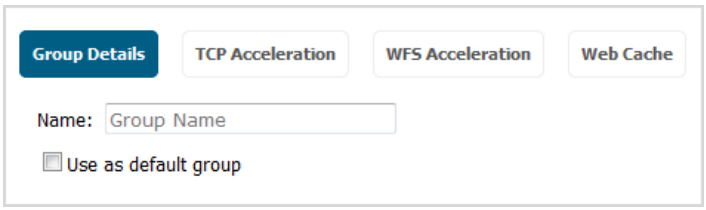
- Adding a Group on page 741
- Setting the Default Group on page 743
- Editing a Group on page 743
- Deleting Groups on page 744

Adding a Group

To add a group:

1 Navigate to **System Setup > WAN Acceleration**.

2 In the **Groups** section, click the **Add** icon. the **New Group** dialog displays.



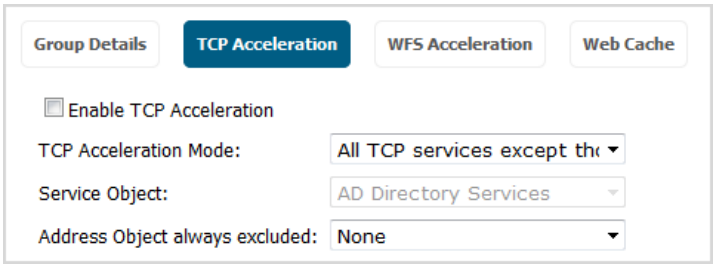
3 Enter a meaningful name for the group in the **Name** field.

4 To designate the group as the default group, select **Use as default group**. This option is not selected by default.

5 If you are:

- Not going to use TCP Acceleration, go to Step 10.

- Using TCP Acceleration, click **TCP Acceleration**.



6 Click **Enable TCP Acceleration**.

7 Select the mode from **TCP Acceleration Mode**:

- **All TCP services except those excluded by default** (this is the default; the Service Object drop-down menu is dimmed); go to Step 9
- **All TCP services except those specified in the Service Object**
- **All TCP services except those specified in the Service Object and those excluded by default**
- **Only TCP services specified in the Service Object** (to enable TCP Acceleration only for the Service Object)

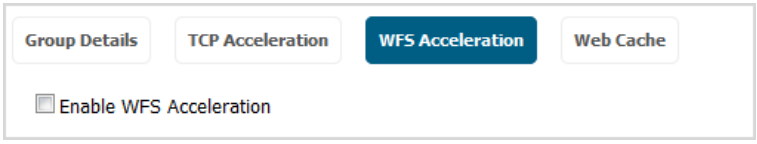
> **TIP:** To see the excluded TCP services, mouse over T**CP Acceleration Mode** to display a popup listing the excluded services.

8 From **Service Object**, select the service object to be excluded or included.

9 To exclude an Address Object from TCP Acceleration, select the address object from **Address Object always excluded**; the default is **None**.

10 If you are:

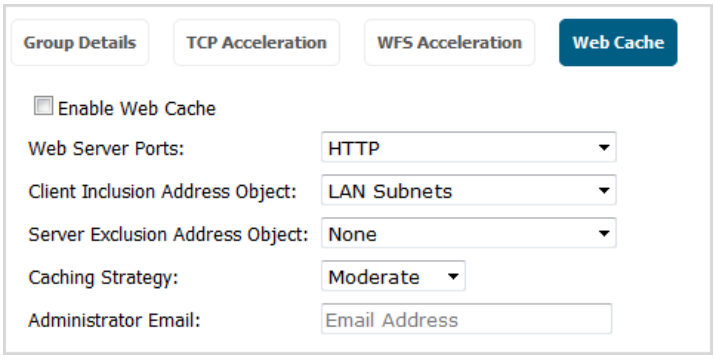
- Not going to use WFS Acceleration, go to Step 12.
- Using WFS Acceleration, click **WFS Acceleration**.



11 Select **Enable WFS Acceleration**.

12 If you are:

- Not going to use web cache, go to Step 19.
- Using web cache, click **Web Cache**.



13 Select **Enable Web Cache**.

14 From **Web Server Ports**, select the Service Object representing the web server ports on which traffic is intercepted and sent to the WXA Web Cache. The default is **HTTP**.

15 From **Client Inclusion Address Object**, select the Address Object or Group that represents those local subnets whose web traffic should be diverted through the WXA Web Cache. The default is **LAN Subnets**.

16 From **Server Exclusion Address Object**, select the Address Object or Group that contains the destination addresses of web servers for which traffic should not be diverted through the WXA Web Cache. The default is **None**: no web sever is excluded and all appropriate traffic is sent through the WXA.

17 From **Caching Strategy**, select the caching strategy that determines the types of objects that are cached and the properties that dictate how long they remain in the cache:

Minimal	Offers basic caching in which the Web Cache caches an object unless the HTTP headers specifically say not to (such as no cache or an expire time occurring in the past)
Moderate	A little less strict, and stores objects in the cache for longer periods. This is the default.
Aggressive	Ignores header options, such as no store and reload, and overrides expire times. CAUTION: **This strategy should be used with great care as it violates the HTTP standard and may lead to unwanted consequences.**

NOTE: Both **Moderate** and **Aggressive** modes include caching of YouTube videos.

18 Optionally, in the **Administrator Email** field, enter the administrator email address.

19 Click **OK**.

Setting the Default Group

Usually, the default group is specified when the group is configured (see Adding a Group on page 741), but you can change the default group at any time.

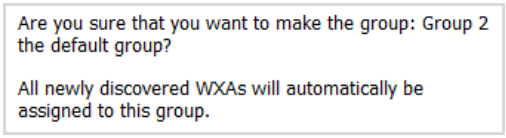
To change the default group:

1 Navigate to **System Setup > WAN Acceleration**.

2 In the **Groups** table, deselect the default group.

3 Select the group to become the default. **Set As Default** becomes available.

4 Click **Set As Default**. A confirmation message displays.



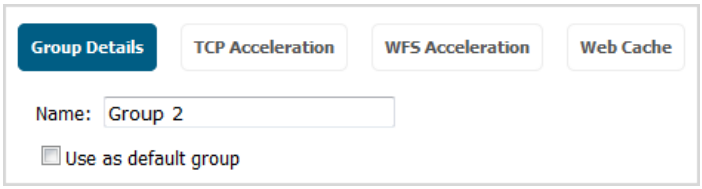
5 Click **YES**. The green indicator now shows in the **Default** column for the default group.

Editing a Group

To edit a group:

1 Navigate to **System Setup > WAN Acceleration**.

2 In the **Groups** table, click the **Edit** icon for the group to be edited. The **Edit Group** dialog displays.



3 Follow Step 3 through Step 19 in Adding a Group on page 741.

Deleting Groups

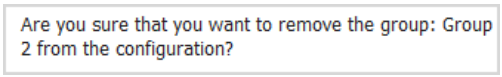
You can delete one or more groups. Groups with associated WXAs or used to govern the acceleration on one or more VPNs, SSL VPNs, or routes cannot be deleted.

To delete a group:

1 Navigate to **System Setup > WAN Acceleration**.

2 In the **Groups** table, select the group to be deleted.

3 Click the **Delete** icon for the group. A confirmation message displays.



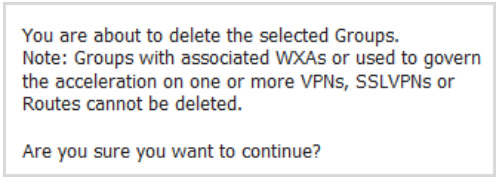
4 Click **YES**.

To delete multiple groups:

1 Navigate to **System Setup > WAN Acceleration**.

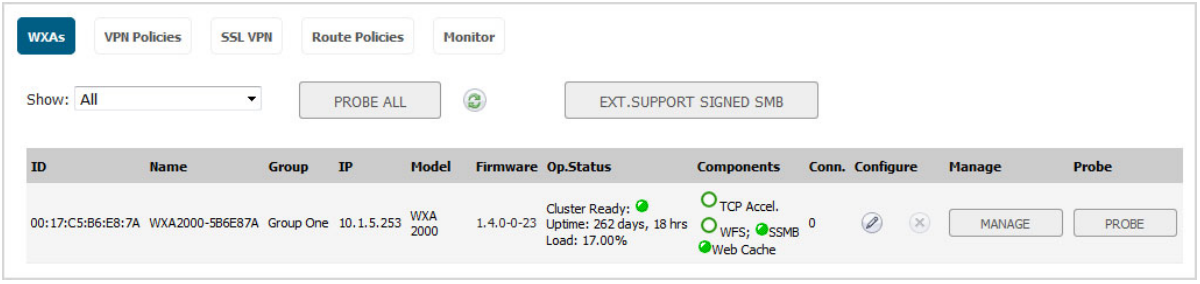
2 In the **Groups** table, select the groups to be deleted. The Delete icon above the table becomes available.

3 Click the **Delete** icon. A confirmation message displays.



4 Click **YES**.

Managing WXAs through the WXAs Table



ID	MAC address of the WXA series appliance.
Name	Name of the WXA series appliance.
Group	Group to which the WXA series appliance belongs.
IP	IP address of the WXA series appliance.
Model	Model of the WXA series appliance.
Firmware	Version of the firmware installed on the WXA series appliance. **NOTE:** Clicking on the firmware version for the WXA series appliance takes you to **MANAGE | Updates | WXA Firmware**. For information about WXA firmware and how to update it, see *SonicOS Update*.
Op. Status	Displays the operational status of the WXA series appliance: • **Cluster Ready** – a green dot indicates the WXA is available for clustering. • **Uptime** – days and hours the WXA has been operational. • **Load** – the rolling average load on the WXA, as a percentage.
Components	• **TCP Acceleration** • **WFS Acceleration** • **WFS Extended Support for Signed SMB** • **Web Cache** Displays the status of the acceleration components: • A green dot indicates the service is running on the WXA. • A white dot indicates the service is running on the WXA and is available to accelerate traffic, but the component is currently disable within the WXA's group settings.
Conn.	Number of connections currently passing through the WXA. The tooltip also displays: • Maximum number of connects supported by this particular WXA model. • Overall licensed number of concurrent connects that can be accessed.
Configure	**Edit** and **Delete** icons. **NOTE:** Active WXA series appliance cannot be deleted.
Manage	**Manage** button, which displays the **Manage WXA** dialog.
Probe	**Probe** button, which probes the WXA and updates the statistics in the table.

Topics:

- Filtering the WXAs Table on page 745
- Probing on page 746
- Refreshing the WXA Table on page 746
- Enabling Extended Support for Signed SMB on page 746

Filtering the WXAs Table

By default, all WXAs are shown in the table. You can restrict the display to only selected groups or unassigned WXAs.

To filter the WXAs Table display:

1 Navigate to **System Setup > WAN Acceleration**.

2 From **Show**, select what to display:

All (default)	Displays all WXAs
For Selected Groups	Displays only those WXAs belonging to the selected group
Unassigned	Displays only WXAs not assigned to any group

3 If you selected **For Selected Group**s, in the **Groups** table, select the group containing the WXAs to display.

Probing

Probing verifies the presence and status of the WXA series appliance(s) and also pushes the latest Group settings to the WXA series appliance. You can probe an individual WXA or all WXAs.

To probe all WXAs:

1 Navigate to **System Setup > WAN Acceleration**.

2 Click **Probe All**. The WXAs table is updated.

To probe an individual WXA:

1 Navigate to **System Setup > WAN Acceleration**.

2 In the WXAs table, click **Probe** in the **Probe** column for the WXA. The display for that WXA is updated.

Refreshing the WXA Table

Refreshing the WXA table refreshes the list of WXAs along with the status of the different acceleration components on each WXA.

To refresh the WXA table:

1 Navigate to **System Setup > WAN Acceleration**.

2 Click the **Refresh** icon.

Enabling Extended Support for Signed SMB

NOTE: Configuring WFS Extended Support for Signed SMB is done outside the configuration of groups.

When you configure WFS Extended Support for Signed SMB, you select the WXA series appliance that is dedicated to accelerating Signed SMB traffic.

Topics:

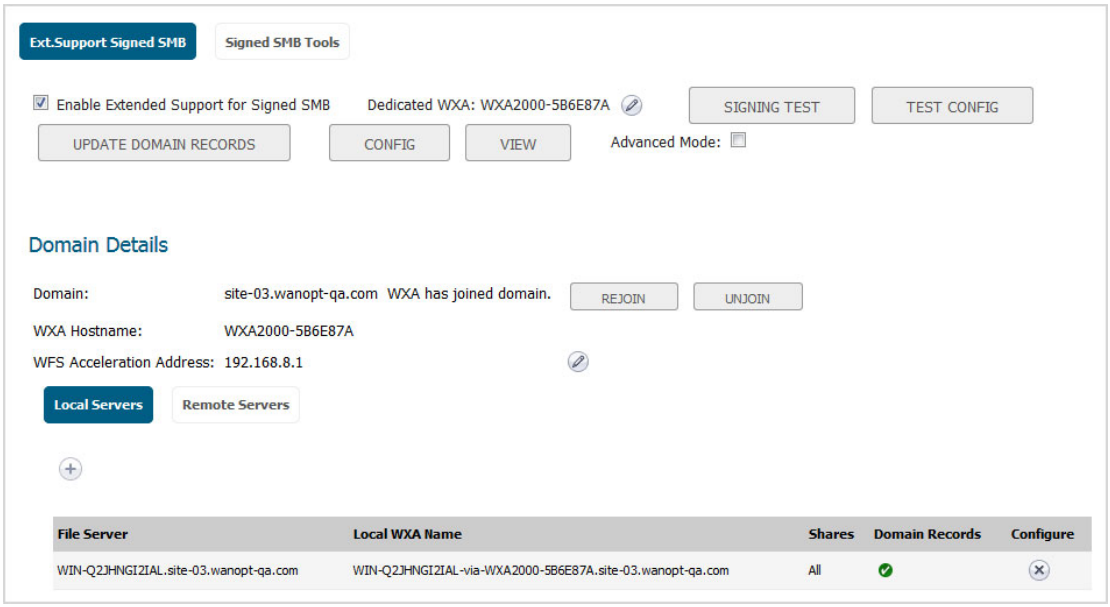
- Ext. Support Signed SMB on page 747
- Advanced Mode on page 751
- Domain Details on page 752
- Local/Remote Servers Table on page 754

Ext. Support Signed SMB

To configure extended support for Signed SMB Acceleration:

1 Navigate to **System Setup > WAN Acceleration**.

2 Click **Ext. Support Signed SMB**. the **Extended Support for Signed SMB Acceleration** dialog displays.

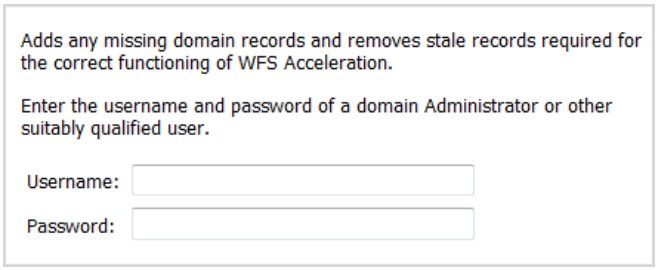


3 Select **Enable Extended Support for Signed SMB**. This option is selected by default.

4 Click the **Edit** icon to select the WXA to be dedicated to extended support.

> **IMPORTANT:** Changing a WXA that is already dedicated to extended support for Signed SMB acceleration, causes any active sessions and file transfers to be terminated, perhaps resulting in loss of data.
>
> The new WXA must join the domain and be configured with the relevant servers and shares. The paths on end users' devices must be changed to reflect the new WXA's configuration. Otherwise, to keep the same paths, the first WXA must unjoin the domain and all domain records must be removed before configuring the new WXA with the same host name and identical settings.

5 Add any missing domain records and remove stale records for SPN aliases and "trusted for delegation" by clicking **Update Domain Records**. The **Update Domain Records** popup displays.

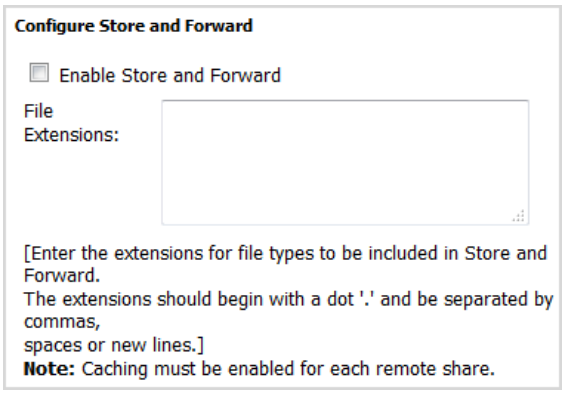


6 Enter the domain Administrator's user name and password in the **Username** and **Password** fields, respectively.

7 Click **Update Records**.

8 If you are not using Store and Forward, go to Step 17.

9 To configure settings for Store and Forward, click **Config**. The **Configure Store and Forward** dialog displays.



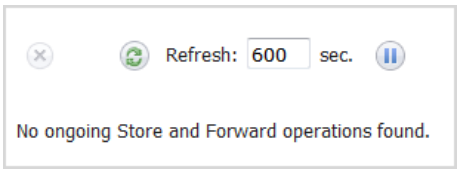
10 Select **Enable Store and Forward**. This option is not selected by default.

11 Enter the extensions for file types to be stored and forwarded in the **File Extensions** field. The extensions should begin with a dot (.) and be separated by commas, spaces, or new lines.

NOTE: Caching must be enabled for each remote share.

12 Click **OK**.

13 To view current file operations for Store and Forward, click **View**. the **Store and Forward** dialog displays.



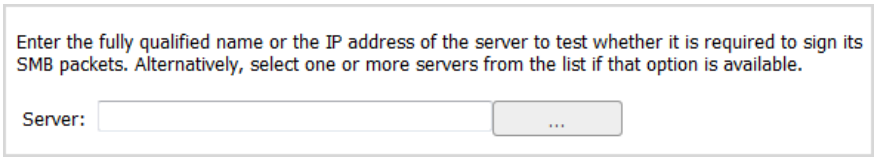
14 To:

- Refresh the display, click the **Refresh** icon.
- Pause the display, click the **Pause** icon.

15 To change the refresh rate, enter the rate in the **Refresh** field. The minimum is 1 second, the maximum is 999 seconds, and the default is **600** seconds.

16 If you do not want to test for signing, go to Step 24.

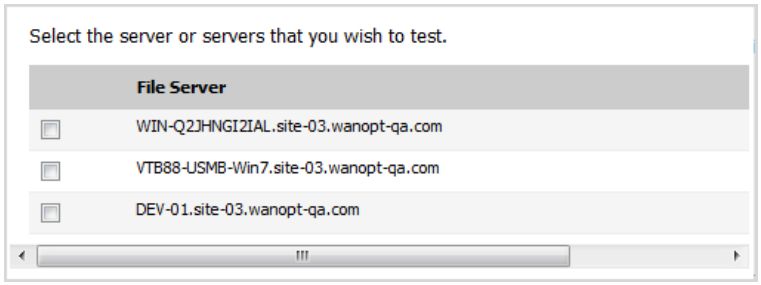
17 To test whether traffic to specified servers requires signing, click **Signing Test**. The **Signing Test** dialog displays.



18 Either:

- Enter the fully qualified name or IP address of the server to test.

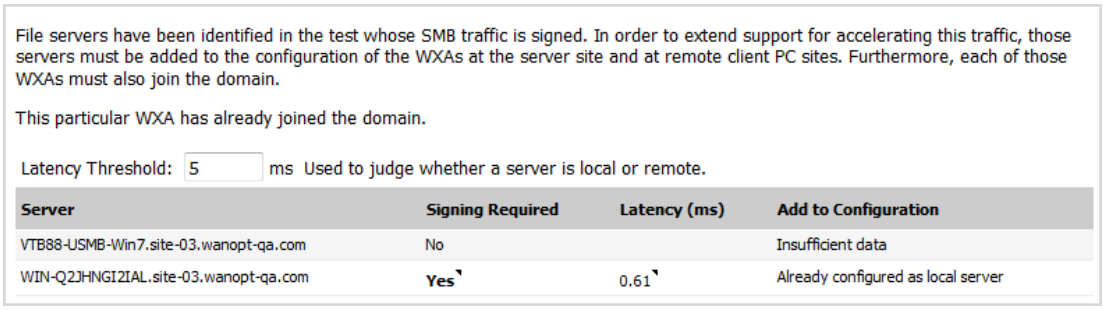
- Display a list of servers by clicking **List** (…). The **Select Servers to Test** popup displays:



19 Select one or more servers.

20 Click **OK**. The popup closes.

21 Click **OK**. Please be patient as it could take a few minutes for the test results to display.



IMPORTANT: If SMB signing is required, the WXA series appliance must join the domain. The test results indicate whether the WXA series appliance has joined the domain.

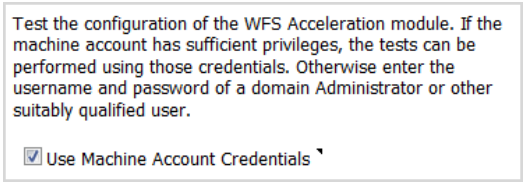
TIP: If signing is required, mousing over the tooltip by:
- **Yes** displays the addresses required for signing
- Latency time displays the times the server has exceeded the time set in **Latency Threshold**.

22 To specify a threshold to determine whether a server is local or remote, enter the time, in milliseconds, in **Latency Threshold**. The minimum time is 1 millisecond, the maximum is 99999999 milliseconds, and the default is **5** seconds.

23 Click **OK**.

24 If you do not want to test the WFS Acceleration module, go to Step 29.

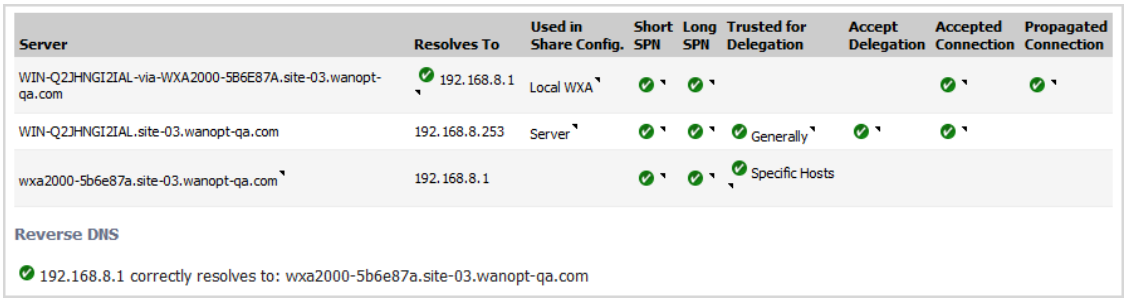
25 To test the WFS Acceleration module's configuration, click **Test Config**. The **Test Configuration** popup displays.



26 If the WXA series appliance has its own machine account with the appropriate permissions, select **Use Machine Account Credentials**. This option is selected by default.

NOTE: If the WXA does not have the appropriate permissions, you need to enter the Administrator user name and password of a domain administrator.

27 Click **Run Tests**. Please be patient as it may take a few minutes for the test results to display.



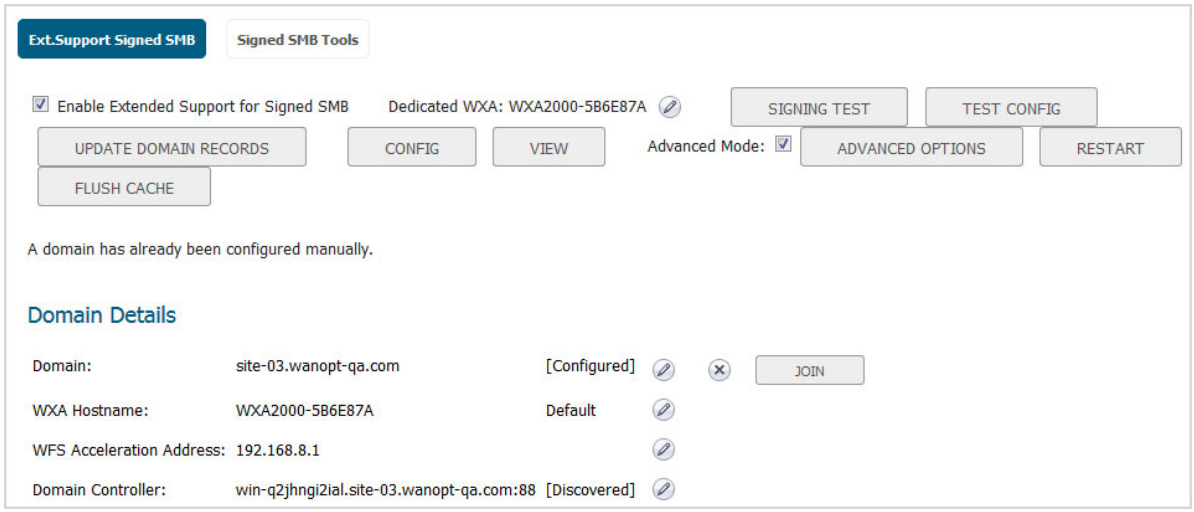
Server	Service Principal Name (SPN) of servers tested
Resolves to	IP address; a green checkmark indicates a correct resolution
Used in Share Config.	Indicates how the SPN is used as a name to identify the appliance
Short SPN	Indicates whether the short SPN is present on the machine
Long SPN	Indicates whether the long SPN is present on the machine
Trusted for Delegation	A green checkmark indicates whether the server is trusted for delegation: • Generally – in general, the server is trusted for delegation • Specific hosts – mousing over the tooltip displays the hosts for which the server is trusted for delegation
Accept Delegation	A green checkmark indicates the server accepts delegation; mousing over the tooltip displays the hosts that can present credentials using the server's short name or long name.
Accepted Connection	A green checkmark indicates that the server has accepted authenticated connections; mousing over the tooltip displays the connections.
Propagated Connection	A green checkmark indicates the server propagated authenticated connections; mousing over the tooltip displays the connections.
Reverse DNS	Displays how a reverse DNS resolves

28 Click **Close**.

29 To use Advanced mode, select **Advanced Mode**. This option is not selected by default.

Advanced Mode

When **Advanced Mode** is selected, the options on the Extended Support for Signed SMB Acceleration dialog change:



To configure Advanced Mode:

1 Navigate to **System Setup > WAN Acceleration**.

2 Click **Ext. Support Signed SMB**. the **Extended Support for Signed SMB Acceleration** dialog displays.

3 Select **Advanced Mode**. The options change.

4 Click **ADVANCED OPTIONS**. The **Advanced Options** dialog displays.

5 From **Client Signing**, select the signing option clients must use:

- **Auto** (default)
- **Mandatory**
- **Disabled**

6 From **Server Signing**, select the signing option servers must use:

- **Auto** (default)
- **Mandatory**
- **Disabled**

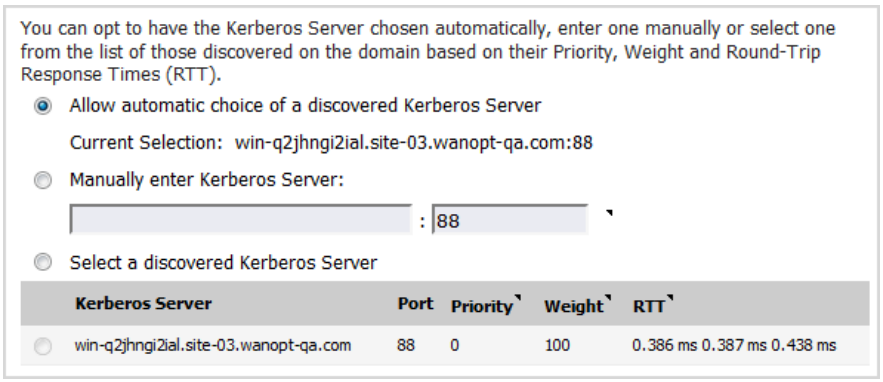
7 Enter the maximum amount of data, in bytes, a client can transmit in the **Max Transmit** field. The default is **4096**.

8 Click **OK**.

9 To restart the WFS Extended Support for Signed SMB service, click **RESTART**.

10 To flush the cache, click **FLUSH CACHE**.

11 In the **Domain Details** section, configure the Kerberos server by clicking the **Edit** icon of **Domain Controller**. The **Configure Kerberos Server** dialog displays.

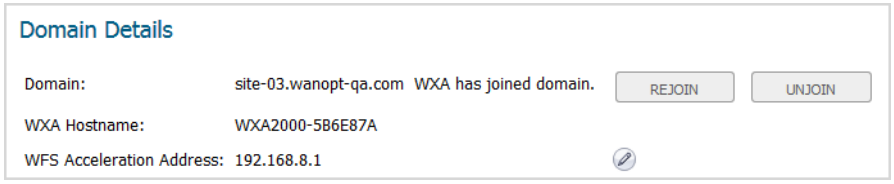


12 Specify how the Kerberos server is to be chosen:

- **Allow automatic choice of a discovered Kerberos Server** – The current selection and its port are displayed.
- **Manually enter Kerberos Server** – The name and port fields become available.
  - Enter the name and port for the Kerberos server used for authentication on the domain.
- **Select a discovered Kerberos Server** – The discovered entries in the **Kerberos Server** table becomes available.
  - Choose one of the entries.

13 Click **OK**.

Domain Details



Domain	Name of the domain and whether the WXA has joined it.
WXA Hostname	Name of the WXA
WFS Acceleration Address	IP address of the WFS Acceleration module

Topics:

- Rejoining a Domain on page 752
- Unjoining a Domain on page 753
- Joining a Domain on page 753
- Deleting a Domain on page 754

Rejoining a Domain

To rejoin a domain:

1 Navigate to **System Setup > WAN Acceleration**.

2 Click **Ext. Support Signed SMB**. the **Extended Support for Signed SMB Acceleration** dialog displays.

3 Click **Rejoin**.

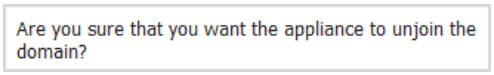
Unjoining a Domain

To unjoin a domain:

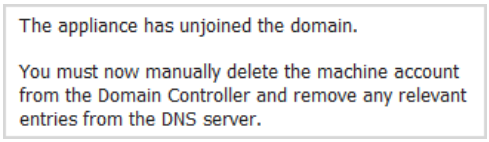
1 Navigate to **System Setup > WAN Acceleration**.

2 Click **Ext. Support Signed SMB**. the **Extended Support for Signed SMB Acceleration** dialog displays.

3 Click **Unjoin**. A confirmation message displays.



4 Click **YES**. A confirmation message displays.



5 Click **OK**. **Rejoin** becomes a **Delete** icon.

6 Manually delete the machine account from the domain controller.

7 Remove any relevant entries from the DNS server.

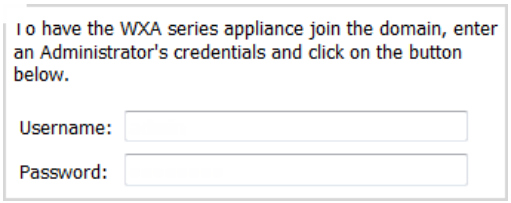
Joining a Domain

To join a domain:

1 Navigate to **System Setup > WAN Acceleration**.

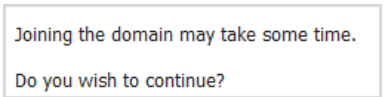
2 Click **Ext. Support Signed SMB**. the **Extended Support for Signed SMB Acceleration** dialog displays.

3 Click **Join**. The **Join Domain** dialog displays.

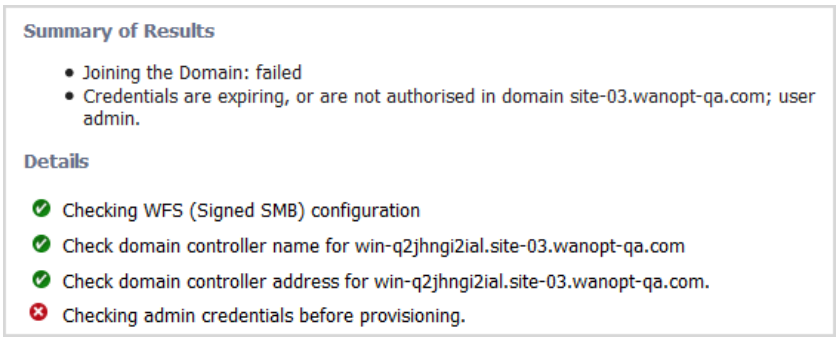


4 Enter the administrator’s username and password in the **Username** and **Password** fields, respectively.

5 Click **Join Domain**.



6 Click **Yes**. The **Join Domain Results** popup displays whether the joining was successful and the details of the process.



7 Click **Close**.

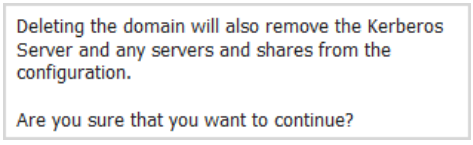
Deleting a Domain

To delete a domain:

1 Navigate to **System Setup > WAN Acceleration**.

2 Click **Ext. Support Signed SMB**. the **Extended Support for Signed SMB Acceleration** dialog displays.

3 Click the **Delete** icon.



4 Click **YES**.

Local/Remote Servers Table



File Server	Name of the file server.
Local WXA Name	Name of the local WXA server.
Shares	
Domain Records	A green checkmark indicates domain records are current, a red X indicates the records need to be updated. To update, click **UPDATE DOMAIN RECORDS**.
Configure	Displays the **Delete** icon.

Topics:

- Deleting a Server on page 755

- Adding a Local File Server on page 755

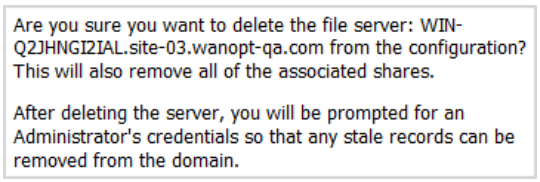
Deleting a Server

To delete a server:

1 Navigate to **System Setup > WAN Acceleration**.

2 Click **Ext. Support Signed SMB**. the **Extended Support for Signed SMB Acceleration** dialog displays.

3 In the **Servers** table, click the **Delete** icon for the server to be deleted. A confirmation message displays.



4 Click **Delete**. The dialog displays.

5 Enter the administrator's username and password in the **Username** and **Password** fields, respectively.

6 Click **OK**.

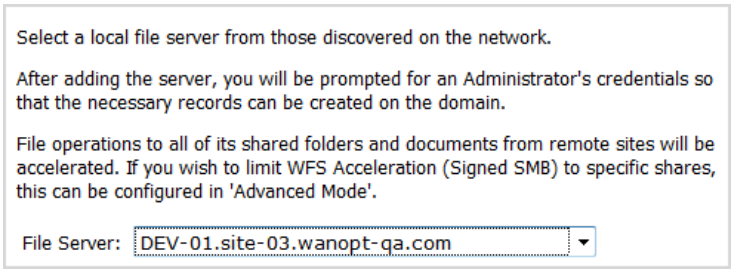
Adding a Local File Server

To add a local file server:

1 Navigate to **System Setup > WAN Acceleration**.

2 Click **Ext. Support Signed SMB**. the **Extended Support for Signed SMB Acceleration** dialog displays.

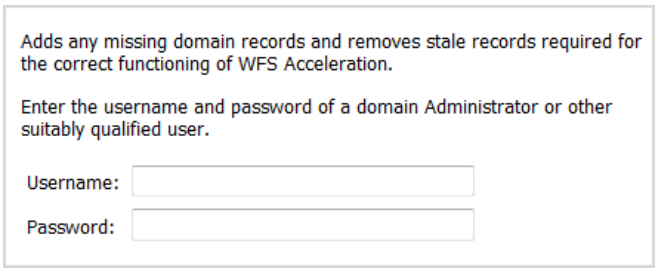
3 Under **Local Servers**, click the **Add** icon. The **Add Local File Server** dialog displays.



4 Select the file server from **File Server**.

5 Click **OK**.

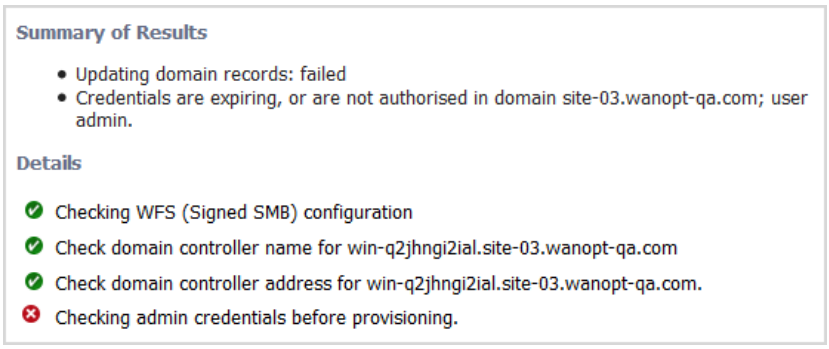
6 The **Update Domain Records** dialog displays.



7 Enter the administrator's username and password in the **Username** and **Password** fields, respectively.

8 Click **UPDATE RECORDS**.

9 Click **Yes**. The **Update Domain Results** popup displays whether the joining was successful and the details of the process.



10 Click **Close**.

Displaying Remote Servers

To display remote servers in the Servers table:

1 Navigate to **System Setup > WAN Acceleration**.

2 Click **Ext. Support Signed SMB**. the **Extended Support for Signed SMB Acceleration** dialog displays.

3 Under **Domain Details**, click Remote Servers. The **Servers** table displays al configured remote servers.

Adding a Remote Server

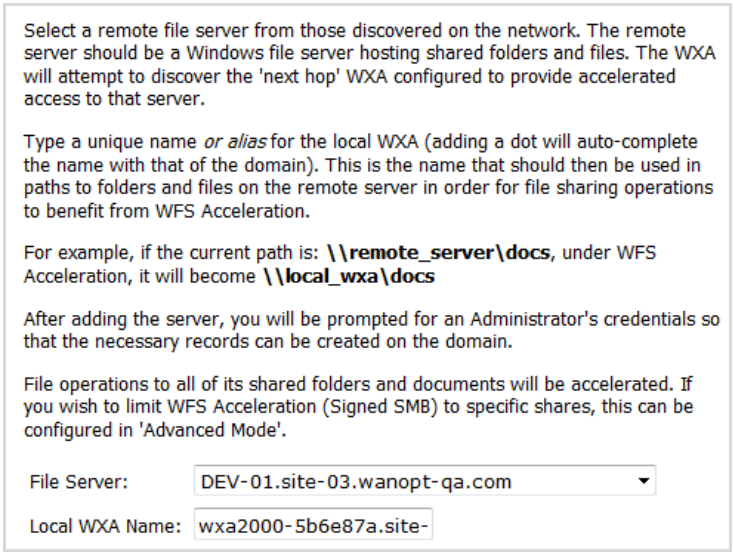
To add a remote server:

1 Navigate to **System Setup > WAN Acceleration**.

2 Click **Ext. Support Signed SMB**. the **Extended Support for Signed SMB Acceleration** dialog displays.

3 Click **Remote Servers**. The **Remote Servers** table displays.

4 Under **Remote Servers**, click the **Add** icon. The **Add Remote File Server** dialog displays.

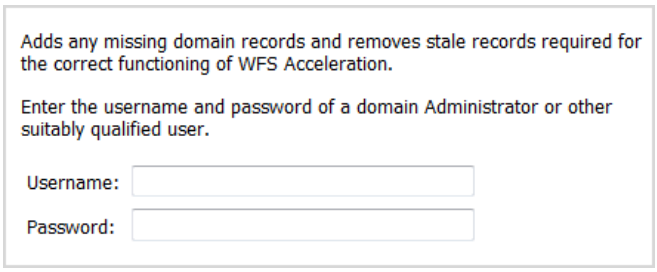


5 Select the file server from **File Server**.

6 Enter the unique name or alias for the WXA server in the **Local WXA Name** field.

7 Click **OK**.

8 The **Update Domain Records** dialog displays.



9 Enter the administrator's username and password in the **Username** and **Password** fields, respectively.

10 Click **UPDATE RECORDS**.

11 Click **Yes**. The **Update Domain Results** popup displays whether the joining was successful and the details of the process.

Using Signed SMB Tools

Topics:

- DNS Name Lookup on page 757
- Available Shares on page 758

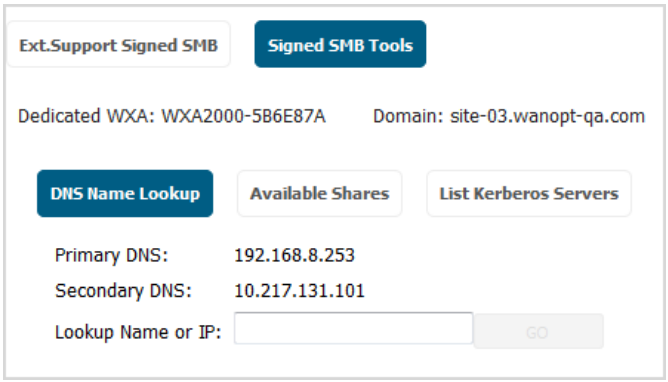
DNS Name Lookup

To look up a DNS name:

1 Navigate to **System Setup > WAN Acceleration**.

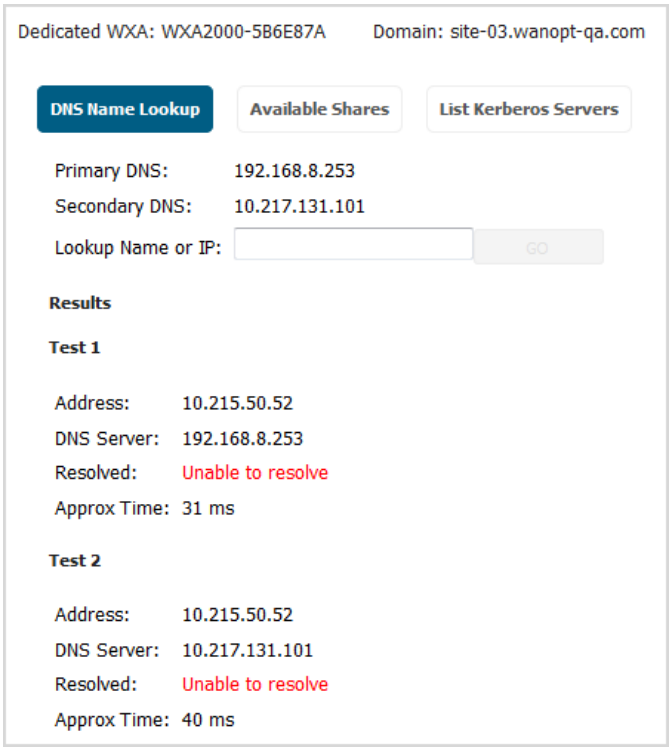
2 Click **Ext. Support Signed SMB**. the **Extended Support for Signed SMB Acceleration** dialog displays.

3 Click **Signed SMB Tools**.



4 Enter the server to look up in the **Lookup Name or IP** field. **Go** becomes available.

5 Click **GO**. The results are displayed:



Available Shares

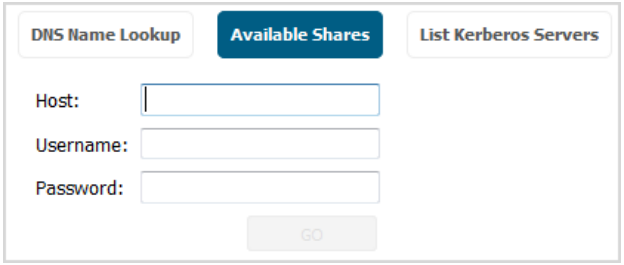
To look up available shares:

1 Navigate to **System Setup > WAN Acceleration**.

2 Click **Ext. Support Signed SMB**. the **Extended Support for Signed SMB Acceleration** dialog displays.

3 Click **Signed SMB Tools**.

4 Click **Available Shares**.



5 Enter the server to look up in the **Host** field.

6 Enter the administrator's user name and password in the Username and Password fields respectively. **Go** becomes available.

7 Click **GO**. The Available Shares popup displays.

8 Click **OK**.

List Kerberos Servers

To list Kerberos servers:

1 Navigate to **System Setup > WAN Acceleration**.

2 Click **Ext. Support Signed SMB**. the **Extended Support for Signed SMB Acceleration** dialog displays.

3 Click **Signed SMB Tools**.

4 Click **List Kerberos Servers**.

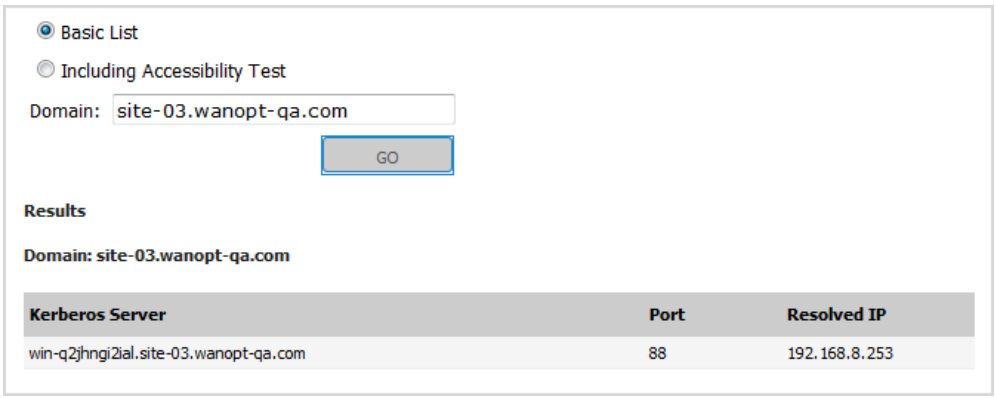


5 Choose how to list the servers:

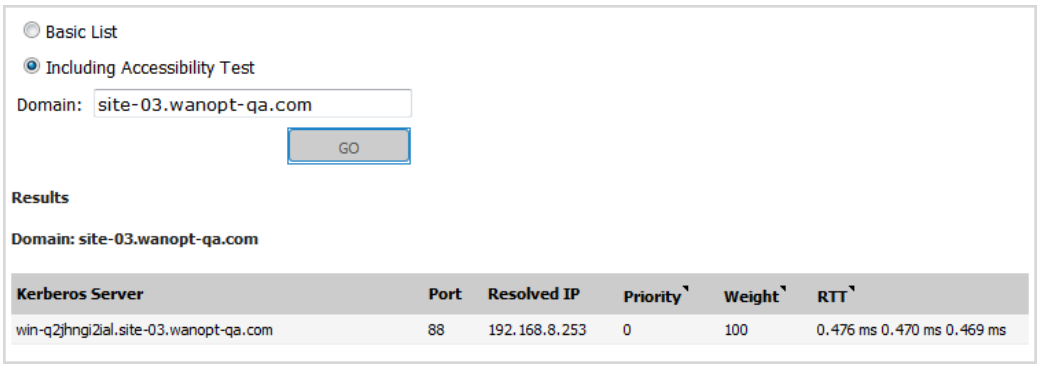
- **Basic List** – Displays only the port and resolved IP of the Kerberos server.
- **Including Accessibility Test** (default) – Includes the server's priority, weight, and RTT.

6 Click **GO**. The results display:

- **Basic List**:



- **Including Accessibility Test**:



Kerberos Server	Name of the Kerberos server.
Port	Port of the Kerberos server.
Resolved IP	IP address to which the server's name resolves.
Priority	Priority of the Kerberos server. Lower values are preferred.
Weight	Relative weight for Kerberos Servers with the same priority. Higher values are preferred.
RTT	Round trip time (RTT) for probes to the Kerberos server.

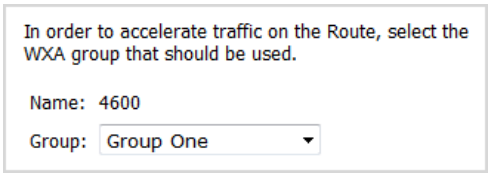
Configuring WXA for VPN Polices



The VPN Policies table displays all the VPN polices with WXA acceleration settings.

To edit the WXA acceleration settings for the VPN policies:

1. Navigate to **System Setup > WAN Acceleration**.
2. Click **VPN Policies**.
3. To filter the policies, select the WXA group from the **Groups** table.
4. Select **For Selected Groups** from **Show**. The default is **All**.
5. Click the **Edit** icon for the policy. The **Edit VPN** dialog displays.



6. Select the WXA group that applies to the VPN policy from **Group**.
7. Click **OK**.

Configuring Acceleration of SSL VPN Traffic

You can enable or disable acceleration of NetExtender SSL VPN traffic from WXAC clients.

NOTE: The WXA must be licensed to support the NetExtender WAN Acceleration Client (WXAC).

To enable WXAC:

1. Navigate to **System Setup > WAN Acceleration**.

2 Click **SSL VPN**.



3 Select the group for which WXAC is to be enabled from **Group** (the default is **None**). **Accept** becomes available.

4 Click **Accept**.

5 The number of active licenses currently in use is displayed under the **Group** drop-down menu.

Displaying and Editing Route Polices for WXA



Source	Gateway where the VPN traffic originated.
Destination	Destination of the VPN traffic.
Comment	Optional comment included when configuring the route.
Group	The group that applies to the route.
Edit	Displays the **Edit** icon.

To filter the routes displayed in the Routes table:

1 Navigate to **System Setup > WAN Acceleration**.

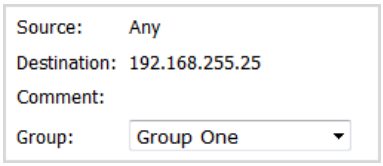
2 Under the **Groups** table, select the routes for a particular group from **Show**. The default is **All**.

To edit a route policy:

1 Navigate to **System Setup > WAN Acceleration**.

2 Under the **Groups** table, click **Route Policies**.

3 Click the **Edit** icon for the desired route. The **Edit Route** popup displays.



4 Select the to which the route applies from **Group**.

5 Click **OK**.

Monitoring Group Connections

You can view the total number of connections passing through the WXA as either a line or bar (stack) chart. You can also view the total connections for a particular group.



To display the connections:

1 Navigate to **System Setup > WAN Acceleration**.

2 Under the **Groups** table, click **Monitor**.

> **TIP:** You can also display the connections by clicking the **MONITOR** for the desired group in the **Groups** table.

3 Select the group for displaying connections from **Group**. The default is **ALL**.

4 From **Chart Type**, select how you want the data displayed:

- **Line** (default)
- **Stack** (bar chart)

5 To zoom in on a selected area of the chart, drag the mouse over the desired area.

6 To return the chart to its default zoom, click **RESET ZOOM**.

Part 9

System Setup | VoIP

- About VoIP
- Configuring SonicWall VoIP Features

51

About VoIP

Topics:

- About VoIP on page 765
  - What is VoIP? on page 765
  - VoIP Security on page 765
  - VoIP Protocols on page 766
  - SonicWall's VoIP Capabilities on page 768

About VoIP

Topics:

- What is VoIP? on page 765
- VoIP Security on page 765
- VoIP Protocols on page 766
- SonicWall's VoIP Capabilities on page 768

What is VoIP?

Voice over IP (VoIP) is an umbrella term for a set of technologies that allow voice traffic to be carried over Internet Protocol (IP) networks. VoIP transfers the voice streams of audio calls into data packets as opposed to traditional, analog circuit-switched voice communications used by the public switched telephone network (PSTN).

VoIP is the major driving force behind the convergence of networking and telecommunications by combining voice telephony and data into a single integrated IP network system. VoIP is all about saving cost for companies through eliminating costly redundant infrastructures and telecommunication usage charges while also delivering enhanced management features and calling services features.

VoIP Security

Companies implementing VoIP technologies in an effort to cut communication costs and extend corporate voice services to a distributed workforce face security risks associated with the convergence of voice and data networks. VoIP security and network integrity are an essential part of any VoIP deployment.

The same security threats that plague data networks today are inherited by VoIP but the addition of VoIP as an application on the network makes those threats even more dangerous. By adding VoIP components to your network, you are also adding new security requirements.

VoIP encompasses a number of complex standards that leave the door open for bugs and vulnerabilities within the software implementation. The same types of bugs and vulnerabilities that hamper every operating system and application available today also apply to VoIP equipment. Many of today's VoIP call servers and gateway devices are built on vulnerable Windows and Linux operating systems.

Security Appliance Requirements for VoIP

VoIP is more complicated than standard TCP/UDP-based applications. Because of the complexities of VoIP signaling and protocols, as well as inconsistencies that are introduced when a Security Appliance modifies source address and source port information with Network Address Translation (NAT), it is difficult for VoIP to effectively traverse a standard Security Appliance. Here are a few of the reasons why.

- **VoIP operates using two separate protocols** - A signaling protocol (between the client and VoIP Server) and a media protocol (between the clients). Port/IP address pairs used by the media protocols (RTP/RTCP) for each session are negotiated dynamically by the signaling protocols. Firewalls need to dynamically track and maintain this information, securely opening selected ports for the sessions and closing them at the appropriate time.
- **Multiple media ports are dynamically negotiated through the signaling session** - negotiations of the media ports are contained in the payload of the signaling protocols (IP address and port information). Firewalls need to perform deep packet inspection on each packet to acquire the information and dynamically maintain the sessions, thus demanding extra Security Appliance processing.
- **Source and destination IP addresses are embedded within the VoIP signaling packets** - A Security Appliance supporting NAT translates IP addresses and ports at the IP header level for packets. Fully symmetric NAT Security Appliances adjust their NAT bindings frequently, and may arbitrarily close the pinholes that allow inbound packets to pass into the network they protect, eliminating the service provider's ability to send inbound calls to the customer. To effectively support VoIP it is necessary for a NAT Security Appliance to perform deep packet inspection and transformation of embedded IP addresses and port information as the packets traverse the Security Appliance.
- **Firewalls need to process the signaling protocol suites consisting of different message formats used by different VoIP systems** - Just because two vendors use the same protocol suite does not necessarily mean they interoperate.

To overcome many of the hurdles introduced by the complexities of VoIP and NAT, vendors are offering Session Border Controllers (SBCs). An SBC sits on the Internet side of a Security Appliance and attempts to control the border of a VoIP network by terminating and re-originating all VoIP media and signaling traffic. In essence, SBCs act as a proxy for VoIP traffic for non-VoIP enabled Security Appliances. SonicWall Security Appliances are VoIP enabled Security Appliances that eliminate the need for an SBC on your network.

NOTE: VoIP is supported on all SonicWall appliances that can run SonicOS 6.2, as long as the VoIP application is RFC-compliant.

VoIP Protocols

VoIP technologies are built on two primary protocols: H.323 and SIP. These protocols can be applied either globally or per firewall rule.

Topics:

- H.323 on page 767
- SIP on page 767

H.323

H.323 is a standard developed by the International Telecommunications Union (ITU). It is a comprehensive suite of protocols for voice, video, and data communications between computers, terminals, network devices, and network services. H.323 is designed to enable users to make point-to-point multimedia phone calls over connectionless packet-switching networks such as private IP networks and the Internet. H.323 is widely supported by manufacturers of video conferencing equipment, VoIP equipment and Internet telephony software and devices.

H.323 uses a combination of TCP and UDP for signaling and ASN.1 for message encoding. H.323v1 was released in 1996 and H.323v5 was released in 2003. As the older standard, H.323 was embraced by many early VoIP players.

An H.323 network consists of four different types of entities:

- **Terminals** - Client end points for multimedia communications. An example would be an H.323 enabled Internet phone or PC.
- **Gatekeepers** - Performs services for call setup and tear down, and registering H.323 terminals for communications. Includes:
    - Address translation.
    - Registration, admission control, and status (RAS).
    - Internet Locator Service (ILS) also falls into this category (although it is not part of H.323). ILS uses LDAP (Lightweight Directory Access Protocol) rather than H.323 messages.
- **Multipoint control units (MCUs)** - Conference control and data distribution for multipoint communications between terminals.
- **Gateways** - Interoperation between H.323 networks and other communications services, such as the circuit-switched Packet Switched Telephone Network (PSTN).

SIP

The Session Initiation Protocol (SIP) standard was developed by the Internet Engineering Task Force (IETF). RFC 2543 was released in March 1999. RFC 3261 was released in June 2002. SIP is a signaling protocol for initiating, managing and terminating sessions. SIP supports 'presence' and mobility and can run over User Datagram Protocol (UDP) and Transmission Control Protocol (TCP).

Using SIP, a VoIP client can initiate and terminate call sessions, invite members into a conferencing session, and perform other telephony tasks. SIP also enables Private Branch Exchanges (PBXs), VoIP gateways, and other communications devices to communicate in standardized collaboration. SIP was also designed to avoid the heavy overhead of H.323.

A SIP network is composed of the following logical entities:

- **User Agent (UA)** - Initiates, receives and terminates calls.
- **Proxy Server** - Acts on behalf of UA in forwarding or responding to requests. A Proxy Server can fork requests to multiple servers. A back-to-back user agent (B2BUA) is a type of Proxy Server that treats each leg of a call passing through it as two distinct SIP call sessions: one between it and the calling phone and the other between it and the called phone. Other Proxy Servers treat all legs of the same call as a single SIP call session.
- **Redirect Server** - Responds to request but does not forward requests.
- **Registration Server** - Handles UA authentication and registration.

SonicWall's VoIP Capabilities

IMPORTANT: If **Wireless-Controller-Only** mode has been selected for **Wireless LAN Controller**, VOIP is disabled.

Topics:

- VoIP Security on page 768
- VoIP Network on page 769
- VoIP Network Interoperability on page 769
- Supported Interfaces on page 770
- Supported VoIP Protocols on page 770
- BWM and QoS on page 773
- How SonicOS Handles VoIP Calls on page 773

VoIP Security

- **Traffic legitimacy** - Stateful inspection of every VoIP signaling and media packet traversing the Security Appliance ensures all traffic is legitimate. Packets that exploit implementation flaws, causing effects such as buffer overflows in the target device, are the weapons of choice for many attackers. SonicWall Security Appliances detect and discard malformed and invalid packets before they reach their intended target.
- **Application-layer protection for VoIP protocols** - Full protection from application-level VoIP exploits through SonicWall Intrusion Prevention Service (IPS). IPS integrates a configurable, high performance scanning engine with a dynamically updated and provisioned database of attack and vulnerability signatures to protect networks against sophisticated Trojans and polymorphic threats. SonicWall extends its IPS signature database with a family of VoIP-specific signatures designed to prevent malicious traffic from reaching protected VoIP phones and servers.
- **DoS and DDoS attack protection** - Prevention of DoS and DDoS attacks, such as the SYN Flood, Ping of Death, and LAND (IP) attack, which are designed to disable a network or service.
    - Validating packet sequence for VoIP signaling packets using TCP to disallow out of sequence and retransmitted packets beyond window.
    - Using randomized TCP sequence numbers (generated by a cryptographic random number generator during connection setup) and validating the flow of data within each TCP session to prevent replay and data insertion attacks.
    - Ensures that attackers cannot overwhelm a server by attempting to open many TCP/IP connections (which are never fully established-usually due to a spoofed source address) by using SYN Flood protection.
- **Stateful monitoring** - Stateful monitoring ensures that packets, even though appearing valid in themselves, are appropriate for the current state of their associated VoIP connection.
- **Encrypted VoIP device support** - SonicWall supports VoIP devices capable of using encryption to protect the media exchange within a VoIP conversation or secure VoIP devices that do not support encrypted media using IPsec VPNs to protect VoIP calls.
- **Application-layer protection -** SonicWall delivers full protection from application-level VoIP exploits through SonicWall Intrusion Prevention Service (IPS). SonicWall IPS is built on a configurable, high performance Deep Packet Inspection engine that provides extended protection of key network services including VoIP, Windows services, and DNS. The extensible signature language used in SonicWall's Deep Packet Inspection engine also provides proactive defense against newly discovered application and

protocol vulnerabilities. Signature granularity allows SonicWall IPS to detect and prevent attacks based on a global, attack group, or per-signature basis to provide maximum flexibility and control false positives.

VoIP Network

- **VoIP over Wireless LAN (WLAN)** - SonicWall extends complete VoIP security to attached wireless networks with its Distributed Wireless Solution. All of the security features provided to VoIP devices attached to a wired network behind a SonicWall are also provided to VoIP devices using a wireless network.

  NOTE: SonicWall's Secure Wireless Solution includes the network enablers to extend secure VoIP communications over wireless networks.

- **Bandwidth Management (BWM) and Quality-of-Service (QoS)** - Bandwidth management (both ingress and egress) can be used to ensure that bandwidth remains available for time-sensitive VoIP traffic. BWM is integrated into SonicWall Quality of Service (QoS) features to provide predictability that is vital for certain types of applications.
- **WAN redundancy and load balancing** - WAN redundancy and load balancing allows for an interface to act as a secondary WAN port. This secondary WAN port can be used in a simple active/passive setup, where traffic is only routed through it if the primary WAN port is down or unavailable. Load balancing can be performed by splitting the routing of traffic based on destination.
- **High availability** - High availability is provided by SonicOS high availability, which ensures reliable, continuous connectivity in the event of a system failure.

VoIP Network Interoperability

- **Plug-and-protect support for VoIP devices** - With SonicOS, VoIP device adds, changes, and removals are handled automatically, ensuring that no VoIP device is left unprotected. Using advanced monitoring and tracking technology, a VoIP device is automatically protected as soon as it is plugged into the network behind a Security Appliance.
- **Full syntax validation of all VoIP signaling packets** - Received signaling packets are fully parsed within SonicOS to ensure they comply with the syntax defined within their associated standard. By performing syntax validation, the Security Appliance can ensure that malformed packets are not permitted to pass through and adversely affect their intended target.
- **Support for dynamic setup and tracking of media streams** - SonicOS tracks each VoIP call from the first signaling packet requesting a call setup, to the point where the call ends. Only based on the successful call progress are additional ports opened (for additional signaling and media exchange) between the calling and called party.

  Media ports that are negotiated as part of the call setup are dynamically assigned by the Security Appliance. Subsequent calls, even between the same parties, uses different ports, thwarting an attacker who might be monitoring specific ports. Required media ports are only opened when the call is fully connected, and are shut down upon call termination. Traffic that tries to use the ports outside of the call is dropped, providing added protection to the VoIP devices behind the Security Appliance.
- **Validation of headers for all media packets** - SonicOS examines and monitors the headers within media packets to allow detection and discarding of out-of-sequence and retransmitted packets (beyond window). Also, by ensuring that a valid header exists, invalid media packets are detected and discarded. By tracking the media streams as well as the signaling, SonicWall provides protection for the entire VoIP session.
- **Configurable inactivity timeouts for signaling and media** - In order to ensure that dropped VoIP connections do not stay open indefinitely, SonicOS monitors the usage of signaling and media streams

associated with a VoIP session. Streams that are idle for more than the configured timeout are shut down to prevent potential security holes.

- **SonicOS allows the administrator to control incoming calls** - By requiring that all incoming calls are authorized and authenticated by the H.323 Gatekeeper or SIP Proxy, SonicOS can block unauthorized and spam calls. This allows the administrator to be sure that the VoIP network is being used only for those calls authorized by the company.
- **Comprehensive monitoring and reporting** - For all supported VoIP protocols, SonicOS offers extensive monitoring and troubleshooting tools:
  - Dynamic live reporting of active VoIP calls, indicating the caller and called parties, and bandwidth used.
  - Audit logs of all VoIP calls, indicating caller and called parties, call duration, and total bandwidth used. Logging of abnormal packets seen (such as a bad response) with details of the parties involved and condition seen.
  - Detailed syslog reports and ViewPoint reports for VoIP signaling and media streams. SonicWall ViewPoint is a Web-based graphical reporting tool that provides detailed and comprehensive reports of your security and network activities based on syslog data streams received from the Security Appliance. Reports can be generated about virtually any aspect of Security Appliance activity, including individual user or group usage patterns and events on specific Security Appliances or groups of Security Appliances, types and times of attacks, resource consumption and constraints.

Supported Interfaces

VoIP devices are supported on the following SonicOS zones:

- Trusted zones (LAN, VPN)
- Untrusted zones (WAN)
- Public zones (DMZ)
- Wireless zones (WLAN)

Supported VoIP Protocols

Topics:

- H.323 on page 770
- SIP on page 771
- SonicWall VoIP Vendor Interoperability on page 771
- CODECs on page 772
- VoIP Protocols on which SonicOS Does Not Perform Deep Packet Inspection on page 772

H.323

SonicOS provides the following support for H.323:

- VoIP devices running all versions of H.323 (currently 1 through to 5) are supported
- Microsoft's LDAP-based Internet Locator Service (ILS)
- Discovery of the Gatekeeper by LAN H.323 terminals using multicast

- Stateful monitoring and processing of Gatekeeper registration, admission, and status (RAS) messages
- Support for H.323 terminals that use encryption for the media streams
- DHCP Option 150. The DHCP Server can be configured to return the address of a VoIP specific TFTP server to DHCP clients
- In addition to H.323 support, SonicOS supports VoIP devices using the following additional ITU standards:
    - T.120 for application sharing, electronic white-boarding, file exchange, and chat
    - H.239 to allow multiple channels for delivering audio, video and data
    - H.281 for Far End Camera Control (FECC)

SIP

SonicOS provides the following support for SIP:

- Base SIP standard (both RFC 2543 and RFC 3261)
- SIP INFO method (RFC 2976)
- Reliability of provisional responses in SIP (RFC 3262)
- SIP specific event notification (RFC 3265)
- SIP UPDATE method (RFC 3311)
- DHCP option for SIP servers (RFC 3361)
- SIP extension for instant messaging (RFC 3428)
- SIP REFER method (RFC 3515)
- Extension to SIP for symmetric response routing (RFC 3581)

SonicWall VoIP Vendor Interoperability

Partial list of devices with which SonicWall VoIP interoperates lists many devices from leading manufacturers with which SonicWall VoIP interoperates.

Partial list of devices with which SonicWall VoIP interoperates

H.323	SIP
Soft-Phones: Avaya Microsoft NetMeeting OpenPhone PolyCom SJLabs SJ Phone Telephones/VideoPhones: Avaya Cisco D-Link PolyCom Sony Gatekeepers: Cisco OpenH323 Gatekeeper Gateway: Cisco	Soft-Phones: Apple iChat Avaya Microsoft MSN Messenger Nortel Multimedia PC Client PingTel Instant Xpressa PolyCom Siemens SCS Client SJLabs SJPhone XTen X-Lite Ubiquity SIP User Agent Telephones/ATAs: Avaya Cisco Grandstream BudgetOne Mitel Packet8 ATA PingTel Xpressa PolyCom PolyCom Pulver Innovations WiSIP SoundPoint SIP Proxies/Services: Cisco SIP Proxy Server Brekeke Software OnDo SIP Proxy Packet8 Siemens SCS SIP Proxy Vonage

CODECs

- **SonicOS supports media streams from any CODEC** - Media streams carry audio and video signals that have been processed by a hardware/software CODEC (COder/DECoder) within the VoIP device. CODECs use coding and compression techniques to reduce the amount of data required to represent audio/video signals. Some examples of CODECs are:
  - H.264, H.263, and H.261 for video
  - MPEG4, G.711, G.722, G.723, G.728, G.729 for audio

VoIP Protocols on which SonicOS Does Not Perform Deep Packet Inspection

SonicWall network Security Appliances do not currently support deep packet inspection for the following protocols; therefore, these protocols should only be used in non-NAT environments.

- Proprietary extensions to H.323 or SIP
- MGCP
- Megaco/H.248
- Cisco Skinny Client Control Protocol (SCCP)
- IP-QSIG
- Proprietary protocols (Mitel's MiNET, 3Com NBX, etc.)

BWM and QoS

One of the greatest challenges for VoIP is ensuring high speech quality over an IP network. IP was designed primarily for asynchronous data traffic, which can tolerate delay. VoIP, however, is very sensitive to delay and packet loss. Managing access and prioritizing traffic are important requirements for ensuring high-quality, real-time VoIP communications.

SonicWall's integrated Bandwidth Management (BWM) and Quality of Service (QoS) features provide the tools for managing the reliability and quality of your VoIP communications.

Quality of Service

QoS encompasses a number of methods intended to provide predictable network behavior and performance. Network predictability is vital to VoIP and other mission critical applications. No amount of bandwidth can provide this sort of predictability, because any amount of bandwidth is ultimately used to its capacity at some point in a network. Only QoS, when configured and implemented correctly, can properly manage traffic, and guarantee the desired levels of network service.

SonicOS includes QoS features that adds the ability to recognize, map, modify and generate the industry-standard 802.1p and Differentiated Services Code Points (DSCP) Class of Service (CoS) designators.

How SonicOS Handles VoIP Calls

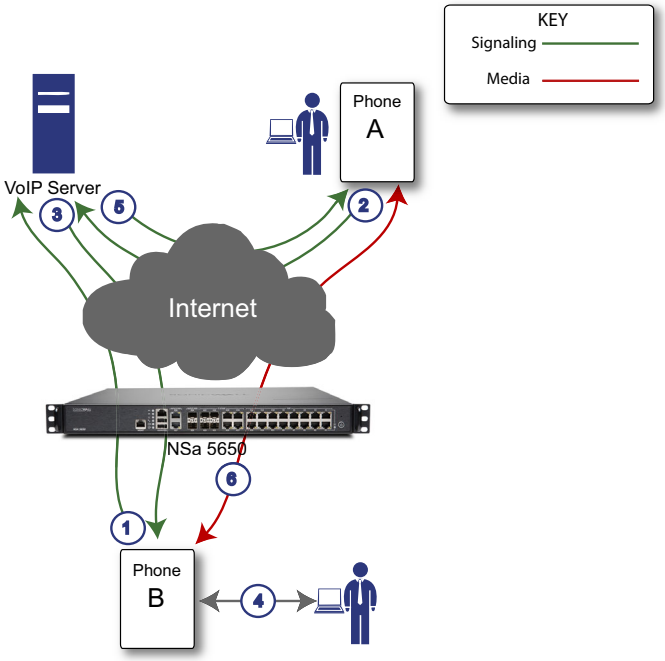
SonicOS provides an efficient and secure solution for all VoIP call scenarios. The following are examples of how SonicOS handles VoIP call flows:

- Incoming Calls on page 773
- Local Calls on page 774

Incoming Calls

Incoming Call Sequence of Events shows the sequence of events that occurs during an incoming call.

Incoming Call Sequence of Events



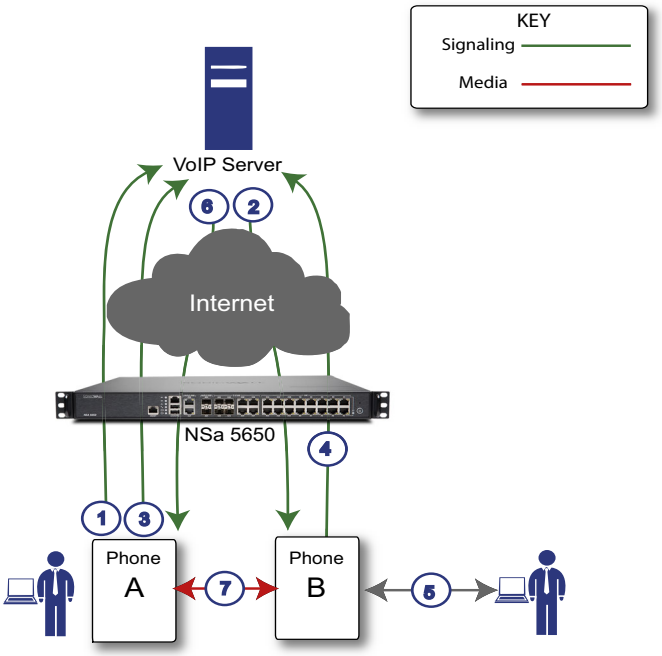
The following describes the sequence of events shown in Incoming Call Sequence of Events:

1. **Phone B registers with VoIP server** - The Security Appliance builds a database of the accessible IP phones behind it by monitoring the outgoing VoIP registration requests. SonicOS translates between phone B's private IP address and the Security Appliance's public IP address used in registration messages. The VoIP server is unaware that phone B is behind a Security Appliance and has a private IP address—it associates phone B with the Security Appliance's public IP address.
2. **Phone A initiates a call to phone B** - Phone A initiates a call to phone B using a phone number or alias. When sending this information to the VoIP server, it also provides details about the media types and formats it can support as well as the corresponding IP addresses and ports.
3. **VoIP Server validates the call request and sends the request to phone B** - The VoIP server sends the call request to the Security Appliance's public IP address. When it reaches the Security Appliance, SonicOS validates the source and content of the request. The Security Appliance then determines phone B's private IP address.
4. **Phone B rings and is answered** - When phone B is answered, it returns information to the VoIP server for the media types and formats it supports as well as the corresponding IP addresses and ports. SonicOS translates this private IP information to use the Security Appliance's public IP address for messages to the VoIP server.
5. **VoIP server returns phone B media IP information to phone A** - Phone A now has enough information to begin exchanging media with Phone B. Phone A does not know that Phone B is behind a Security Appliance, as it was given the public address of the Security Appliance by the VoIP Server.
6. **Phone A and phone B exchange audio/video/data through the VoIP server** - Using the internal database, SonicOS ensures that media comes from only Phone A and is only using the specific media streams permitted by Phone B.

Local Calls

Local VoIP Call Sequence of Events shows the sequence of events that occurs during a local VoIP call.

Local VoIP Call Sequence of Events



The following describes the sequence of events shown in Local VoIP Call Sequence of Events:

1 **Phones A and B register with VoIP server** - The Security Appliance builds a database of the accessible IP phones behind it by monitoring the outgoing VoIP registration requests. SonicOS translates between the phones' private IP addresses and the Security Appliance's public IP address. The VoIP server is unaware that the phones are behind a Security Appliance. It associates the same IP address for both phones, but different port numbers.

2 **Phone A initiates a call to phone B by sending a request to the VoIP server -** Even though they are behind the same Security Appliance, phone A does not know Phone B's IP address. Phone A initiates a call to phone B using a phone number or alias.

3 **VoIP Server validates the call request and sends the request to phone B -** The VoIP server sends the call request to the Security Appliance's public IP address. The Security Appliance then determines phone B's private IP address.

4 **Phone B rings and is answered** - When phone B is answered, the Security Appliance translate its private IP information to use the Security Appliance's public IP address for messages to the VoIP server.

5 **VoIP Server returns phone B media IP information to phone A -** Both the called and calling party information within the messages are translated by SonicOS back to the private addresses and ports for phone A and phone B.

6 **Phone A and phone B directly exchange audio/video/data -** The Security Appliance routes traffic directly between the two phones over the LAN. Directly connecting the two phones reduces the bandwidth requirements for transmitting data to the VoIP server and eliminates the need for the Security Appliance to perform address translation.

52

Configuring SonicWall VoIP Features

Topics:

- Configuration Tasks on page 776
  - Configuring VoIP on page 776
  - Configuring VoIP Logging on page 782

Configuration Tasks

Configuring the SonicWall Security Appliance for VoIP deployments builds on your basic network configuration in the SonicWall Management Interface. This section assumes the Security Appliance is configured for your network environment.

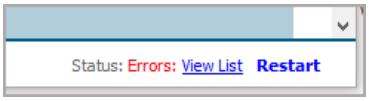
NOTE: For general information on VoIP, see About VoIP on page 765.

Topics:

- Configuring VoIP on page 776
- Configuring VoIP Logging on page 782

Configuring VoIP

IMPORTANT: If **Wireless-Controller-Only** mode has been selected for **Wireless LAN Controller**, any attempt to enable SIP and/or H.323 options displays an error message in the lower right corner of the browser window:



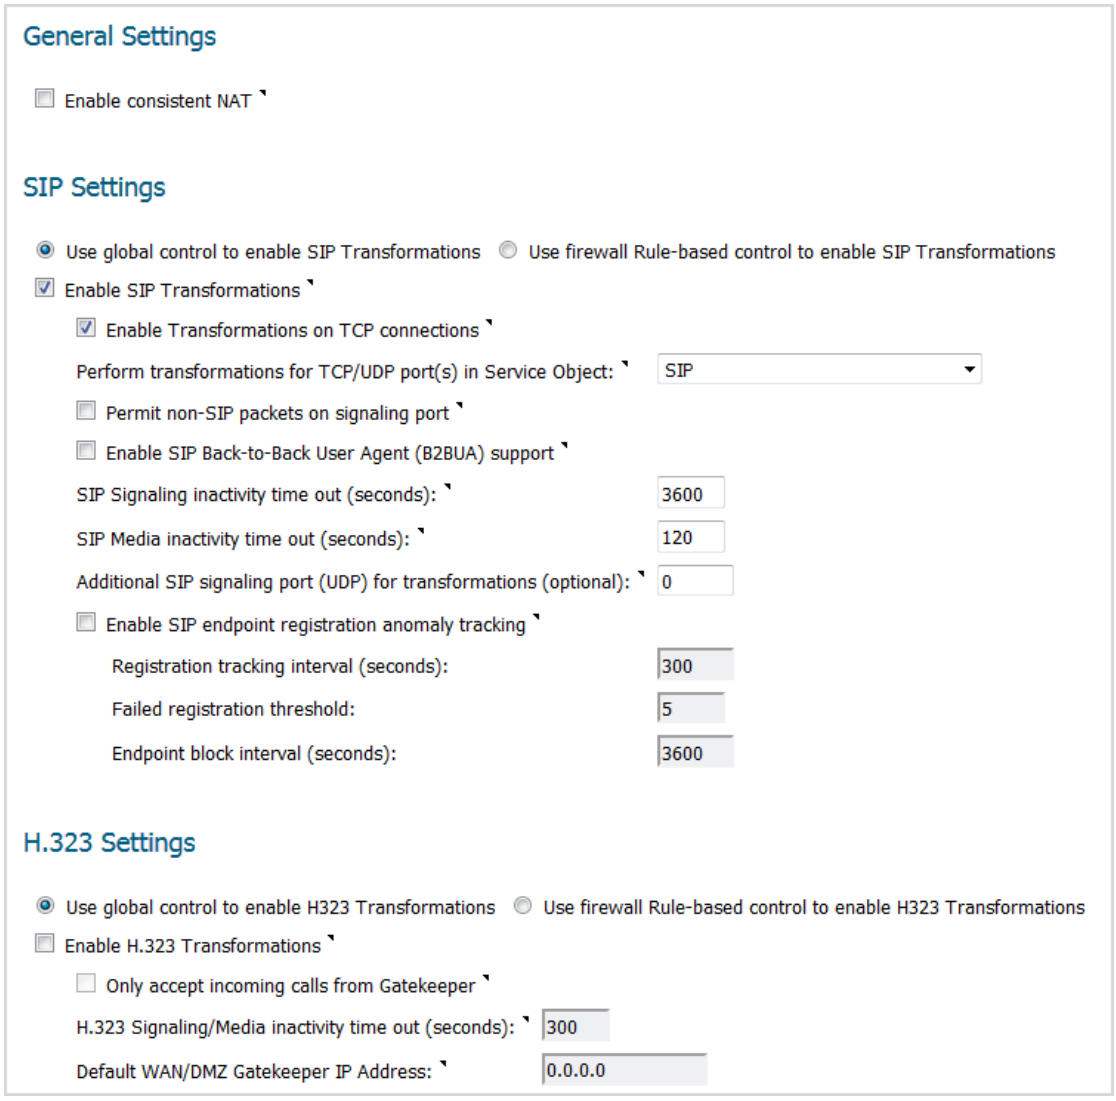
Clicking the `View List` link displays an error log:



You configure VoIP through settings on **MANAGE | System Setup > VOIP**. This page is divided into three sections:

- **General Settings**
- **SIP Settings**

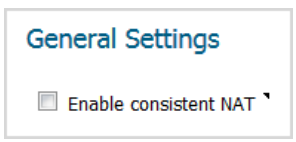
- **H.323 Settings**



Topics:

- General Settings on page 777
- SIP Settings on page 778
- H.323 Settings on page 780

General Settings



There is one option under **General Settings**: **Enable Consistent NAT**.

Consistent NAT enhances standard NAT policy to provide greater compatibility with peer-to-peer applications that require a consistent IP address to connect to, such as VoIP. Consistent NAT uses an MD5 hashing method to

consistently assign the same mapped public IP address and UDP Port pair to each internal private IP address and port pair.

For example, NAT could translate the private (LAN) IP address and port pairs, `192.116.168.10/50650` and `192.116.168.20/50655` into public (WAN) IP/port pairs, as shown in IP address and port pairs:

IP address and port pairs

Private IP/Port	Translated Public IP/Port
192.116.168.10/50650	64.41.140.167/40004
192.116.168.20/50655	64.41.140.167/40745

With Consistent NAT enabled, all subsequent requests from either host `192.116.168.10` or `192.116.168.20` using the same ports illustrated in IP address and port pairs result in using the same translated address and port pairs. Without Consistent NAT, the port and possibly the IP address change with every request.

NOTE: Enabling Consistent NAT causes a slight decrease in overall security, because of the increased predictability of the address and port pairs. Most UDP-based applications are compatible with traditional NAT. Therefore, do not enable Consistent NAT unless your network uses applications that require it.

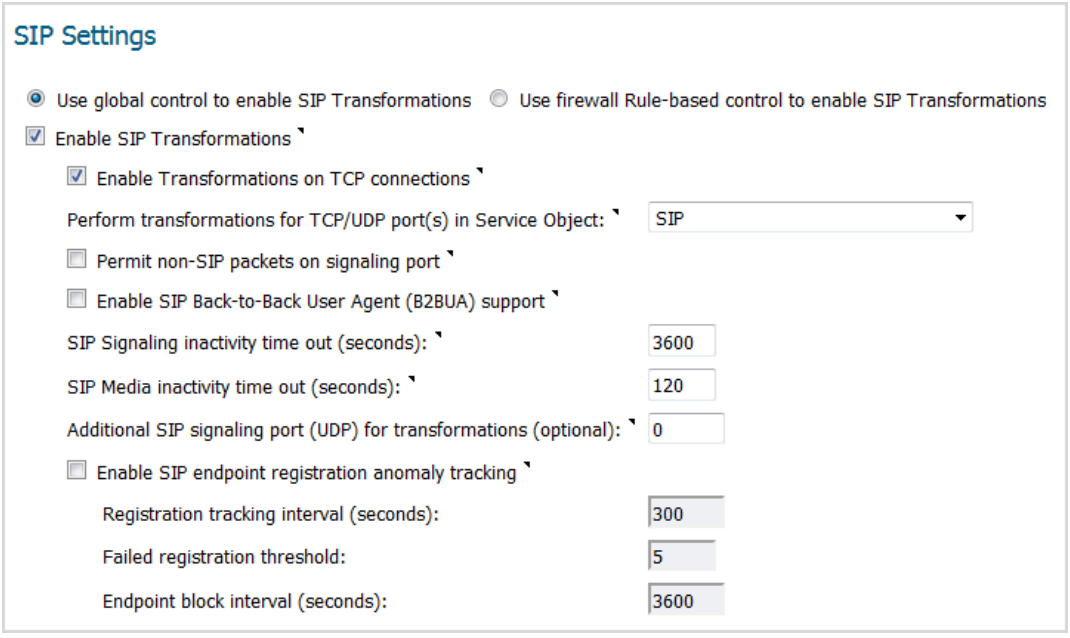
IMPORTANT: For Consistent NAT to work properly, the minimum time interval between calls must be at least 200 msec.

Enabling Consistent NAT

To enable consistent NAT:

1. Select the **Enable Consistent NAT** option. This option is not selected by default.
2. Click **Accept**.

SIP Settings



By default, SIP clients use their private IP address in the SIP (Session Initiation Protocol) Session Definition Protocol (SDP) messages that are sent to the SIP proxy. If your SIP proxy is located on the public (WAN) side of

the firewall and the SIP clients are located on the private (LAN) side of the firewall, the SDP messages are not translated and the SIP proxy cannot reach the SIP clients.

Enabling SIP

To enable SIP:

1 Navigate to **MANAGE | System Setup > VOIP**.

2 In the **SIP Settings** section, choose whether to enable SIP transformation globally or by firewall rule:

- **Use global control to enable SIP Transformations**. This option is selected by default.
- **Use firewall Rule-based control to enable SIP Transformations**. Be sure to configure a firewall rule to control SIP transformations as described in *SonicOS Policies*.

3 If you are not configuring SIP transformations, go to Step 12.

4 **Enable SIP Transformations** is not selected by default. Select this option to:

- Transform SIP messages between LAN (trusted) and WAN/DMZ (untrusted).

  You need to check this setting when you want the Security Appliance to do the SIP transformation. If your SIP proxy is located on the public (WAN) side of the Security Appliance and SIP clients are on the LAN side, the SIP clients by default embed/use their private IP address in the SIP/Session Definition Protocol (SDP) messages that are sent to the SIP proxy; hence, these messages are not changed and the SIP proxy does not know how to get back to the client behind the Security Appliance.
- Enable the Security Appliance to go through each SIP message and change the private IP address and assigned port.
- Control and open up the RTP/RTCP ports that need to be opened for SIP session calls to happen.

NAT translates Layer 3 addresses, but not Layer 7 SIP/SDP addresses, which is why you need to select **Enable SIP Transformations** to transform the SIP messages.

TIP: In general, you should select **Enable SIP Transformations** unless there is another NAT traversal solution that requires this feature to be turned off. SIP Transformations works in bi-directional mode, meaning messages are transformed going from LAN to WAN and vice versa.

When **Enable SIP Transformations** is selected, the other options become available.

5 To perform SIP transformations on TCP-based SIP sessions, select **Enable SIP Transformation on TCP connections**. This option is selected by default.

6 Select a Service Object from Perform transformations to **TCP/UDP port(s) in Service Object**. The default is **SIP**.

7 Selecting **Permit non-SIP packets on signaling port** enables applications such as Apple iChat and MSN Messenger, which use the SIP signaling port for additional proprietary messages. This option is not selected by default.

IMPORTANT: Enabling this checkbox might open your network to malicious attacks caused by malformed or invalid SIP traffic.

8 If the SIP Proxy Server is being used as a B2BUA, enable the **Enable SIP Back-to-Back User Agent (B2BUA) support** setting. This option is disabled by default and should be enabled only when the Security Appliance can see both legs of a voice call (for example, when a phone on the LAN calls another phone on the LAN).

TIP: If there is no possibility of the firewall seeing both legs of voice calls (for example, when calls are only made to and received from phones on the WAN), the **Enable SIP Back-to-Back User Agent (B2BUA) support** setting should be disabled to avoid unnecessary CPU usage.

9 Use the **SIP Signaling inactivity time out (seconds)** and **SIP Media inactivity time out (seconds)** options to define the amount of time a call can be idle (no traffic exchanged) before the firewall blocks further traffic. A call goes idle when placed on hold. Specify the maximum idle time when:

- There is no signaling (control) message being exchanged in **SIP Signaling inactivity time out.** The minimum time is 30 seconds, the maximum time is 1000000 seconds (~1.2 days) and the default is **3600** seconds (60 minutes).
- No media (for example, audio or video) packets are being exchanged in the **SIP Media inactivity time out**. The minimum time is 30 seconds, the maximum time is 3600 seconds (1 hour), and the default time is **120** seconds (2 minutes).

10 Use the **Additional SIP signaling port (UDP) for transformations** setting to specify a non-standard UDP port to carry SIP signaling traffic. Normally, SIP signaling traffic is carried on UDP port 5060. However, a number of commercial VOIP services use different ports, such as 1560. When this setting is non zero (**0** is the default; the maximum value is 65535), the Security Appliance performs SIP transformation on these non-standard ports.

TIP: Vonage's VoIP service uses UDP port 5061.

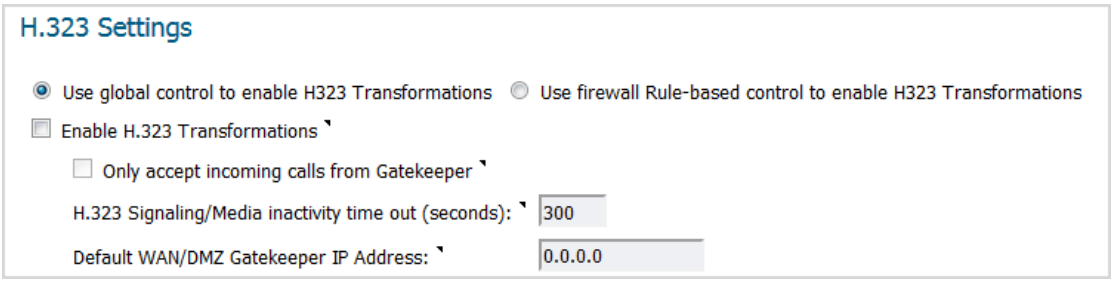
11 To track SIP endpoint registration anomalies, select the **Enable SIP endpoint registration anomaly tracking** option. This option is not selected by default. When it is selected, these options become available:

- **Registration tracking interval (seconds)** – Specify the interval between checking for anomalies. The default is **300** seconds (5 minutes).
- **Failed registration threshold** – Specify the number of failed registrations before checking for anomalies. The default is **5** failures.
- **Endpoint block interval (seconds)** – The default is **3600** (60 minutes).

12 Either:

- Click **Accept**.
- Go to H.323 Settings on page 780.

H.323 Settings



Configuring H.323 Settings

To configure H.323 settings:

1 Navigate to **MANAGE | System Setup > VOIP | H.323 Settings**.

2 Choose whether to enable H.323 transformation globally or by firewall rule:

- **Use global control to enable H.323 Transformations**. This option is selected by default.

- **Use firewall Rule-based control to enable H.323 Transformations**. Be sure to configure a firewall rule to control H.323 transformations as described in *SonicOS Policies*.

3 If you are not configuring H.323 transformations, go to Step 5.

1 Select **Enable H.323 Transformation** to allow stateful H.323 protocol-aware packet content inspection and modification by the firewall. This option is disabled by default. When the option is selected, the other H.323 options become active.

The firewall performs any dynamic IP address and transport port mapping within the H.323 packets, which is necessary for communication between H.323 parties in trusted and untrusted networks/zones.

Disable the **Enable H.323 Transformation** to bypass the H.323 specific processing performed by the firewall.

2 Select **Only accept incoming calls from Gatekeeper** to ensure all incoming calls go through the Gatekeeper for authentication. The Gatekeeper refuses calls that fail authentication.

3 In the **H.323 Signaling/Media inactivity time out (seconds)** field, specify the amount of time a call can be idle before the firewall blocks further traffic. A call goes idle when placed on hold. The default time is **300** seconds (5 minutes), the minimum time is 60 seconds (1 minute), and the maximum time is 122400 seconds (34 hours).

4 The **Default WAN/DMZ Gatekeeper IP Address** field has a default value of **0.0.0.0**. Enter the default H.323 Gatekeeper IP address in this field to allow LAN-based H.323 devices to discover the Gatekeeper using the multicast address `225.0.1.41`. If you do not enter an IP address, multicast discovery messages from LAN-based H.323 devices go through the configured multicast handling.

5 Click **Accept**.

Topics:

- Configuring Bandwidth on the WAN Interface on page 781
- Configuring VoIP Access Rules on page 781

Configuring Bandwidth on the WAN Interface

NOTE: For information on Bandwidth Management (BWM) and configuring BWM on the WAN interface, see *SonicOS Policies*.

Configuring VoIP Access Rules

By default, stateful packet inspection on the firewall allows all communication from the LAN to the Internet and blocks all traffic to the LAN from the Internet. Additional network access rules can be defined to extend or override the default access rules.

If you are defining VoIP access for client to use a VoIP service provider from the WAN, you configure network access rules between source and destination interface or zones to enable clients behind the firewall to send and receive VoIP calls.

TIP: Although custom rules can be created that allow inbound IP traffic, the firewall does not disable protection from Denial of Service attacks, such as the SYN Flood and Ping of Death attacks.

NOTE: You must select Bandwidth Management on **MANAGE | System Setup| Network > Interfaces** for the **WAN** interface before you can configure bandwidth management for network access rules.

For how to add access rules for VoIP traffic on the SonicWall Security Appliance, see *SonicOS Policies*.

Configuring VoIP Logging

You can enable the logging of VoIP events, which are displayed on **INVESTIGATE | Logs | Event Logs**. To enable logging of VoIP, see *SonicOS Investigate*.

Part 10

System Setup | Virtual Assist

- Configuring Virtual Assist

53

Configuring Virtual Assist

Topics:

- About Virtual Assist on page 784
- Maximizing Virtual Assist Flexibility on page 785

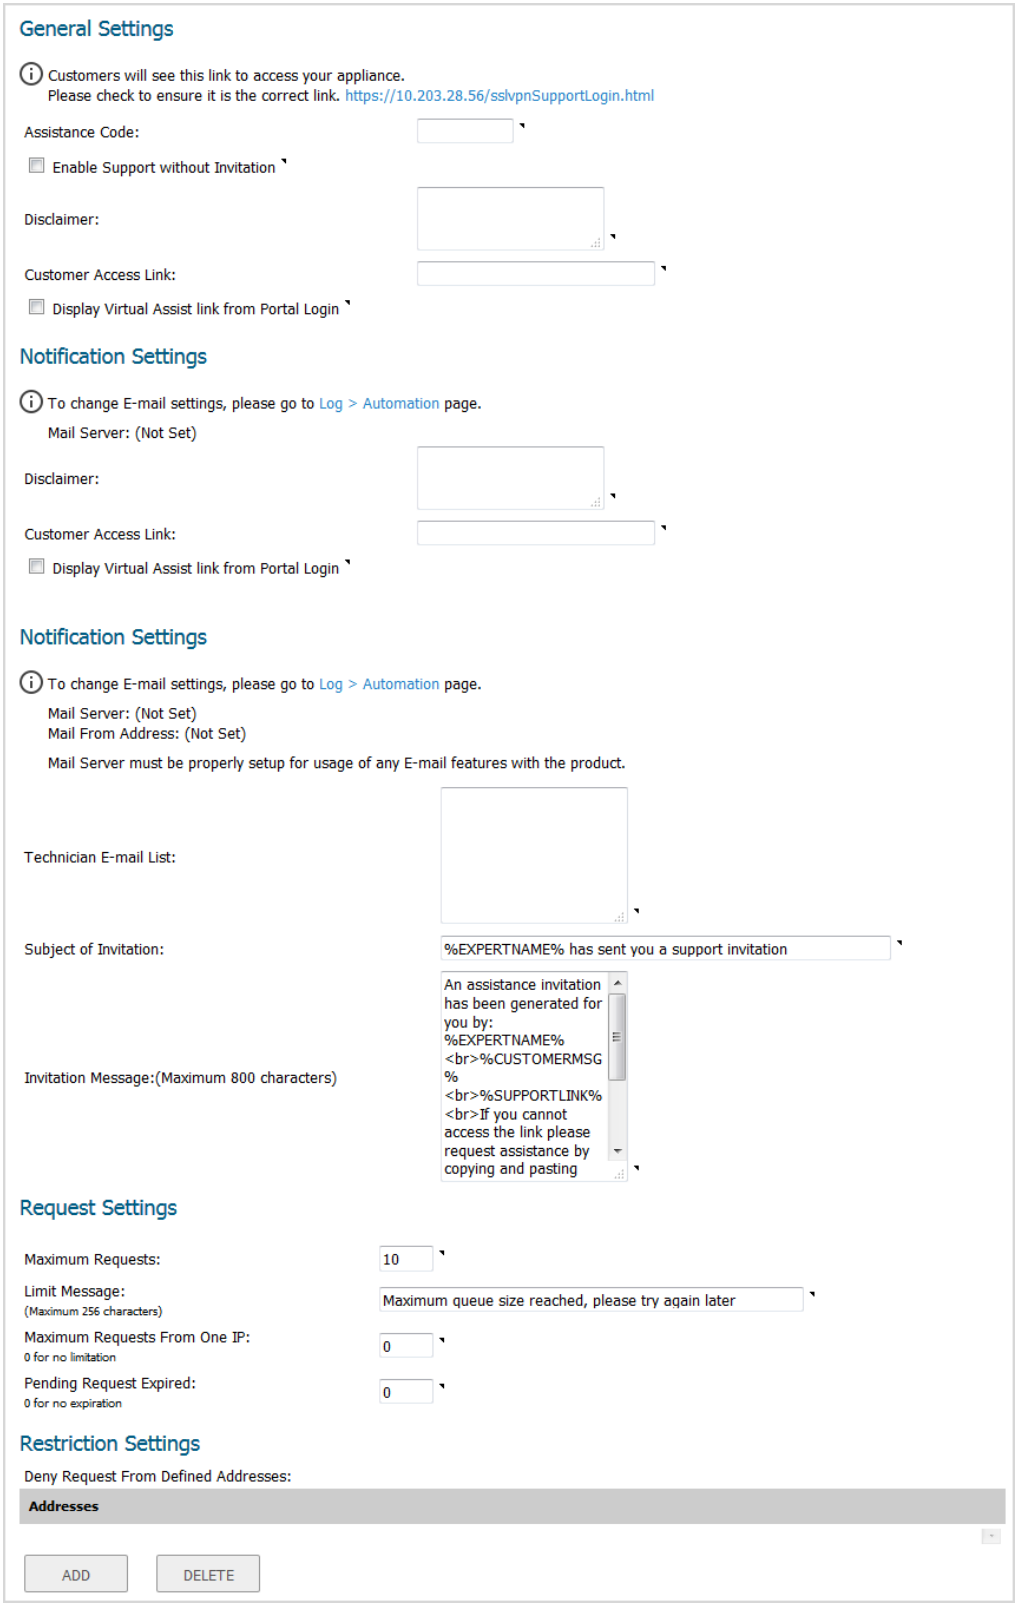
About Virtual Assist

Virtual Assist allows you to support customer technical issues without having to be on-site with the customer. This capability serves as an immense time-saver for support personnel, while adding flexibility in how they can respond to support needs. You can allow or invite customers to join a queue to receive support, then virtually assist each customer by remotely taking control of a customer's computer to diagnose and remedy technical issues.

NOTE: The technician or administrator providing Virtual Assist must be located inside the local network of the SonicWall Security Appliance.

Maximizing Virtual Assist Flexibility

You control the Virtual Assist through settings on **System Setup > Virtual Assist**.



Configuring Virtual Assist

To maximize the flexibility of the Virtual Assist feature, you should take the time to properly adjust all of the settings.

Topics:

- Providing Access to Users on page 786
- Customizing Notifications on page 787
- Managing Requests on page 788
- Blocking Requests from Certain IP Addresses on page 789

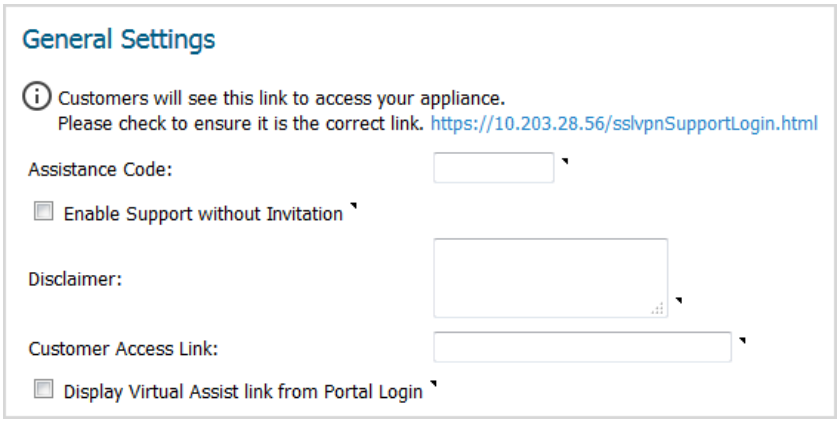
Providing Access to Users

You need to decide how to provide access for customers to gain support through Virtual Assist:

- Enable virtual assist support without the need for an invitation.
- P By setting a global assistance code for customers, you can restrict who enters the system to request help. The code can be a maximum of eight (8) characters, and can be entered in the Assistance Code field. Customers receive the code through an email provided by the technician or an administrator.

To provide access to users:

1 Navigate to **MANAGE | System Setup > Virtual Assist**.



2 To provide a global code for customers to enter before being able to request help, enter up to 8 alphanumeric characters in the **Assistance Code** field. To indicate a code is not required, leave this field blank.

TIP: The Assistance Code can be used to restrict who can enter the system to request help.

3 To allow customers to request help through the support login web page without being invited by a technician:

a Leave the **Assistance Code** field blank.

b Select **Enable Support without Invitation**.

NOTE: If this option is not selected, a customer can receive assistance only by being invited through email by a technician. Select this option for customers to be able to request help from the login page.

4 To create a written message that customers must read and agree to before receiving support, enter the disclaimer in the **Disclaimer** field.

5 To give access to your SSL VPN Security Appliance from outside your network, enter a URL in the **Customer Access Link** field. If you leave this field blank, the support invitation to customers uses the URL the technician uses to access the Security Appliance.

> **TIP:** Configure this option if the SSL VPN Security Appliance is accessed with a different URL from outside your network.

6 To redirect customers to the support login page if they navigate to the technician login page, select **Display Virtual Assist link from Portal Login**.

7 Click **Accept**.

8 To ensure the access link your customer sees is correct, click the link in the **General Settings** information note. The access link you configured in Step 5 displays; for example.



Customizing Notifications

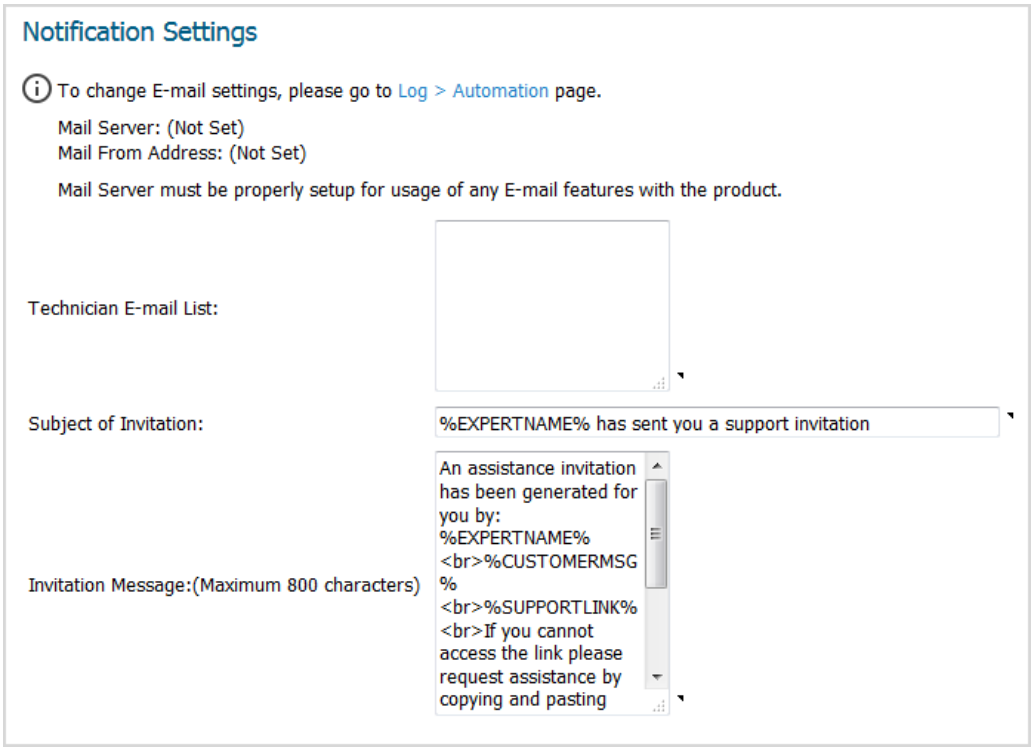
In the **Notification Settings** section, you customize various aspects of the invitation and technician notification.

To customize invitation and technician notification:

> **IMPORTANT:** Before configuring the notification settings, configure the email server and email address on **MANAGE | Logs & Reporting | Log Settings | Automation**; to display this page quickly, click the link in the informational note in the **Notification Settings** section. For information about setting up the email server, see *SonicOS Logs and Reporting*.

1 Navigate to **MANAGE | System Setup > Virtual Assist.**

2 Scroll to **Notification Settings**.



3 Create a list of technician email addresses to receive a notification email when an uninvited customer enters the support queue in the **Technician E-mail List** field. Add up to 10 emails to this list, with each separated by a semicolon.

4 To customize the subject line of support invitation emails, enter the desired text in the **Subject of Invitation** field, using the variables listed in Variables. A sample invitation subject line is provided.

Variables

For the	Use
Technician Name	%EXPERTNAME%
Customer Message in the Invitation	%CUSTOMERMSG%
Link for Support	%SUPPORTLINK%
Link to SSL-VPN	%ACCESSLINK%

5 To customize the body of the invitation email, enter the desired text in the **Invitation Message** field, using the variables listed in Variables. The message can contain a maximum of 800 characters. A sample invitation subject is provided.

6 Click **ACCEPT**.

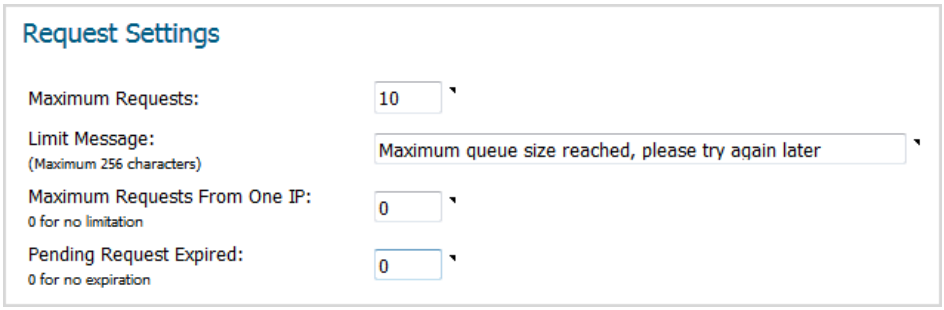
Managing Requests

You manage and limit support requests in the **Request Settings** section.

To manage and limit support requests:

1 Navigate to **MANAGE | System Setup > Virtual Assist**.

2 Scroll to **Request Settings**.



3 To limit the number of customers that can be awaiting assistance in the queue at one time, enter a limit in the **Maximum Requests** field. When this limit is reached, new requests are blocked. The default queue size is **10** requests.

4 To display a message to customers when there are currently no available spots in the queue, as the maximum requests limit has been reached, enter the message in the **Limit Message** field. You can create a message up to 256 characters. A sample message is given.

5 To limit the number of requests coming from a single IP, enter the limit in the **Maximum Requests From One IP** field. This prevents the same customer from requesting Virtual Assist support multiple times at once and thus be put in the queue multiple times. Enter **0** (default) for no limitation.

6 To avoid customers waiting indefinitely for Virtual Assist support during high-volume periods, you can set limit how long a customer can remain in the queue without receiving support by entering the limit, in minutes, in the **Pending Request Expired** field. Enter **0** (default) if you do not wish to set a limit.

7 Click **ACCEPT**.

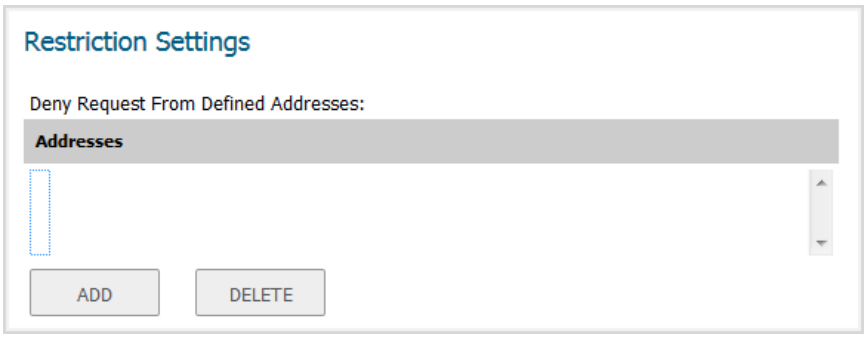
Blocking Requests from Certain IP Addresses

If you encounter requests from unwanted or illegitimate sources, you can block requests from defined IP addresses.

To block requests from IP addresses:

1 Navigate to **MANAGE | System Setup > Virtual Assist**.

2 Scroll to **Restriction Settings**.



3 Click **ADD**. The **Admin Address** dialog displays.



4 Select the type of source address from **Source Address Type**:

- **IP Address** – default
- **IP Network** – the options change; go to Step 7.



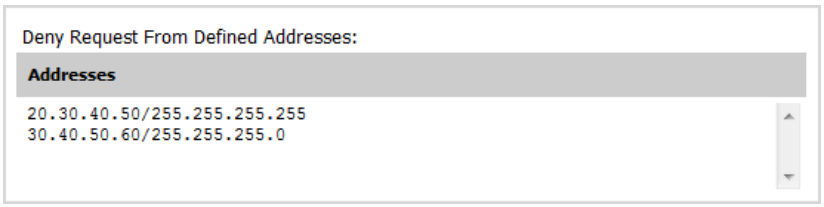
5 Enter the IP address to be blocked in the **IP Address** field.

6 Go to Step 9.

7 Enter the network address to be blocked in the **Network Address** field.

8 Enter the subnet mask for the address in the **Subnet Mask** field.

9 Click **OK**. Entry is added to the **Deny Request From Defined Addresses** table.



10 Click **ACCEPT**.

Deleting Blocked Addresses

To delete entries from the Deny Request From Defined Addresses field:

1 Navigate to **MANAGE | System Setup > Virtual Assist**.

2 Scroll to **Restriction Settings**.

3 Select the entry to delete.

4 Click **DELETE**.

Part 11

System Setup | Appendices

- Configuring Open Authentication, Social Login, and LHM
- BGP Advanced Routing
- IPv6
- SonicWall Support

A

Configuring Open Authentication, Social Login, and LHM

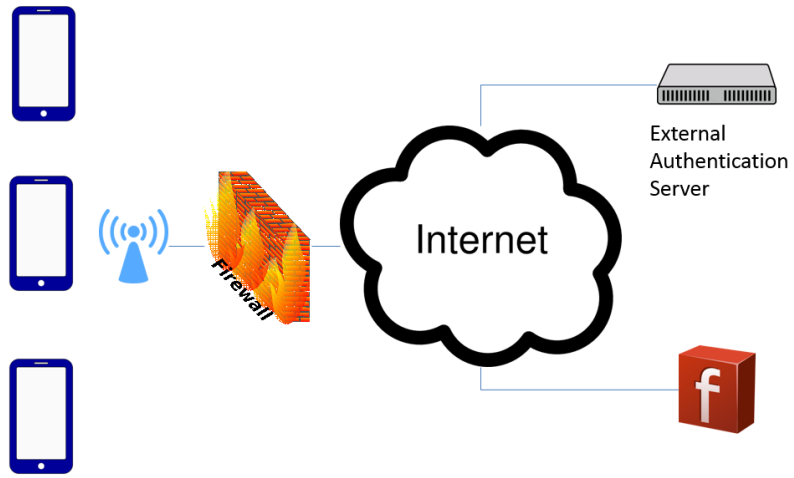
Topics:

- About OAuth and Social Login on page 792
- About Lightweight Hotspot Messaging (LHM) on page 796
- Configuring Facebook for Social Login on page 798
- Configuring Open Authentication and Social Login on page 800
- Verifying the Social Login Configuration on page 807
- Using Social Login, LHM, and ABE on page 807

About OAuth and Social Login

Social Login is a form of single sign-on authentication that utilizes existing user credentials from social networking services such as Facebook, Twitter, or Google+ to then sign in to a third-party website instead of creating a new login account specifically for that website. The Open Authentication (OAuth) Social Login feature can be used with guest service on wireless zones, LAN zones, or DMZ zones using pass-through authentication; see External Authentication Server Login Topology. Pass-though authentication is a method of performing authentication to a domain controller that resides within a trusted domain. Wireless guest services are widely used in public WiFi hot spots and corporate WiFi services set up for guests.

External Authentication Server Login Topology



Topics:

- What are OAuth and Social Login? on page 793
- Benefits of OAuth and Social Login on page 793
- How Do OAuth and Social Login Work? on page 794
- Supported Platforms on page 795

What are OAuth and Social Login?

OAuth is an open standard for authorization. OAuth provides client applications “secure delegated access” to server resources on behalf of a resource owner, and specifies a process for owners to authorize third-party access to their server resources without sharing their credentials.

Social Login, also known as social sign-in, is a form of single sign-on (SSO) using existing login information from a social networking service such as Facebook, Twitter, or Google+ to sign into a third-party website instead of creating a new login account specifically for that website.

Benefits of OAuth and Social Login

Topics:

- OAuth
- Social Login

OAuth

OAuth is a popular mechanism that assists users in sharing data between applications. You can take advantage of OAuth by using it as a login provider for your web application.

Other Advantages

- Limiting customer profiles on the net
- Fewer passwords to track
- Not required to submit a password where trust might be an issue
- You can still prevent access from the OAuth provider
- Lower risk of ID theft. Authentication is assumed by the provider
- Lower risk of bug failure with authentication using previously proven APIs
- Less storage requirement on your data servers

Disadvantages

- You cannot tailor user profiles for your own applications
- User confusion in creating accounts with OAuth providers when they do not have existing accounts

Social Login

Social login is designed to simplify the login process and to realize a higher conversion rate for registrations.

Other Advantages

- Quick registration
- Remember fewer logins
- Target-rich content
- Use of multiple identities
- Collection of visitor data
- Detailed or personalized user experience
- Familiar login environments
- Fewer failed logins
- Ease of use for mobile

Disadvantages

- Low trust level
- Non-Social users excluded
- Data accuracy can be falsified
- Blocked content from Social networks
- Security issues

How Do OAuth and Social Login Work?

The Open Authentication (OAuth) and Social Login features can both be used with internal wireless services and SonicPoints as a wireless zone guest service. Guests can log in to the Internet using your company's corporate WiFi. Wireless guest services are widely used in public WiFi hot spots and corporate WiFi services set up for guests.

Both OAuth and Social Login use wireless guest services that include Internet access and can be configured to use either or both of these methods of connection:

- No Redirect on page 794
- Redirect to a Landing Page on page 795

No Redirect

No Redirect provides open Internet access to guests with no required encryption, so guests are allowed to connect to the provided WiFi freely.

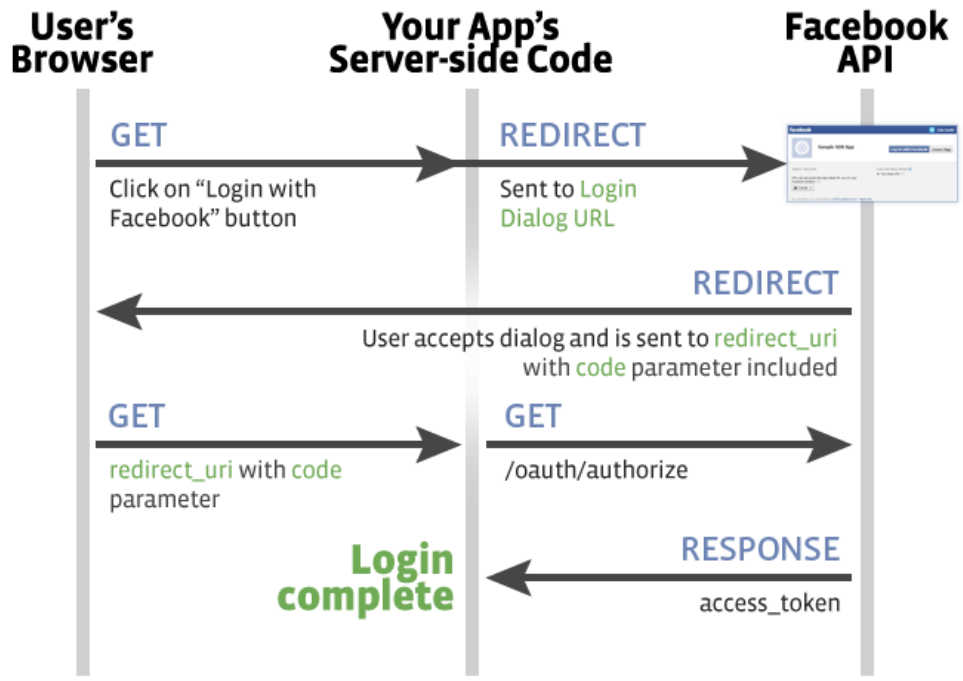
No redirect can also provide WPA/WEP passphrase or password access where guests would need a passcode to use the available WiFi. The passcode could be provided through other means, such as on a receipt.

Redirect to a Landing Page

A landing web page provides the most widely used hot spot access. While the layer 2 WiFi access is open, guests are directed to a landing web page when accessing the first layer; see Oauth flow. Some other redirect access options include:

- No authentication on the landing page
- Guests can create a new login account and then sign in with it
- Guests can sign up using a code sent to them by SMS to a mobile phone, email, or other method
- Scanning a QR code with a mobile app
- Using a social login

Oauth flow



Supported Platforms

Open Authentication and Social Login is supported on SonicWall firewalls:

- Running SonicOS 6.2.7 and higher
- Under GMS Management running GMS 8.3

Requirements for Development and Production

- A Facebook account
    - Enable Facebook For Developers

- External server
  - Public accessible
  - Has a domain name
  - PHP support
  - SSL Certificate
- SonicWall firewall
  - Can be reached by the external server (by IP or FQDN)
  - Wireless (internal or SonicPoint)

About Lightweight Hotspot Messaging (LHM)

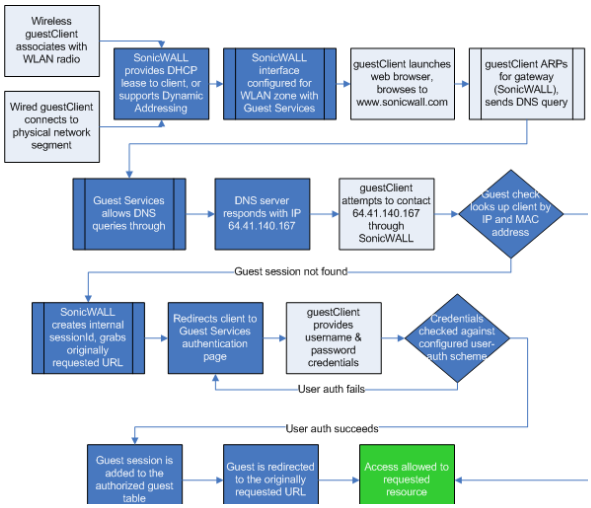
Lightweight Hotspot Messaging (LHM) leverages the SonicWall Guest Service model, wherein users can be classified and authorized for differentiated network access through a SonicWall Security Appliance. For example, the SonicWall can be configured such that any user connecting through an interface belonging to a guest-services-enabled WLAN (wireless LAN) Zone only has access to the Internet (Untrusted network), but does not have access to the LAN (Trusted network). This allows a single firewall to offer simultaneous access to trusted and guest users.

LHM extends the Guest Services model by breaking apart the authentication and authorization processes, thereby allowing the authentication to occur external to the SonicWall. This allows for extensive customization of the authentication interface, and also allows for any kind of imaginable authentication scheme to be used.

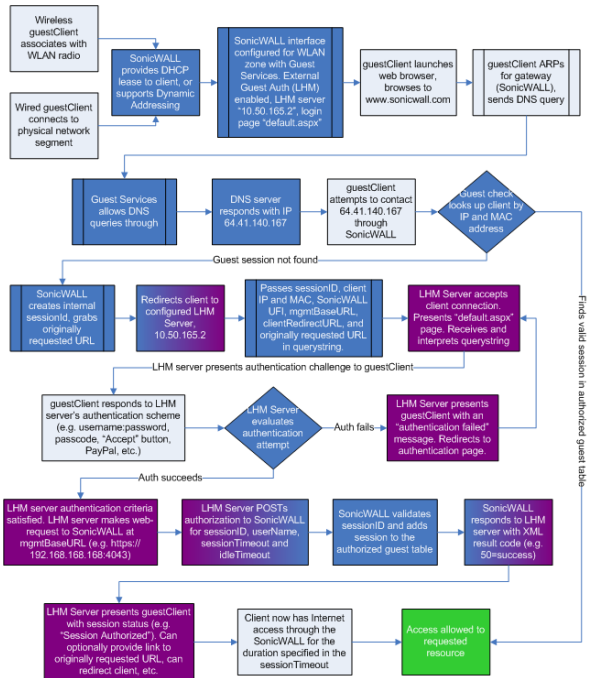
A side by side view of the original Guest Services authorization flow and the LHM authorization flow is shown in Comparison of authorization flows:

Comparison of authorization flows

Original Guest Services Authorization Flow

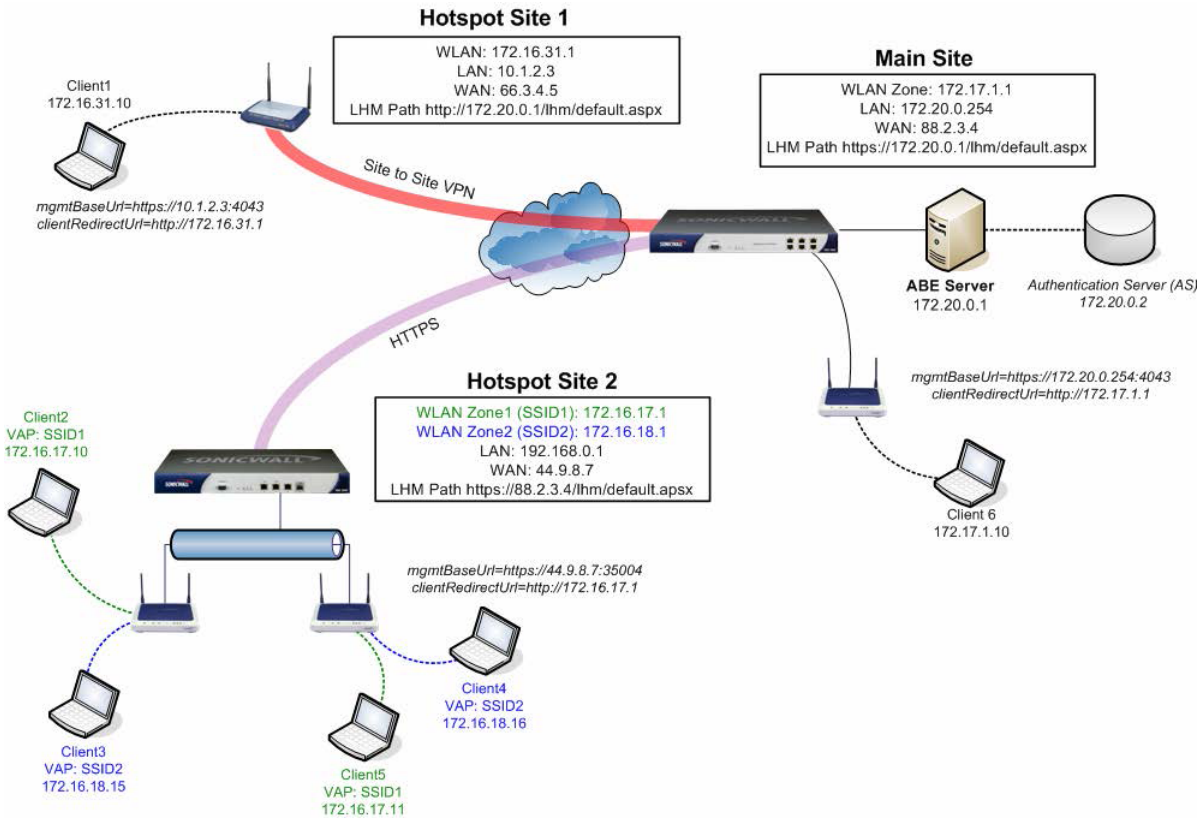


LHM Authorization Flow



LHM defines the method and syntax for communications between a SonicWall wireless access device (such as a SOHO W firewall, a TZ Wireless series firewall, or a SonicPoint with a governing SonicWall Security Appliance) and an Authentication Back-End (ABE) for authenticating Hotspot users and providing them parametrically bound network access. LHM configuration example depicts a generic configuration.

LHM configuration example



LHM allows network operators to provide centralized management of multiple Hotspot locations by providing an interface between SonicWall's Wireless Guest Services and any existing ABE. LHM is an adaptation of the generalized WISPr and GIS specifications.

LHM was designed to satisfy the requirements of a particularly common operational environment rather than a broad set of environments. Specifically, LHM allows for Hotspot user-management and authentication to occur entirely on the network operator's ABE, supporting any method of account creation and management, and any extent of site customization and branding. This approach enables integration into any existing environment without dependencies upon particular billing, accounting or database systems, and also provides the network operator with unrestricted control of the site's design, from look-and-feel to redirection.

To provide authentication of IPv4 Traffic and IPv6 Traffic, SonicOS supports two external Lightweight Hotspot Messaging (LHM) web servers, one for IPv4 and the other for IPv6. Different traffic types redirect to their corresponding LHM server.

Configuring Facebook for Social Login

Topics:

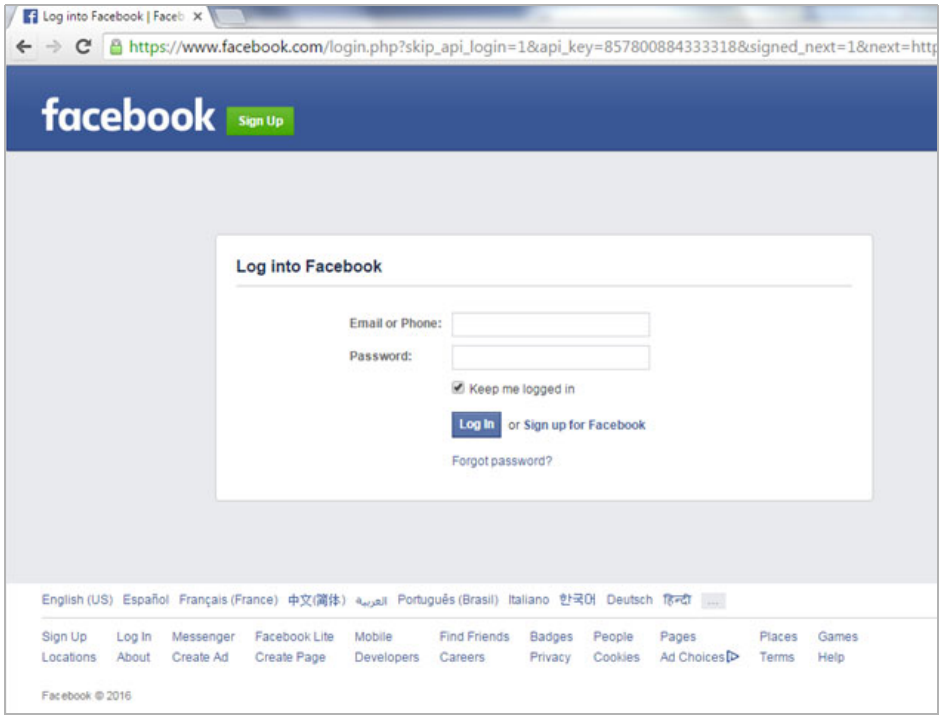
- Facebook Settings on page 799
- Client OAuth Settings on page 800
- Guest Status (demo) on page 800

Facebook Settings

To login to Facebook for Developers:

1 Open a Web browser

2 Log in to your Facebook for Developers account at https://developers.facebook.com/.

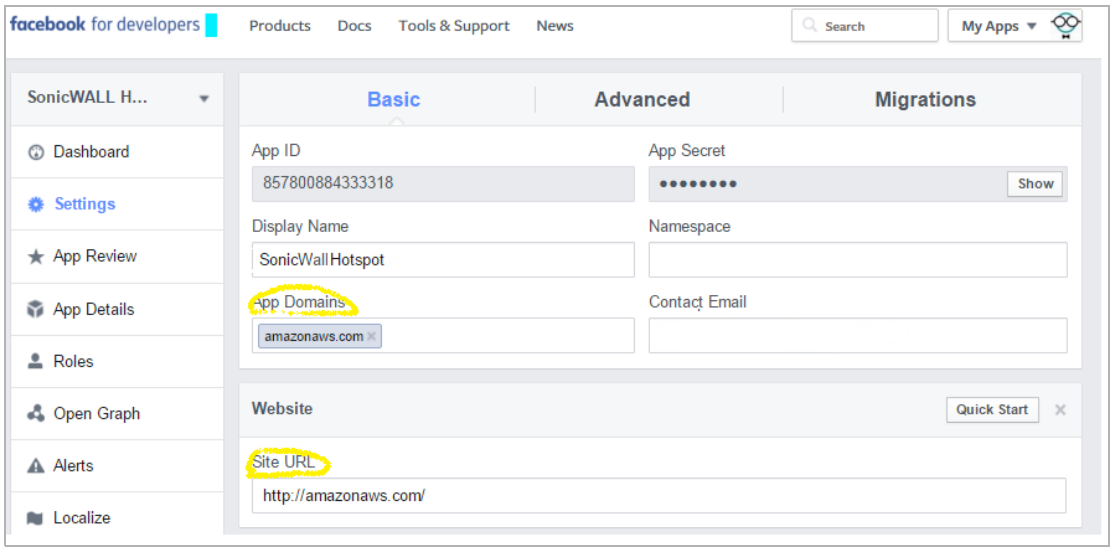


3 Complete the login process or sign up for a new developer's account.

4 Click **Settings** in the left column.

See Example of settings for Facebook for developers to fill the form, but adjust the Facebook **Settings** to work with your LHM server.

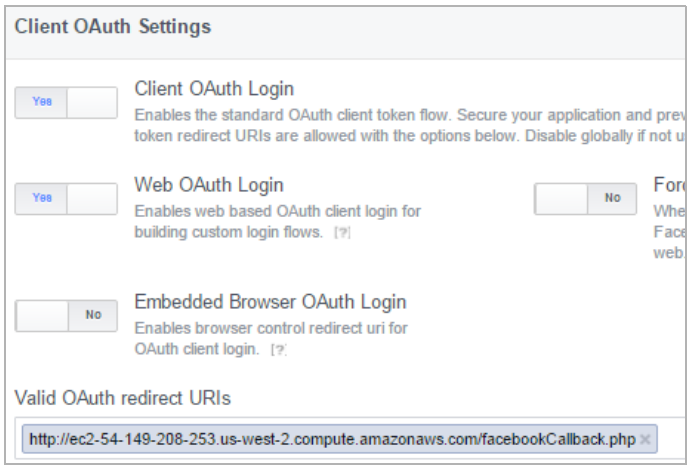
Example of settings for Facebook for developers



Client OAuth Settings

You should adjust your Client OAuth settings at Facebook for Developers, https://developers.facebook.com/ (**Products > Facebook Login > Settings**), similar to those shown in Example of OAuth Facebook settings.

Example of OAuth Facebook settings



Guest Status (demo)

When a wireless client is allowed access to the SonicWall WiFi, the owner's account name and information is sent to SonicOS. You can collect and store this information in your own databases.

Configuring Open Authentication and Social Login

Topics:

- About Configuring Guest Services on page 800
- About Configuring Social Login on page 801
- Configuring Social Login in SonicOS on page 801

About Configuring Guest Services

Although SonicOS provides its own guest account management, you can use your own IT infrastructure to better accommodate your business requirements. This configuration can be done by setting up external guest authentication or a social login. **Guest Services** is provided in the **Add/Edit Zone** dialogs of the SonicOS wireless zone, LAN zone, or DMZ zone (**MANAGE | System Setup > Network > Zones**).

About Configuring Social Login

This feature simplifies cumbersome logins for end users as well as provides reliable demographic information to web developers.

To prepare for configuring Social Login:

1 Create a wireless zone, LAN zone, or DMZ zone as described in Adding a New Zone on page 397, and set up or edit a network zone with security capabilities.

2 In SonicOS, the external server can also be created or selected as a Lightweight Hotspot Messaging (LHM) server IP or FQDN address object.

Configuring Social Login in SonicOS

Setting up your Security Appliance properly requires some configuration. The Security Appliance blocks most Internet applications, but a few should be allowed for this feature to function correctly.

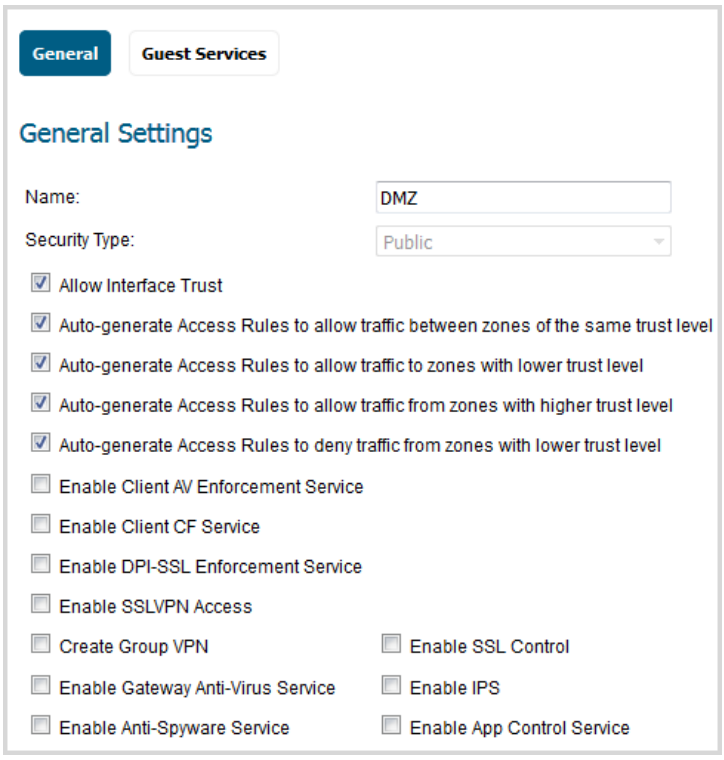
IMPORTANT: An LHM server should be in service before configuring Social Login.

To configure your Security Appliance for Social Login:

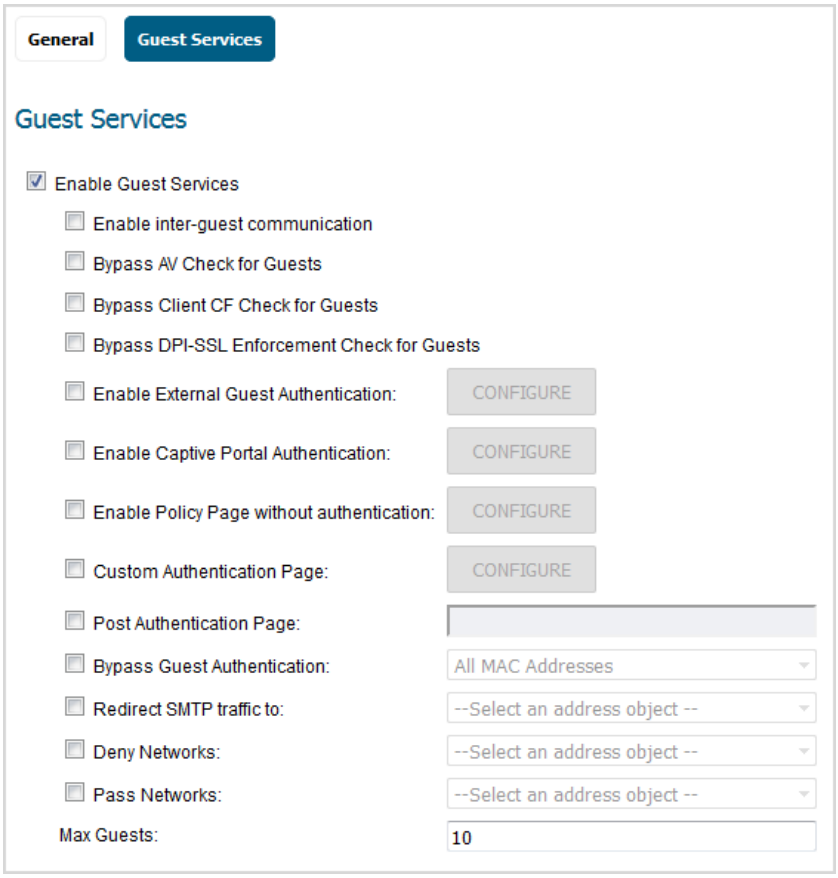
1 Navigate to **MANAGE | System Setup > Network > Zones** to set up or edit a network zone with wireless security capabilities. For more information on adding the network zone, Adding a New Zone on page 397.

NOTE: The external server can also be created or selected as an Lightweight Hotspot Messaging (LHM) server IP or FQDN address object.

2 Click the WLAN **Edit** icon to access the WLAN network zone. The **Edit Zone** dialog displays.



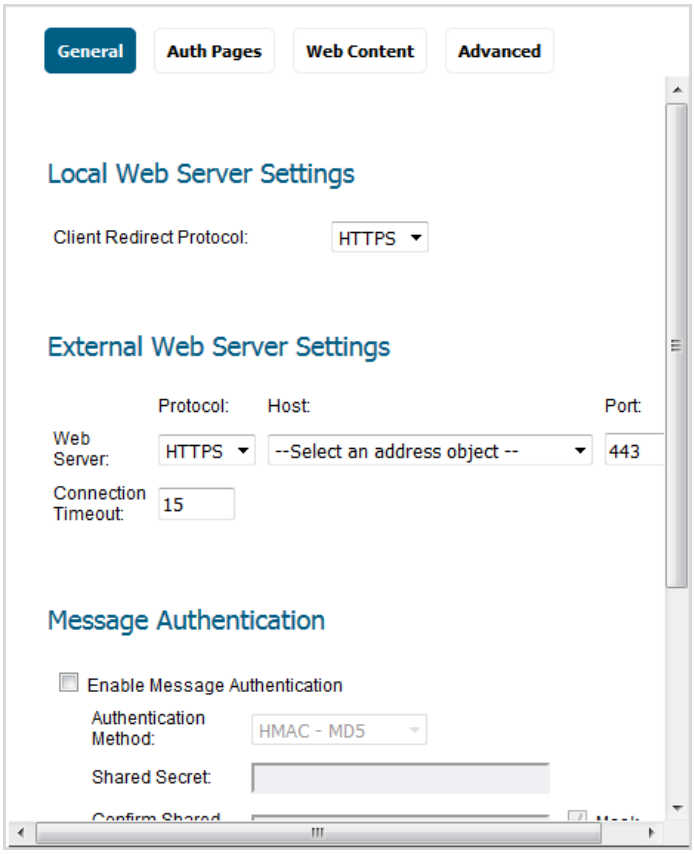
3 Click **Guest Services**.



4 Select **Enable Guest Services**. The other options activate.

5 Select **Enable External Guest Authentication**. **CONFIGURE** activates and the next four options become unavailable.

6 Click **Configure**. The **External Guest Authentication** dialog displays.



7 From **Client Redirect Protocol** under **Local Web Server Settings**, select either:

- **HTTPS** (default)
- **HTTP**

SonicOS automatically creates the necessary pass-through authentication network domains for allowing authentication process traffic between the authentication server and the user. The automatically added address object groups are named Default Social Login Pass Group. This address object group is appended to the currently configured pass networks, if any, or it is added into a new group called Social Login Pass Group.

8 For the **External Web Server Settings**, you should have an LHM server already in service:

- Select a protocol: **HTTPS** (default) or **HTTP**.
- Select an Address Object associated with the LHM server from **Host**.
- Enter the TCP port of operations for the selected protocol on the LHM server in **Port**; the default is **80**.
- Enter duration, in seconds, before the LMH server is considered unavailable on a redirect attempt in **Connection Timeout**; the default is **15** minutes. On timeout, the client is presented with the `Server Down` message configured on **Web Content**.

9 To enable message authentication, under **Message Authentication**, select **Enable Message Authentication**. The subordinate options become available. This option is not selected by default.

TIP: Use HMAC digest and embedded querystring in communication with the LHM server. This is useful if you are concerned about message tampering when HTTP is used to communicate with the LHM server. Optional.

a From **Authentication Method**, select:

- **HMAC - MD5** (default)
- **HMAC - SHA1**
- **HMAC - SHA256**

b Enter the shared secret for the hashed MAC in the **Shared Secret** field.

> **TIP:** If a shared secret is used, it also needs to be configured on the LHM server scripts.

c Repeat the shared secret in the **Confirm Shared Secret** field.

d To see the shared secret in both fields, deselect **Mask Shared Secret**. This option is selected by default.

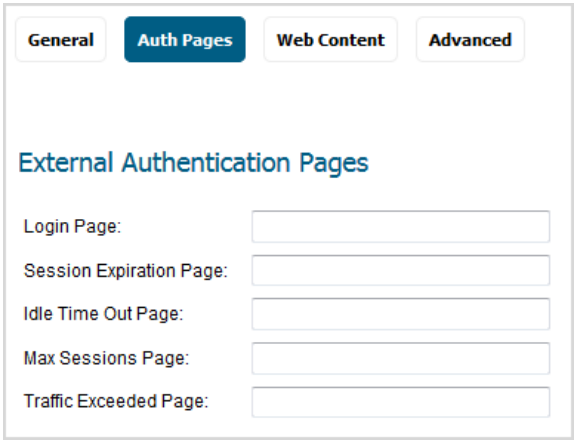
10 In the **Social Network Login** section, select **Enable Social Network Login**. The social network options activate.

11 Select one or more social networks to enable for open authentication:

- **Facebook**
- **Google**
- **Twitter**

SonicOS automatically creates the necessary pass-through authentication network domains for allowing authentication process traffic between the authentication server and the user. The automatically added address object groups are named **Default Social Login Pass Group**. This address object group is appended to the currently configured pass networks, if any, or it is added into a new group called **Social Login Pass Group**.

12 Click **Auth Pages**.



> **TIP:** These pages may each be a unique page on the LHM server, or they may all be the same page with a separate event handler for each status message. Examples are provided as follows to work with the newly developed scripts.

13 Enter a **Login Page** location, such as `login.php`, but based on your developer's input pages. These scripts are hosted by your own LHM server, so you should be able to make sure they function correctly.

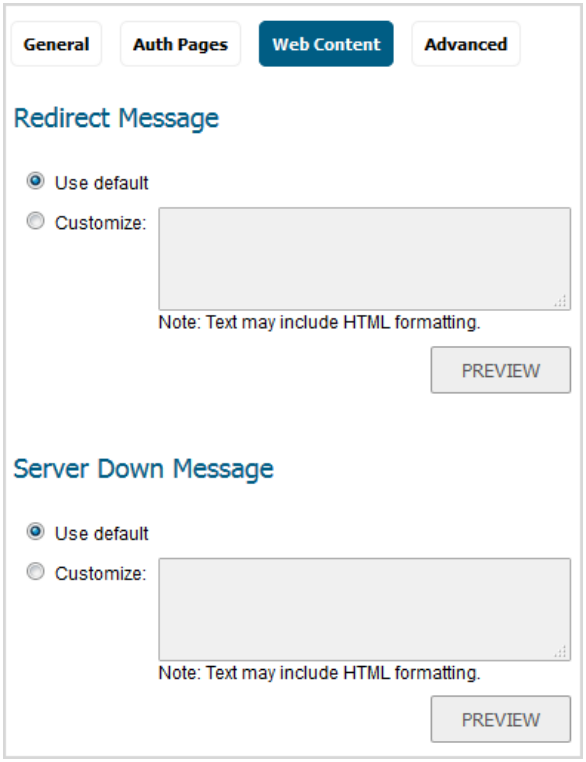
14 Enter the location of the remaining pages:

- **Session Expiration Page** – The page to which the client is redirected when the session expires. After a session expires, the user must create a new LHM session.

- **Idle time Out Page** – The page to which the client is redirected when the idle timer is exceeded. After the idle timer is exceeded, the user can log in again with the same credentials as long as there is time left for the session.
- **Max Sessions Page** – The page to which the client is redirected when the maximum number of sessions has been reached.
- **Traffic Exceeded Page** – The page to which the client is redirected when the maximum traffic has been reached.

15 If you have finished configuring the options, go to Step 27

16 Optionally, click **Web Content**.

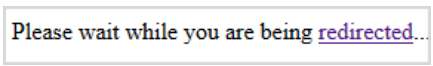


17 Under **Redirect Message**, specify either the default or customized message that is presented to the client (usually for no more than one second) explaining that the session is being redirected to the LHM server. This interstitial page is used (rather than going directly to the LHM server) so that the SonicWall Security Appliance can verify the availability of the LHM server. Choose either:

- **Use default** (default); go to Step 20.
- **Customize** – The customize field and **PREVIEW** become active.

18 Enter the custom message in the **Customize** field. The text may include HTML formatting.

19 To see a preview of your customized message or the default message, click **PREVIEW**. The **External Guest Redirect** message displays, for example, the default message:

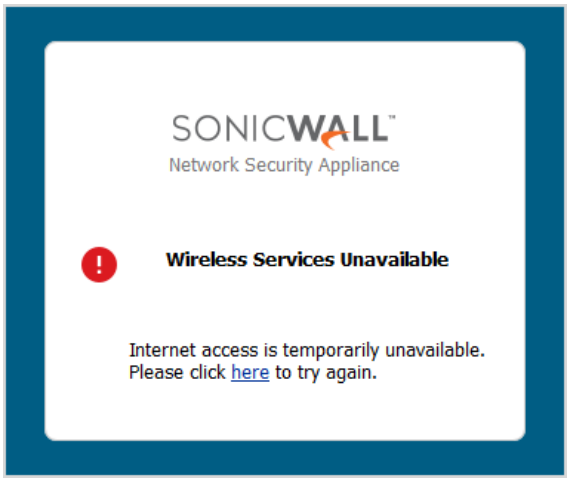


20 Under **Server Down Message**, you can specify a default or customized message that is presented to the client when the Redirector determines that the LHM server in unavailable. Choose either:

- **Use default** (default); go to Step 21.

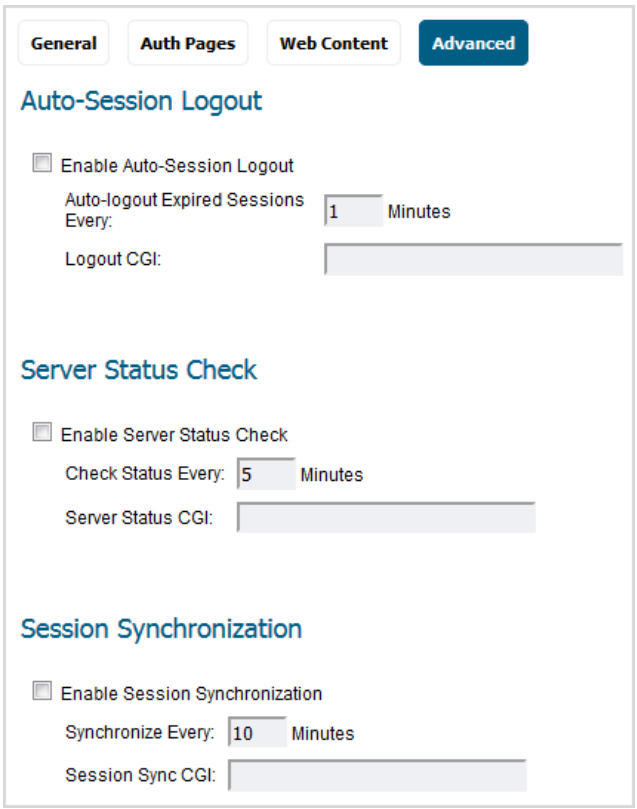
- **Customize** – Enter the custom message in the **Customize** field. The text may include HTML formatting.

21 To see a preview of your customized message or the default message, click **PREVIEW**. The **Wireless Services Unavailable** message displays, for example, the default message:



22 If you have finished configuring the options, go to Step 27.

23 Optionally, click **Advanced**.



24 To specify the time increment and the page to which the SonicWall Security Appliance POSTs when a session is logged out (either automatically or manually), in the Auto Session Logout section, select **Enable Auto-Session Logout**. The two suboptions become available. This option is not selected by default.

a To specify the time increment for logging out auto-sessions, specify the number of minutes in **Auto-logout Expired Sessions Every Minutes** field. The default is **1** minute.

b Enter the logout common gateway interface (CGI) in **Logout CGI**.

25 To specify the time increment and the page to which the Security Appliance POSTs to determine the availability of components on or behind the LHM server (such as a back-end database), in the **Server Status Check** section, select **Enable Server Status Check**. The two suboptions become available. This option is not selected by default.

a To specify the time increment for checking the server status, specify the number of minutes in **Check Status Every Minutes** field. The default is **1** minute.

b Enter the server status CGI in **Server Status CGI**.

26 To specify the time increment and the page to which the Security Appliance POSTs the entire Guest Services session table, in the **Session Synchronization** section, select **Enable Session Synchronization**. The two suboptions become available. This option is not selected by default.

a To specify the time increment for posting the Guest Services session table, specify the number of minutes in **Synchronization Minutes** field. The default is **10** minutes.

b Enter the session synchronization CGI in **Session Synch CGI**.

27 Click **OK.**

Verifying the Social Login Configuration

You can verify the correct configuration of Open Authentication and Social Login by viewing **MANAGE | Policies > Objects**. For more information about objects, see *SonicOS Policies*.

To verify settings:

1 Navigate to **MANAGE | Policies > Objects > Address Objects**.

2 Select **Address Groups**, which should show:

- Domains have been added automatically.
- Facebook, Google, and/or Twitter login traffic can pass through successfully.

Using Social Login, LHM, and ABE

Topics:

- About ABE on page 808
- Session Life Cycle on page 809
- Message Format on page 816
- LHM RESTful API on page 822
- Frequently Asked Questions (FAQs) on page 822
- LHM Script Library on page 829

About ABE

The ABE consists of a Web Server (WS) to host content for user interaction and an (optional) Authentication Server (AS) to provide directory services authentication. The AS can be any kind of user authentication mechanism, including, but not limited to RADIUS, LDAP, or AD; the only requirement is that the WS can communicate with the AS for authentication purposes. The WS and AS can be administered on a single server or on separate servers.

LHM also provides the ability for the AS to use the SonicWall Security Appliance's internal user database for user authentication. For details on messaging, see Message Format on page 816, Local Authentication Request on page 817, and Local Authentication Reply on page 817.

The ABE needs to communicate with the Hotspot SonicWall to exchange result codes and session information. All communications are HTTPS and can occur either directly (such as to the LAN, WAN, X0 interface of the SonicWall Security Appliance) or over a VPN tunnel to one of the SonicWall Security Appliance's management interface addresses. The LHM management interface is automatically derived through a route (path) lookup, and only the management interface(s) accepts LHM management messaging through automatically added Access Rules.

LHM communications occur on a specific LHM management port that must be defined on the SonicWall Security Appliance, and the LHM management port must be different from the standard HTTPS Management port.

To allow the ABE to communicate with the SonicWall, and to redirect clients to the appropriate interface on the SonicWall, two parameters are constructed by the SonicWall and passed (among others) through the client redirect to the ABE. The following communication parameters must be used for all communications between the ABE and the SonicWall.

- *mgmtBaseUrl* - The IP address and the port that the ABE uses to communicate with the SonicWall. It is composed of the HTTPS protocol designator, the IP of the selected LHM management interface, and the LHM port (such as `https://10.1.2.3:4043`).
- *clientRedirectUrl* - The IP address (and optionally the port) on the SonicWall to which clients are redirected during various phases of the session, namely the LAN management IP on the TZW, or the WLAN IP on a SonicOS device (such as `http://172.16.31.1`).

The parameter values are passed to the ABE by the SonicWall during Session Creation (see Session Creation on page 809) and during the Session State Sync (see Message Format on page 816), and should be used by the ABE as the base in the construction of all relevant URLs. The following are the pages on the SonicWall Security Appliance that is referenced by the ABE:

- *wirelessServicesUnavailable.html* – ABE is unavailable message. This redirect is typically sent by the SonicWall, but can also be referenced by the ABE. Text is configurable.
- *externalGuestRedirect.html* – Initial redirect message provided by the SonicWall on session creation. Text is configurable
- *externalGuestLogin.cgi* – The page to which the ABE posts session creation data.
- *externalGuestLogoff.cgi* – The page to which the ABE posts session termination data.
- *localGuestLogin.cgi* – The page to which the ABE posts for authenticating user credentials against the SonicWall's internal user database.
- *createGuestAccount.cgi* – The page to which the ABE posts to create a guest account in the SonicWall's internal user database.
- *externalGuestUpdateSession.cgi* – The page to which the ABE posts to update the *sessionLifetime* and *idleTimeout* parameters of an existing session (see Session Update on page 816).

For communications from the SonicWall to the ABE, URLs (including host, port, and page/resource) hosted on the ABE is fully configurable at the SonicWall Security Appliance. The host can be specified using either an IP address or fully qualified domain name (FQDN). When using FQDN, the name is resolved upon first use and is stored by the SonicWall as an IP address.

Session Life Cycle

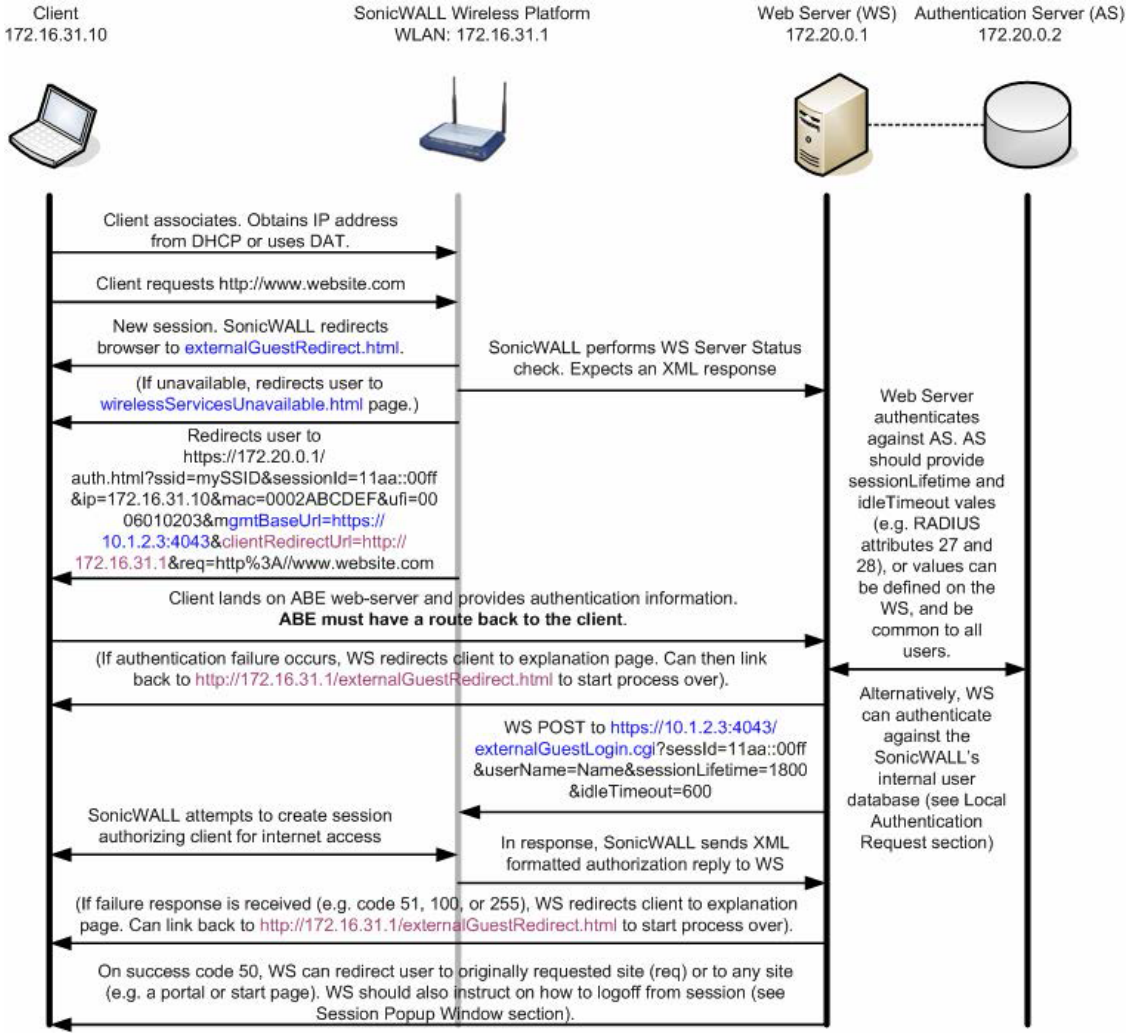
The following sections describe the phases of a session life cycle, as well as the Session Popup Window and Web Server (WS) Status Check components:

- Session Creation on page 809
- Session Popup Window on page 812
- Idle Timeout on page 812
- Session Timeout on page 813
- User Logout on page 813
- Administrator Logout (Optional) on page 814
- Web Server Status Check on page 814
- Session State Sync on page 815
- Message Authentication on page 815
- Session Update on page 816

Session Creation

Session creation occurs when a wireless client attempts access, and the SonicWall Security Appliance has no active session information for that client based upon MAC address.

Session Creation Flow



1 Wireless client associates with SonicWall. Obtains IP Address from internal DHCP server, or uses static addressing with Dynamic Address Translation (DAT) feature.

2 Client requests web-resource, `http://www.website.com`.

- SonicWall Security Appliance determines that this is a new session.

3 SonicWall Security Appliance redirects client to internally hosted `externalGuestRedirect.html` page. The `externalGuestRedirect.html` page provides administrator-configurable text explaining that the session is being redirected for authentication.

4 During this redirect, the Security Appliance checks the availability of the ABE through a JavaScript redirect attempt to the configured target redirect page.

- If the redirect to the WS fails to occur within a specified period (the value is configurable on the SonicWall, between 1 and 30 seconds) the Security Appliance redirects the session to the internal `wirelessServicesUnavailable.html` page.

5 In addition to the JavaScript availability check, an optional full Web Server Status Check is available from the SonicWall (see Web Server Status Check on page 814). This option can be configured to run at a configurable interval between 1 and 60 minutes. If an error response code of 1, 2, or 255 occurs, the Security Appliance logs the response and redirects the browser to the internal `wirelessServicesUnavailable.html` page. This page provides administrator-configurable text explaining recourse.

6 If available, the Security Appliance redirects client to authentication portal hosted on AS at: `https://172.20.0.1/auth.html?ssid=mySSID&sessionId=11aa::00ff&ip=172.16.31.10&mac=0002ABCDEF&ufi=0006010203&mgmtBaseUrl=https://10.1.2.3:4043&clientRedirectUrl=http://172.16.31.1&req=http%3A//www.website.com`

- *ssid* – The ESSID (wireless network name) of the wireless network to which the redirected client was associated.
- *sessionId* — A 32 byte hex representation of a 16 byte MD5 hash value generated by the SonicWall, which is used by the SonicWall and the WS for indexing clients (such as `11aa3e2f5da3e12ef978ba120d2300ff`).
- *ip* — The client IP address.
- *mac* — The client MAC address.
- *req* — The originally requested web-site is passed as an argument to the authentication server)
- *ufi* — The SonicWall Unique Firewall Identifier. To be used for site identification, if desired.
- *mgmtBaseUrl* — The protocol, IP address, and port on the SonicWall with which the IP subsequently communicates.
- *clientRedirectUrl* — The protocol, IP address (and optionally port) on the SonicWall that the ABE uses for client redirection.
- *req* – The client's originally requested URL, if any, URL encoded.

7 Client provides authentication information (such as username, password, token, and so on).

NOTE: The WS must be able to reach the Client, for example, by VPN, NAT or route.

8 WS validates user against AS.

- AS provides session specific information, namely, Session Timeout and Idle Timeout values.
- Session specific values can optionally be applied globally by the WS rather than obtained from the AS; some value simply needs to be passed to the Security Appliance.
- Timeout values are presented in seconds and can range from 1 to 863,913,600 (equal to 9999 days).

9 If authentication fails, the WS should redirect the client to a page explaining the failure. A link should be provided back to `http(s)://172.16.31.1/externalGuestRedirect.html` to restart the process.

10 If successful, the WS connects to the Security Appliance either through HTTPS or through VPN and POSTs `https://10.1.2.3:4043/externalGuestLogin.cgi?sessId=11aa::00ff&userName=Name&sessionLifetime=1800&idleTimeout=600`.

- The Security Appliance attempts to create the session and sends a result to the WS in the same connection. Results are described in Message Format on page 816.

11 If a failure response is received (such as code 51, 100, or 255), WS should redirect client to a page explaining the failure. A link can be provided back to: `http(s)://172.16.31.1/externalGuestRedirect.html` to start process over.

12 If successful (code 50), WS can redirect user to the originally requested site (`req`) or to any site (such as a portal or start page). WS should also instruct on how to logoff from session (such as bookmark a page, popup window, URL).

Session Popup Window

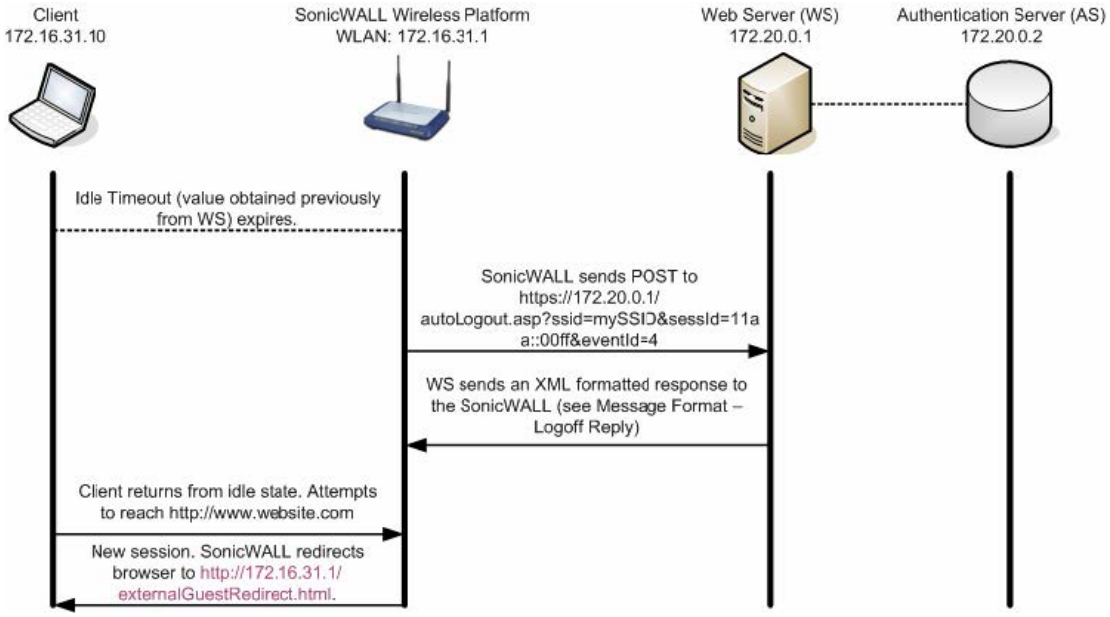
It is recommended that sessions be managed through a Session Popup window. This should be a browser window instantiated at the time of Session Creation providing session time information (such as lifetime, idle timeout value, timer countdowns) and **Logout**. Sample code is provided.

- Clicking **Logout** ends the session and triggers a User Logout event.
- Attempting to close the window should provide a warning message that closing the window ends the session.
- Closing the window ends the session and triggers a User Logout event.

Idle Timeout

Idle timeout occurs when the idle timeout (specified in Session Creation on page 809, Step 8) is exceeded.

Idle Timeout Flow



1 Idle timer (as set during Session Creation on page 809) expires.

2 Because the client's browser might not be open at this time, we do not initiate this process with a redirect. Instead, SonicWall sends a POST to the WS at: `https://172.20.0.1/autoLogout.asp?ssid=mySSID&sessId=11aa::00ff&eventId=4` (see Message Format on page 816 for Logoff event IDs).

   - The resource to which the POST is sent is configurable on the Security Appliance from **MANAGE | System Setup | Network > Zones**. Edit the WLAN zone (on the **Edit Zone** dialog: **Guest Services > External Guest Authentication > Advanced > Auto-Session Logout > Logout CGI**.
   - The WS hosted page must expect and interpret the *sessId* and *eventId* values.

3 The WS sends an XML result to the WS in the same connection. Results are described in Logoff Reply on page 818.

4 If the client returns from the idle state and attempts to reach a web resource, the Security Appliance redirects the user to the internal `externalGuestRedirect.html` page, starting the session creation process over (see Session Creation on page 809).

NOTE: To conserve resources, it is recommended that the idle timeout be set to a maximum of 10 minutes.

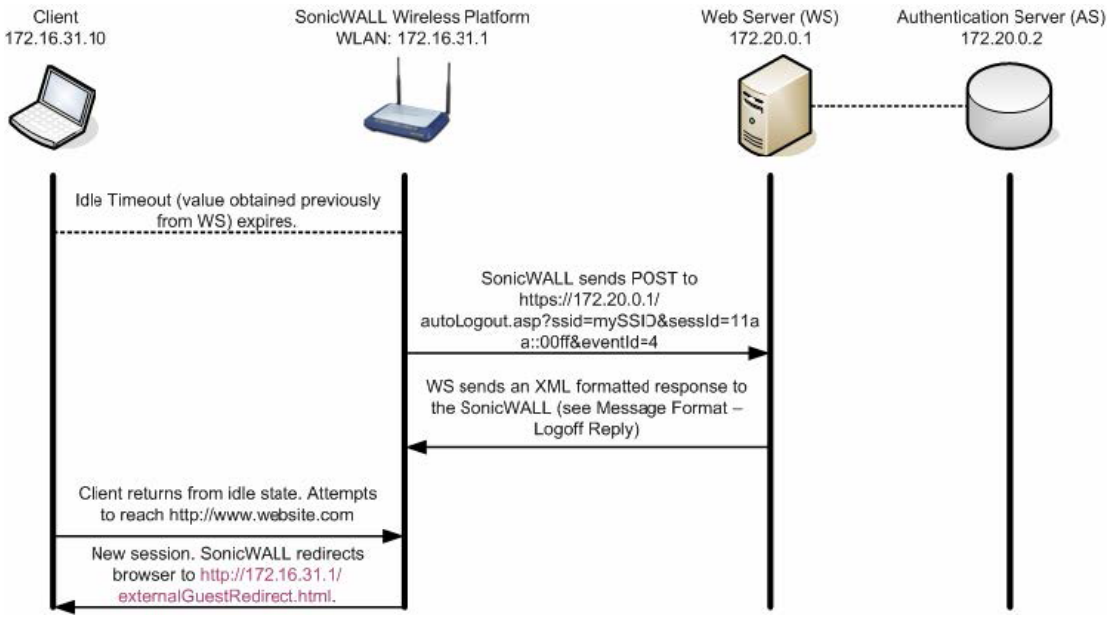
Session Timeout

The event occurs when the Session lifetime expires. The exchange is the same as the Idle Timeout above, except the Session Timeout *`eventId`* value is 3 instead of 4 for an Idle Timeout.

User Logout

Event occurs when the user actively ends the session by closing their Session Popup window or by using **Logout** provided on the Session Popup window. The Session Popup window is the preferred method for user logout; however, the same result can be achieved without this method by allowing the session's lifetime to expire. The latter removes the dependency on the Session Popup window, but manages resources less efficiently.

User Logout Flow



1 Client logs out using **Logout**, or closes the session popup window.

2 The WS sends a POST to: `https://10.1.2.3:4043/externalGuestLogoff.cgi?sessId=11aa::00ff&eventId=1` (for Logoff event IDs, see Message Format on page 816).

- `sessId` — The value generated during Session Creation (see Session Creation on page 809) by the Security Appliance, which is used by the Security Appliance and the WS for indexing clients.
- `eventId` — Describes the logoff request event.

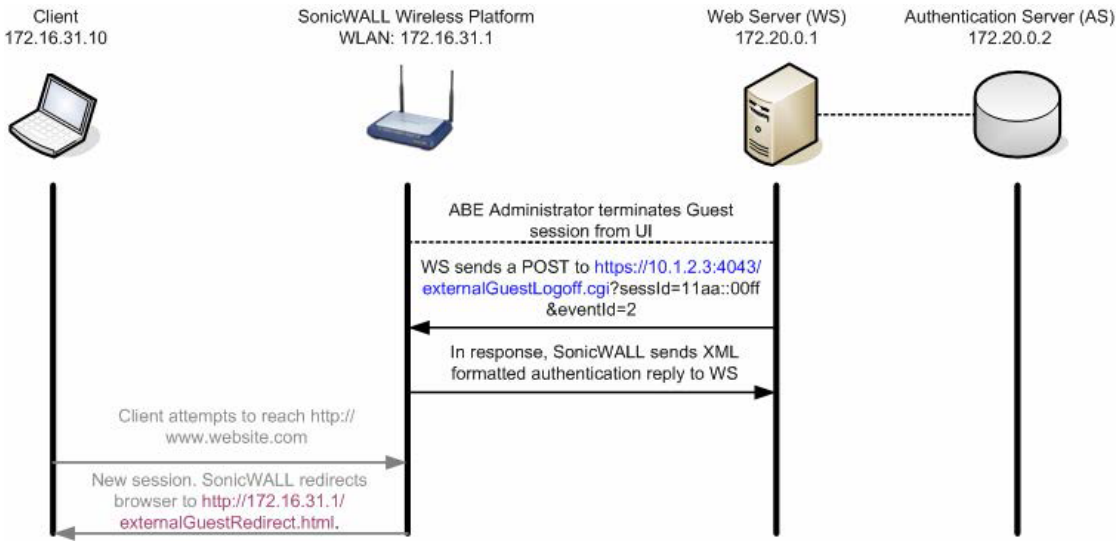
3 SonicWall Security Appliance responds with a result to the WS in the same connection. Results are described in Logoff Reply on page 818.

4 If the client attempts to reach a web resource, the Security Appliance redirects the user to the internal `http://172.16.31.1/externalGuestRedirect.html` page, starting the Session Creation process over (see Session Creation on page 809).

Administrator Logout (Optional)

The event occurs when the ABE administrator logs out from a Guest session from the management interface. It is not possible at this time to terminate ABE-established Guest Sessions from the SonicOS Management Interface itself. ABE-established Guest Sessions are represented as such (or distinctly from internal WGS Guest Sessions) on the SonicOS Management Interface and are not editable.

Administrator Logout Flow



1 The ABE administrator terminates the Guest session from the management UI.

2 The WS sends a POST to the Security Appliance: `https://10.1.2.3:4043/externalGuestLogoff.cgi?sessId=11aa::00ff&eventId=2`. (for Logoff event IDs, see Message Format on page 816).

- `sessId` — The value generated during Session Creation by the Security Appliance, which is used by the Security Appliance and the WS for indexing clients.
- `eventId` — Describes the logoff request event.

3 The SonicWall sends a result to the WS in the same connection. Results are described in Logoff Reply on page 818.

4 If the client returns from the idle state and attempts to reach a web resource, the Security Appliance redirects the user to the internal `http://172.16.31.1/externalGuestRedirect.html` page, starting the Session Creation process over (see Session Creation on page 809).

Web Server Status Check

To provide more granular ABE status than simple Web Server (WS) availability (as is provided by the mandatory Step 4 of Session Creation Flow on page 810, the JavaScript redirect), the SonicWall can optionally send a secure HTTP GET operation to the WS in order to determine server operational status. The target URL is configurable, as is the interval of the query (between 1 and 60 minutes). The WS responds back in an XML format listing the server's current state. For details. see Message Format on page 816.

If an error response code (1, 2, or 255) is received (indicating that the WS itself is available, but that some other ABE error condition has occurred), the SonicWall logs the response and redirects all subsequent authentication requests to an internal `wirelessServicesUnavailable.html` page. This page provides administrator-configurable text explaining recourse.

The Security Appliance continues to attempt to query the ABE at the configured interval and resumes redirection to the WS (rather than to the `wirelessServicesUnavailable.html` page) when a response code of 0 (Server Up) is received.

Session State Sync

At a configurable interval (between 1 and 60 minutes), the Security Appliance optionally sends a secure HTTP POST operation to the WS containing an XML list of all currently active guest sessions. The CGI post provides the `sessionList` as an XML list of all active guest sessions. For details, see Message Format on page 816.

The feature itself is enabled through a checkbox on the Security Appliance, but is disabled by default. The target URL is configurable.

Message Authentication

This feature ensures that the CGI data exchanged between both the Security Appliance and ABE originated from the SonicWall Security Appliance/ABE device, and that it has not been tampered with. If enabled, an additional CGI parameter, named `hmac`, is added to all CGI data exchanged. The following is an example of what the redirect URL now looks like with message authentication enabled:

```
https://10.1.2.3/login.asp?sessionId=faad7f12ac26d5c2fe3236de2c149a22&ip=172.16.31.2&mac=00:90
:4b:6a:37:32&ufi=0006B1020148&mgmtBaseUrl=https://10.0.61.222:4043/&clientRedirectUrl=http://1
92.168.168.168:80/&req=http%3A//www.google.com/&hmac=cd2399aeff26d5c2fe3236d211549acc
```

NOTE: The SonicWall Security Appliance URL encodes the following characters within the value of the `req` (and only the `req`) variable:

```
% = %25
: = %3A
= %20 (space)
? = %3F
+ = %2B
& = %26
= = %3D
```

In the preceding example, the HMAC signature was generated using the following data:

```
HMAC(
   faad7f12ac26d5c2fe3236de2c149a22 +
   172.16.31.2 +
   00:90:4b:6a:37:32 +
   0006B1020148 +
   https://10.0.61.222:4043/ +
   https://10.0.61.222:4043/ +
   http%3A//www.google.com/
)
```

If message authentication is enabled, then the SonicWall device expects an HMAC signature as part of the CGI post data originating from the ABE. If the SonicWall detects that the HMAC is missing or incorrect, then an error code of 251 is returned, and the requested operation (such as guest login, account creation) is aborted.

Session Update

Session update allows for the ABE to update the Session Lifetime and Idle Timeout values of existing session on the Security Appliance. This allows, for example, for additional time to be purchased by guest users and added to an existing session.

- The Session Update can be sent from the ABE to the SonicWall at any time during a session's lifetime.
- The *userName* and *sessionLifetime* values must be specified in the message
- The *sessID* value may be specified. If included, the update pertains to the specified session. If omitted, the update pertains to all sessions matching the specified *userName*.

For details, see Message Format on page 816.

Message Format

Topics:

- External Authentication Request on page 816
- Local Authentication Request on page 817
- Local Authentication Request on page 817
- Local Authentication Reply on page 817
- Logoff Request on page 818
- Logoff Reply on page 818
- Web Server Status Check on page 814
- Session State Sync on page 815
- Session State Sync Reply on page 820
- Local Account Creation Request on page 820
- Local Account Creation Reply on page 821
- Update Session Request on page 821
- Update Session Reply on page 821

NOTE: The XML Schema location is subject to change.
The SonicWall Security Appliance IP address and port is defined in the *mgmtBaseUrl* variable.

External Authentication Request

The WS sends a secure HTTP POST operation to: `https://SonicWall.ip.add.ress:port/externalGuestLogin.cgi`. The post parameters include these arguments:

- *sessId*: Session ID
- *userName*: The full user ID
- *sessionLifetime*: The session lifetime of the user (in seconds)
- *idleTimeout*: The maximum idle timeout (in seconds)

External Authentication Reply

The Security Appliance returns an XML response in this format:

```
<?xml version="1.0" encoding="UTF-8">
<SonicWallAccessGatewayParam
   xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
   xsi:noNamespaceSchemaLocation="http://www.SonicWall.com/
   SonicWallAccessGatewayParam.xsd">
   <AuthenticationReply>
      <ResponseCode>{response code}</ResponseCode>
      <ReplyMessage>{reply message}</ReplyMessage>
   </AuthenticationReply>
</SonicWallAccessGatewayParam>
```

The `{response code}` includes one of the values listed in External Authentication Response Codes.

External Authentication Response Codes

Response Code	Response Meaning
50	Login succeeded
51	Session limit exceeded
100	Login failed -- access reject
251	Msg. Auth failed -- Invalid HMAC
253	Invalid session ID
254	Invalid or missing CGI parameter
255	Internal error

Local Authentication Request

The WS sends a secure HTTP POST operation to: `https://SonicWall.ip.add.ress:port/localGuestLogin.cgi`. The post parameters includes these arguments:

- *`sessId`*: Session ID
- *`userName`*: The full user ID
- *`passwd`*: The guest's clear-text password

Local Authentication Reply

The SonicWall returns an XML response in this format:

```
<?xml version="1.0" encoding="UTF-8">
<SonicWallAccessGatewayParam
   xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
   xsi:noNamespaceSchemaLocation="http://www.SonicWall.com/
   SonicWallAccessGatewayParam.xsd">
   <AuthenticationReply>
      <ResponseCode>{response code}</ResponseCode>
      <ReplyMessage>{reply message}</ReplyMessage>
   </AuthenticationReply>
</SonicWallAccessGatewayParam>
```

The `{response code}` includes one of the values listed in Local Authentication Response Codes.

Local Authentication Response Codes

Response Code	Response Meaning
50	Login succeeded
51	Session limit exceeded
52	Invalid username/password
100	Login failed -- access reject
251	Msg. Auth failed -- Invalid HMAC
253	Invalid session ID
254	Invalid or missing CGI parameter
255	Internal error

Logoff Request

The WS sends a secure HTTP POST operation to: `https://SonicWall.ip.add.ress:port/externalGuestLogoff.cgi`. The post parameters includes the following arguments:

- *`sessId`*: GW Session ID
- *`eventId`*: Logoff event ID. Must be one of the following:

Logoff Event ID	Event Meaning
1	Guest logged out manually
2	Admin logged off the specified guest
3	Guest session expired
4	Guest idle timeout expired

Logoff Reply

The Security Appliance returns an XML response in this format:

```
<?xml version="1.0" encoding="UTF-8">
<SonicWallAccessGatewayParam
   xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
   xsi:noNamespaceSchemaLocation="http://www.SonicWall.com/
   SonicWallAccessGatewayParam.xsd">
   <LogoffReply>
        <ResponseCode>{response code}</ResponseCode>
        <ReplyMessage>{reply message}</ReplyMessage>
   </LogoffReply>
</SonicWallAccessGatewayParam>
```

The *`{response code}`* includes one of the values listed in Logoff response codes:

Logoff response codes

Response Code	Response Meaning
150	Logoff succeeded
251	Msg. Auth failed -- Invalid HMAC
253	Invalid session ID

Logoff response codes

Response Code	Response Meaning
254	Invalid or missing CGI parameter
255	Internal error

Web Server Status Check

The WS returns an XML response in this format:

```
<?xml version="1.0" encoding="UTF-8">
<SonicWallAccessGatewayParam
   xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
   xsi:noNamespaceSchemaLocation="http://www.SonicWall.com/
   SonicWallAccessGatewayParam.xsd">
   <ServerStatus >{status code}</ ServerStatus >
</SonicWallAccessGatewayParam>
```

The *{response code}* includes one of the values listed in Web server status check response codes.

Web server status check response codes

Response Code	Response Meaning
0	Server Up
1	DB down
2	Configuration error
255	Internal error

Session State Sync

Periodically, the GW sends a secure HTTP POST operation to the AS containing an XML list of all currently active guest sessions. Both the target URL and time period are configurable by the GW administrator.

The CGI post parameters include this argument:

- *sessionList*: XML list of all active GW guest sessions.

The session list returns an XML response in this format:

```
<?xml version="1.0" encoding="UTF-8">
<SonicWallAccessGatewayParam
   xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
   xsi:noNamespaceSchemaLocation="http://www.SonicWall.com/
   SonicWallAccessGatewayParam.xsd">
   <SessionSync>
      <SessionCount>{Session Count}</SessionCount>
      <SessionList>
      <Session>
      <Ssid>{ESSID}</Ssid>8
      <ID>{Session ID}</ID>
      <UserName>{User Name}</UserName>
      <IP>{IP Address}</IP>
      <MAC>{MAC Address}</MAC>
      <Idle>
      {Time Idle (expressed in seconds)}
      </Idle>
      <SessionRemaining>
      {Session Remaining (expressed in seconds)}
      <SessionRemaining>
      <BaseMgmtUrl>
```

```
        {https://ip.add.re.ss:port}
        </BaseMgmtUrl>
        <RxBytes>
        {total bytes received}
        </RxBytes>
        <TxBytes>
        {total bytes transmitted}
        </TxBytes>
        </Session>
    </SessionList>
</SessionSync>
</SonicWallAccessGatewayParam>
```

Session State Sync Reply

The WS returns an XML response in this format:

```
<?xml version="1.0" encoding="UTF-8">
<SonicWallAccessGatewayParam
   xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
   xsi:noNamespaceSchemaLocation="http://www.SonicWall.com/
   SonicWallAccessGatewayParam.xsd">
   <SessionSync>
      <ResponseCode>{response code}</ResponseCode>
   </SessionSync>
</SonicWallAccessGatewayParam>
```

The *{response code}* includes one of the values listed in Session state sync reply response codes.

Session state sync reply response codes

Response Code	Response Meaning
200	Sync successful
201	Sync failed
255	Internal error

Local Account Creation Request

The WS sends a secure HTTP POST operation to: `https://SonicWall.ip.add.ress:port/createGuestAccount.cgi`. The post parameters include these arguments:

- *userName*: The full user ID (maximum length: 32)
- *passwd*: The guest's clear-text password (maximum length: 64)
- *comment*: Optional (maximum length: 16). Default=**NULL**
- *enforceUniqueLogin*: Optional: 1=true, 0=false. Default=**1**
- *activateNow*: Optional: 1=true, 0=false. Default=**0**
- *autoPrune*: Optional: 1=true, 0=false. Default=**1**
- *accountLifetime*: The account lifetime of the user (expressed in seconds)
- *sessionLifetime*: The session lifetime of the user (expressed in seconds)
- *idleTimeout*: The max idle timeout (expressed in seconds)

Local Account Creation Reply

The Security Appliance returns an XML response in this format:

```
<?xml version="1.0" encoding="UTF-8">
<SonicWallAccessGatewayParam
   xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
   xsi:noNamespaceSchemaLocation="http://www.SonicWall.com/
   SonicWallAccessGatewayParam.xsd">
   <AccountCreationReply>
      <ResponseCode>{response code}</ResponseCode>
      <ReplyMessage>{reply message}</ReplyMessage>
   </AccountCreationReply>
</SonicWallAccessGatewayParam>
```

The *{response code}* includes one of the values listed in Local Account Creation Reply Response Codes.

Local Account Creation Reply Response Codes

Response Code	Response Meaning
10	Account creation succeeded
11	Max account limit
12	Account Exists
251	Msg. Auth failed -- Invalid HMAC
254	Invalid or missing CGI parameter
255	Internal error

Update Session Request

The POST from the ABE may be made to the Security Appliance at `externalGuestUpdateSession.cgi` in this format:

```
https://10.1.2.3:4043/externalGuestUpdateSession.cgi?sessId=11aa::00ff&userName=guest&sessionLifetime=600&idleTimeout=180
```

The post parameters include these arguments:

- *sessID*: The value may be specified. If the value is not specified, then all guest sessions matching the specified username are updated.
- *userName*: The value must be specified as it defines the name of the user session (or potentially sessions if no session ID is provided) that is updated.
- *sessionLifetime*: The value must be specified as it defines the number of seconds to assign to the session. It can be any number from 1 to 863,913,600.
- *idleTimeout*: The value may be specified. It:
  - Defines the number of seconds to assign to the session.
  - Can be any number from 1 to 863,913,600.
  - Must be less than or equal to the *sessionLifetime*.

  If an *idleTimeout* is not provided, the session's existing *idleTimeout* value is maintained.

Update Session Reply

The Security Appliance returns an XML response in this format:

```
<?xml version="1.0" encoding="UTF-8">
<SonicWallAccessGatewayParam
   xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
   xsi:noNamespaceSchemaLocation="http://www.SonicWall.com/
   SonicWallAccessGatewayParam.xsd">
   <UpdateSessionReply>
      <ResponseCode>{response code}</ResponseCode>
      <ReplyMessage>{reply message}</ReplyMessage>
   </ UpdateSessionReply >
</SonicWallAccessGatewayParam>
```

The *{response code}* includes one of the values listed in Update session reply response codes.

Update session reply response codes

Response Code	Response Meaning
210	Session Update succeeded
211	Session Update failed
251	Msg. Auth failed -- Invalid HMAC
254	Invalid or missing CGI parameter
255	Internal error

LHM RESTful API

SonicOS supports LHM RESTful API.

Lightweight Hotspot Messaging (LHM) defines the method and syntax for communications between a SonicWall wireless access device (such as a SOHO W, TZ-series W, or a SonicPoint with a governing SonicWall Security Appliance) and an Authentication Back-End (ABE) for authenticating Hotspot users and providing them parametrically bound network access.

A RESTful API is an application program interface (API) that uses HTTP requests to GET, PUT, POST and DELETE data. A RESTful API, also referred to as a RESTful web service, is based on representational state transfer (REST) technology, an architectural style and approach to communications often used in web services development.

Frequently Asked Questions (FAQs)

Topics:

- Do the LHM server scripts have to be written in ASP? on page 823
- Why were these new scripts written in ASP.NET? on page 823
- How can I use LHM to provide Guest Services access to wired users? on page 823
- Can I use LHM to provide access using LDAP, RADIUS, a button, the time of day, tasseography, a survey, relative barometric pressure, a pass code, and so on as the authenticator? on page 824
- Can SonicWall write the script for me that does that? on page 824
- I want to use the sample scripts SonicWall provided. What do I need to do to use them? on page 824
- Where can the LHM server reside? on page 825
- Why are my Guest Clients unable to reach the LHM Server, or why are the pages on the LHM server not loading? on page 825
- How does the LHM exchange between the SonicWall and the LHM server work (concise version, typical environment)? on page 825

- What do all the LHM settings mean? How do I configure them? on page 826
- Can I change the LHM Management port from its default of TCP 4043? on page 828
- Do I need to use the HMAC option? If I do want to use it, how do I use it? on page 828
- Does SonicWall provide any support for these scripts? on page 828
- I've written a new script, I've made some great enhancements to your scripts, or I've just made your scripts work a whole lot better than you did; is SonicWall interested? on page 829
- LHM Script Library on page 829

Do the LHM server scripts have to be written in ASP?

No. The LHM server scripts can be written using any platform capable of handling web requests and XML, the two core components of LHM. This includes Perl, PHP, ASP, ASP.NET, and J2EE.

Why were these new scripts written in ASP.NET?

ASP.NET was chosen for the new scripts because of its prevalence, and because it does lots of things well, not the least of which being the ease with which it handles XML.

How can I use LHM to provide Guest Services access to wired users?

Although Guest Services (previously known as WGS, or Wireless Guest Services) were designed for wireless (hotspot) users, Guest Services can also be employed for wired users by placing the wired interface (or interfaces, as the case may be on the PRO 1260 with PortShield) into a Wireless Zone with SonicPoint Enforcement disabled. All Guest Services options then apply to wired users, including among others, LHM, Dynamic Address Translations, Allow/Deny Networks.

What is the difference between authentication and authorization?

Authentication describes the process of a user providing a response to some kind of challenge. The challenge can be just about anything, although traditionally it is a `username:password`. LHM breaks this dependence of the traditional model by abstracting the authentication. The role of authenticator is fulfilled by the LHM server, and the methods of authentication are bound only by imagination. Consider the following methods of authentication:

- Provide a valid username and password
- Guess the number the computer generated
- Complete this questionnaire
- Pass a quiz with a score of at least 80 percent
- Click **I Accept**

After authentication, the client can then be authorized to do something.

Authorization is the process of granting access to something. For authorization to be useful, the authorizer must have a means of stopping the client from getting to guarded resources. In the case of LHM, the SonicWall is the client's gateway (either wired or wireless), so it can very effectively act as authorizer. After the SonicWall

receives the OK from the authenticator for a client, it creates the Guest Services session and allows the client access to the Internet.

Can I use LHM to provide access using LDAP, RADIUS, a button, the time of day, tasseography, a survey, relative barometric pressure, a pass code, and so on as the authenticator?

Yes.

Can SonicWall write the script for me that does that?

We have provided a series of sample scripts as examples and for you to freely modify, but we do not provide custom scripts. We can, however, put you in touch with someone who can provide custom scripts. There are many SonicWall partners who have web development teams on staff who can provide these services.

I want to use the sample scripts SonicWall provided. What do I need to do to use them?

You need:

- Microsoft Windows 2000, XP, 2003 platform running IIS 5.0 or higher, running the latest service packs and Hotfixes.
- The Microsoft .NET 1.1 (or higher) Framework: http://www.microsoft.com/downloads/details.aspx?FamilyId=262D25E3-F589-4842-8157-034D1E7CF3A3&displaylang=en
- The latest .NET Framework Service Pack: http://www.microsoft.com/downloads/details.aspx?familyid=A8F5654F-088E-40B2-BBDB-A83353618B38&displaylang=en

To use the scripts:

1. Copy the LHM script (or scripts) you wish to use to the `wwwroot` directory (usually in `C:\inetpub\wwwroot`).
2. Configure Guest Services on your SonicWall to use External Guest Authentication, as described in the What do all the LHM settings mean? How do I configure them? on page 826.

   Some scripts need write privileges, particularly those that use databases. Depending on your configuration, two or three separate “users” need to have write access to the script directories that require writing.

   - The first account (all platforms) is **IUSR_MACHINENAME** (where *machinename* = the name of the local machine).
   - The second account on:
     - Windows XP is **ASPNET** (ASP.NET machine account).
     - Other platforms is **IWAM_MACHINENAME** (where *machinename* = the name of the local machine).

- If database read/write access continues to fail even after assigning these permissions, it might be necessary to add read/write privileges for the **NETWORK SERVICE** account.

NOTE: Versions on .NET Framework prior to 1.1 had user permission problems on domain controllers (http://support.microsoft.com/default.aspx?scid=kb;en-us;Q315158). It is strongly recommended that 1.1 (or higher) be installed.

3 After your environment is set up, you need to customize the scripts. This has been maed as simple as possible by placing all the interesting configurable bits in the `myvars.aspx` file. All entries are well commented, and their purposes and syntax should be evident. Further customization to the scripts themselves can be performed, but is generally not necessary.

Where can the LHM server reside?

The LHM Server can be virtually anywhere in the network, as long as it is reachable by the Guest Clients. It can be located at a centralized network operations center where it can administer LHM for multiple hotspots, or it can be co-located with a single SonicWall Security Appliance.

Why are my Guest Clients unable to reach the LHM Server, or why are the pages on the LHM server not loading?

Guest clients communicate directly with the LHM server; the communication is not proxied by the SonicWall Security Appliance. In other words:

- The Guest Client's subnet must be able to reach the LHM server.
- The LHM server must know how to reach the Guest Client's subnet (by route, NAT, or VPN).
- Firewall Access Rules must be configured to allow the Guest Client subnet to reach the LHM server.

How does the LHM exchange between the SonicWall and the LHM server work (concise version, typical environment)?

1 The Guest Client associates, gets a DHCP lease, and launches a web browser.

2 DNS is allowed through the SonicWall Security Appliance. The URL FQDN resolves to its IP address.

3 The SonicWall Security Appliance checks if the Guest Client has an authenticated session.

- If it's new, SonicWall Security Appliance redirects the client to the internal redirect (`Please wait while you are being redirected`) page.

4 The internal redirect page attempts to redirect the Guest Client to the LHM server.

- If it fails, it redirects the client to the internal server-down (`Wireless Internet access is temporarily unavailable. Please click here to try again`) page.

5 The Guest Client is redirected to the LHM server. In the redirect URL, the Security Appliance embeds `querystring` information describing the embryonic session (such as the `sessionID`, the client's MAC and IP address, the Security Appliance's LHM management IP and port, the UFI, the originally requested URL).

- The LHM server script grabs the `querystring` information.
- The client directly retrieves the LHM landing page from the LHM server.

6 Depending on the authorization model used (such as `username:password`, `passcode`, **I Accept** button), the LHM server decides that the Guest Client is worthy of access.

7 The LHM server initiates a web-request to the SonicWall Security Appliance at the configured management port (such as TCP 4043) to the `externalGuestLogin.cgi` page.

- The LHM server POSTs the `sessionID` (which it obtained in Step 5) along with the `username` (which it either got from the user or made up) and the `session-lifetime` and `idle-timeout` (both of which it determines).

8 The Security Appliance validates the `sessionID`, tries to create the session, and then responds to the POST with a result code describing whether or not it was able to authorize (create) the Guest session.

9 The LHM Server interprets the result code and reports the results (such as `Session Authorized - You may now start browsing`, `Session creation failed - Rats`, `Max sessions`) to the Guest Client.

What do all the LHM settings mean? How do I configure them?

Rather than going into the full detail provided in About Lightweight Hotspot Messaging (LHM) on page 796, let's just explain what the settings mean and how you might configure them:

The LHM configuration on a wireless SonicOS is done on the **Edit Zone — WLAN** dialog.

Topics:

- General on page 826
- Auth Pages on page 827
- Web Content on page 827
- Advanced on page 827

General

Local Web Server Settings

Client Redirect Protocol	The protocol (HTTP or HTTPS) used by the SonicWall Security Appliance when performing the initial internal client redirect through the `Please wait while you are being redirected` page. (This message is configurable from the **Redirect Message** area on the **Web Content** tab.) This step is prior to redirection to the LHM server.

External Web Server Settings

Web Server Protocol	The protocol (HTTP or HTTPS) running on the LHM server.
Web Server Host	The IP or resolvable FQDN of the LHM server.
Web Server Port	The TCP port of operations for the selected protocol on the LHM server.
Connection Timeout	The duration of time, in seconds, before the LMH server is considered unavailable on a redirect attempt. On timeout, the client is presented with the `Server Down` message configured on the **Web Content** tab.

Message Authentication

Enable Message Authentication	Use HMAC digest and embedded querystring in communication with the LHM server. This is useful if you are concerned about message tampering when HTTP is used to communicate with the LHM server. Optional.
Authentication Method	Select **MD5** or **SHA1**.
Shared Secret	The shared secret for the hashed MAC. If used, it also needs to be configured on the LHM server scripts.

Auth Pages

External Authentication Pages

NOTE: These pages may each be a unique page on the LHM server, or they may all be the same page with a separate event handler for each status message. Examples are provided as follows to work with the newly developed scripts.

Login Page	The first page to which the client is redirected (such as `lhm/accept/default.aspx`).
Session Expiration Page	The page to which the client is redirected when the session expires (such as `lhm/accept/default.aspx?cc=2`). After a session expires, the user must create a new LHM session.
Idle Timeout Page -	The page to which the client is redirected when the idle timer is exceeded (such as `lhm/accept/default.aspx?cc=3`). After the idle timer is exceeded, the user can log in again with the same credentials as long as there is time left for the session.
Max Session Page	The page to which the client is redirected when the maximum number of sessions has been reached (such as `lhm/accept/default.aspx?cc=4`).

Web Content

Redirect Message

The default or customized message that is presented to the client (usually for no more than one second) explaining that the session is being redirected to the LHM server. This interstitial page is used (rather than going directly to the LHM server) so that the Security Appliance can verify the availability of the LHM server.

Server Down Message

The default or customized message that is presented to the client when the Redirector determines that the LHM server in unavailable.

Advanced

The parameters are optional.

Auto Session Logout	The time increment and the page to which the SonicWall Security Appliance POSTs when a session is logged out (either automatically or manually).

Server Status Check	The time increment and the page to which the SonicWall POSTs to determine the availability of components on or behind the LHM server (such as a back-end database).
Session Synchronization	The time increment and the page to which the SonicWall POSTs the entire Guest Services session table. This allows the LHM server to synchronize the state of Guest Users for accounting, billing, or heuristics.

Can I change the LHM Management port from its default of TCP 4043?

Yes. This is easily done in SonicOS by modifying the port values of the External Guest Authentication Service Object.

Do I need to use the HMAC option? If I do want to use it, how do I use it?

The HMAC function is optional. It ensures that messages sent by the SonicWall to the LHM server and the LHM server to the SonicWall Security Appliance have not been tampered with. HMAC achieves this by calculating a keyed (password-aided) message authentication code on the information being passed between the two peers, and by adding that calculated digest to the data. Upon receiving the data, the other side calculates the digest itself, and compares it to the transmitted MAC; if the two match, the data was delivered intact. You should consider using the HMAC option if you are in an insecure environment or if you are concerned with security.

If you choose to use HMAC, you may implement your own HMAC routines, but the simplest method is to use the SonicWall-written `SonicSSL.dll` library, along with the `libeay32.dll`, which is freely available as part of OpenSSL; both are available from SonicWall by request.

To use HMAC:

1. Copy the `libeay32.dll` file to the path on the LHM (IIS) server (for example, into the `C:\Windows\system32` folder).
2. Copy the `SonicSSL.dll` file to any location on the same server.
3. Register the `SonicSSL.dll` file with the command `regsvr32 SonicSSL.dll`.

After this is done, the LHM scripts are able to use the `Server.CreateObject(SonicSSL.Crypto)` object for HMAC calculations. The HMAC functions are included in the scripts described in LHM Script Library on page 829.

IMPORTANT: The SonicWall Security Appliance URL Encodes (converts certain characters from their ASCII notation to hex notation) the `req` (originally requested URL) portion of the `querystring`, but the SonicWall method of URL encoding is slightly different from the Microsoft method (as employed by `Request.QueryString`, for example). Because of this difference in methods, it is possible for the string upon which the HMAC is being performed to be different between the Security Appliance and the LHM server. The provided scripts compensate for this by manually encoding the `req` portion of the `querystring` in a fashion consistent with the SonicWall method.

Does SonicWall provide any support for these scripts?

The scripts are provided as examples, and they are not supported by SonicWall Technical Support, nor can SonicWall support assist with the configuration of your LHM back-end environment. Future consultative support services might address this.

I've written a new script, I've made some great enhancements to your scripts, or I've just made your scripts work a whole lot better than you did; is SonicWall interested?

Yes! We are always looking for new ways to use LHM, and for people to contribute to the library of available scripts. We consider LHM scripts written on any platform, using any authentication method. Send an email to `products@SonicWall.com` describing your script, and do consider it for addition to our library. Submitting a script gives SonicWall permission to freely modify and/or redistribute the submitted script.

LHM Script Library

The SonicWall LHM Script library was established to serve as a resource for people using or wishing to use LHM for Guest Services. The goal is to attract multiple contributors and consumers, helping the library to grow to house a large, varied, and useful collection of scripts that anyone can modify or use as-is.

The first contribution to the library comprises six scripts: some in response to common user requests (`accept`, `guestbook`, and `adauth`), and some more uncommon (`lhmquiz`, `random`, and `paypal`). They were written outside of a Visual Studio .NET development environment, so their styles can be diverse. Common to all the scripts, however, are:

- Modularization of the configurable variables, such as the paths to files, server IP addresses, use of a popup logout window, salt values, and timer settings. These configurable values are gathered into the `myvars.aspx` file so per-environment editing can be done in one place rather than having to search for configurable elements.
- Extensive commentary explaining step-by-step what is being done.

A `chooser.aspx` landing page has been provided at the top-level of the scripts directory. This script was designed for demonstration environments to allow for the selection of a lower-level (specific) script without having to reconfigure the LHM settings on the SonicWall Security Appliance to point to a specific script. In other words, LHM on the Security Appliance can be configured to point to the top-level `chooser.aspx` script, which then enumerate all the sub-directories (lower-level scripts such as `random`, `accept`, `adauth`). The top-level `chooser.aspx` script opens the target lower-level `default.aspx` script in a new window, and passes the original `querystring` in its entirety.

All of the scripts begin with the `default.aspx` page, and client redirection is performed automatically as needed. The LHM configuration on the SonicWall should, therefore, point to the `default.aspx` page at the appropriate path (such as `lhm/accept/default.aspx` or `lhm/adauth/default.aspx`). Some scripts have separate administrative function page; these are noted in the script descriptions.

A `logout.aspx` page is also provided with each script. The use of this page is controllable with the `logoutPopup` variable in `myvars`. Setting a value of `1` enables the use of the popup logout window. The window is invoked by the LHM authentication process after a successful response code (50) is received from the Security Appliance. The script passes the `sessID`, `mgmtBaseUrl`, and `sessTimer` variables to the `logout.aspx` window so that the window can track the session time, and can POST a logout event back to the Security Appliance (at the `mgmtBaseUrl`) for the correct session (`sessID`) when/if the user wants to manually terminate the session.

About the Use of the Logout Popup Window

- The use of the logout popup is not necessary. Sessions timeout by themselves after their configured lifetime expires. The popup window simply provides users a mechanism to manually terminate their own sessions.
- The window launches with a javascript popup, so popup blockers block the window.

- Closing the window does not interrupt the session. Only **Logout** can end a session.
- Because the countdown timer runs client-side, steps have been taken to prevent refreshing the page. Refreshing the page resets the client-side countdown timer, but it does not affect the actual session timer. The F5 key and right-click mouse event are captured and suppressed, which does not work on all browsers.
- The use of the logout popup should agree with the nature of the scripts authentication scheme:
    - Some scripts have non-exclusive login processes, meaning that the user can login repeatedly (such as the `Accept` and `ADAuth` scripts). The use of the logout popup on these non-exclusive scripts is encouraged.
    - Some scripts are non-exclusive, but gather data that should be kept unique (such as the `Guestbook` and `LHMQuiz` scripts). The use of the logout popup on these scripts is acceptable, but can lead to redundant data being gathered.
    - Some scripts are exclusive, meaning that after the user authenticates, it is not possible to repeat the authentication process without some kind of cost (such as the `PayPal` script or the `Random` script where `useDB` is enabled). The use of the logout popup is discouraged on these scripts because the user has no simple means of logging back in.

The scripts also provide hidden output for a .NET procedure error, where the text is hidden by matching it to the color of the background. In the event of some kind of failure or error condition, error output may be provided and made visible by hitting CTRL-A on the web-page to select all of the text.

The following is a description of each of the scripts, what they do, and how they do it. As new scripts are added to the library, similar descriptions accompany them to help with understanding, customization, and integration.

Topics:

- Accept Script on page 830
- ADAuth Script on page 842
- Guestbook Script on page 855
- LHMQuiz Script on page 870
- PayPal Script on page 888
- Random Script on page 909
- Chooser.aspx Script on page 929

Accept Script

Authentication Model	The Guest Client clicks **I Accept**.	
Purpose	Present an acceptable use policy, terms of service, or welcome screen to the client.	
myvars Variables	`logoutPopup`	Controls the use of the logout popup window. Set to: • **0** to disable the popup window. • **1** to enable the popup window.
	`sessTimer`	The session timer in seconds.
	`idleTimer`	The idle timer in seconds.

	`username`	The username applied to the guest sessions. Because the script does not obtain a username from the client, it can be: • Explicitly set here for all clients. • Set to `useMAC` to set the username to the MAC address.
	`strHmac`	The shared secret for the optional HMAC function.
	`hmacType`	The digest type to use if HMAC is in use: **MD5** or **SHA1**.
	`logo`	The names of the logo (image) file to use on page headers.
Session Flow	1 The Guest Client clicks **I Accept**. 2 The LHM post string is assembled with the `sessionID`, the username (either default of MAC), the default session lifetime, and idle lifetime. 3 The script performs the LHM post to the SonicWall Security Appliance to authorize the session.	
Additional Considerations	Only the basic LHM configuration is required.	

Topics:

- default.aspx on page 831
- logout.aspx on page 836
- myvars.aspx on page 841

default.aspx

```
<%@ Page Language="VB" Debug="true" %>
<%@ Import Namespace="System" %>
<%@ Import Namespace="System.Net" %>
<%@ Import Namespace="System.IO" %>
<%@ Import Namespace="System.Xml" %>
<%@ Import Namespace="System.Text" %>
<%@ Import Namespace="System.Security" %>
<%@ Import namespace="System.Security.Cryptography.X509Certificates"%>

<!-- #INCLUDE file="myvars.aspx" -->

<script language="VB" runat="server">

'This class allows SSL certs signed by unknown CAs to be accepted.
'This is necessary for the POST to the SonicWall authorizing the LHM session.
Public Class acceptAllCerts
   Implements System.Net.ICertificatePolicy
   Public Function CheckValidationResult(ByVal srvPoint As ServicePoint, _
   ByVal cert As X509Certificate, ByVal request As WebRequest, ByVal problem As Integer) _
      As Boolean Implements ICertificatePolicy.CheckValidationResult
         Return True
   End Function
End Class

'Sample LHM redirect querystring:
'http://127.0.0.1/lhm/accept/default.aspx?sessionId=0b712fd83b9f5313db5af1cea6b1004f&ip=10.50.
165.231&mac=00:0e:35:bd:c9:37&ufi=0006b11184300&mgmtBaseUrl=https://10.50.165.193:4043/&client
RedirectUrl=https://10.50.165.193:444/&req=http%3A//www.google.com/ig

Dim ip as String
Dim sessionId as String
Dim mac as String
Dim ufi as String
Dim mgmtBaseUrl as String
Dim clientRedirectUrl as String
Dim req as String
```

```
Dim hmac as String
Dim customCode as String

Sub Page_Load(src as Object, e as EventArgs)

   LHMResult.Text=""
   catchError.Text=""

   ip=Request.QueryString("ip")
   sessionId=Request.QueryString("sessionId")
   mac=Request.QueryString("mac")
   ufi=Request.QueryString("ufi")
   mgmtBaseUrl=Request.QueryString("mgmtBaseUrl")
   clientRedirectUrl=Request.QueryString("clientRedirectUrl")
   req=Request.QueryString("req")
   hmac=Request.QueryString("hmac")
   customCode=Request.QueryString("cc")

   'customCode grabs the "cc=" querystring value sent by the SonicWall. This allows you to use
the same
   'page (e.g. this page) for the "Session Expiration" (?cc=2), "Idle Timeout" (?cc=3) and "Max
Sessions" (?cc=4) page.
   If customCode <> "" Then
      Select Case customCode
      Case "2"
         LHMResult.Text="<br><H3><font color=""red"">Your LHM session has expired. You may try
to initiate a new session.</font></H3>"
      Case "3"
         LHMResult.Text="<br><H3><font color=""red"">You have exceeded your idle timeout.
Please log back in.</font></H3>"
      Case "4"
         LHMResult.Text="<br><H3><font color=""red"">The maximum number of sessions has been
reached. Please try again later.</font></H3>"
      End Select
   End If

   'Set the userName to the grabbed client MAC address if so configured in myvars
   If userName = "useMAC" Then
      userName = mac
   End If

   'Use the override class in myvars.aspx to accept untrusted certificates from the SonicWall
   'This is necessary for the POST to the SonicWall authorizing the LHM session.
   System.Net.ServicePointManager.CertificatePolicy = New acceptAllCerts

   'Note - the routine below for handling the hmac requires the use of the SonicSSL.dll and
libeay.dll libraries.
   'The DLL must be copied to the IIS server, and the SonicSSL dll must be registered with
"regsvr32 sonicssl.dll"
   If hmac <> "" Then

      'SonicWall URL Encode routine is different from Microsoft - this is the SonicWall method
      req=Replace(req,"%","%25")
      req=Replace(req,":","%3A")
      req=Replace(req," ","%20")
      req=Replace(req,"?","%3F")
      req=Replace(req,"+","%2B")
      req=Replace(req,"&","%26")
      req=Replace(req,"=","%3D")

      Dim strHmacText as String
      Dim objCrypto as Object
      Dim strHmacGenerated
      Dim loginError as String

      'Initialize the Crypto object
      objCrypto = Server.CreateObject("SonicSSL.Crypto")
```

```
      'The text to be encoded
      strHmacText = sessionId & ip & mac & ufi & mgmtBaseUrl & clientRedirectUrl & req

      'Calculate the hash with a key strHmac, the return value is a string converted form the
output sha1 binary.
      'The hash algorithm (MD5 or SHA1) is configured in myvars and in the Extern Guest Auth
config on the SonicWall
      If hmacType = "MD5" Then
         strHmacGenerated = objCrypto.hmac_md5(strHmacText, strHmac)
      Else
         strHmacGenerated = objCrypto.hmac_sha1(strHmacText, strHmac)
      End If

      If strHmacGenerated <> hmac Then
         Dim hmacFail as String
         hmacFail = "<font color=""red"">The HMAC failed validation. Please notify an
attendant.</font><br><br>"
         hmacFail+="<font color=""9CBACE"">Received HMAC: " & hmac & "<br>Calculated HMAC: " &
strHmacGenerated & "<br>"
         hmacFail+="Make sure the digest functions on the SonicWall and LHM server match.<br>"
         hmacFail+="Also make sure the shared secret on the SonicWall and myvars match</font>"
         catchError.Text=hmacFail
      End If

   End If

End Sub

Sub btnSubmit_Click(Sender As Object, E As EventArgs)

   'Let the user know that we are setting up the session, just in case it takes more than a
second
   LHMResult.Text = "Authorizing session. Please wait."

   'The LHM cgi on the SonicWall - this does not change
   Dim loginCgi as String = "externalGuestLogin.cgi"

   'Assemble the data to post back to the SonicWall to authorize the LHM session
   Dim loginParams as String = "sessId=" & sessionId & "&userName=" &
Server.URLEncode(userName) & "&sessionLifetime=" & sessTimer & "&idleTimeout=" & idleTimer

   'Combine mgmtBaseUrl from the original redirect with the login cgi
   Dim postToSNWL as String = mgmtBaseUrl & loginCgi

   'Convert the loginParams to a well behaved byte array
   Dim byteArray As Byte() = Encoding.UTF8.GetBytes(loginParams)

   Try
      'Create the webrequest to the SonicWall
      Dim toSNWL as WebRequest = WebRequest.Create(postToSNWL)

      'Calculate the length of the byte array
      toSNWL.ContentLength = byteArray.Length

      'Set the method for the webrequest to POST
      toSNWL.Method = "POST"

      'Set the content type
      toSNWL.ContentType = "application/x-www-form-urlencoded"

      'Open the request stream
      Dim dataStream As Stream = toSNWL.GetRequestStream()

      'Write the byte array to the request stream
      dataStream.Write(byteArray, 0, byteArray.Length)
```

```
      'Close the Stream object
      dataStream.Close()

      'Get the response
      Dim snwlReply As WebResponse = toSNWL.GetResponse()

      'Display the status - looking for 200 = OK.
      'Response.Write(CType(snwlReply, HttpWebResponse).StatusCode)

      'Grab the response and stuff it into an xml doc for possible review
      Dim snwlResponse as XmlDocument = New XmlDocument()
      snwlResponse.Load(snwlReply.GetResponseStream())

      'Set the xPath to the SNWL reply, and get the response
      Dim codePath as String = "SonicWallAccessGatewayParam/AuthenticationReply/ResponseCode"

      'Response.Write(snwlResponse.SelectSingleNode(codePath).InnerXml)

      'Response code 50 - Login Succeeded
      If snwlResponse.SelectSingleNode(codePath).InnerXml = "50"

         'Do we want to provide a logout popup window?
         If logoutPopup = "1" Then
            'Popup hack using Javascript for logout window
            Dim sb As New System.Text.StringBuilder()
            sb.Append("<script language='javascript'>")
            sb.Append("window.open('logout.aspx?sessId=")
            sb.Append(Server.URLEncode(CStr(sessionId)))
            sb.Append("&mgmtBaseUrl=")
            sb.Append(Server.URLEncode(CStr(mgmtBaseUrl)))
            sb.Append("&sessTimer=")
            sb.Append(Server.URLEncode(CStr(sessTimer)))
            sb.Append("','logOut','toolbar=no,")
            sb.Append("addressbar=no,menubar=no,")
            sb.Append("width=400,height=250');")
            sb.Append("<")
            sb.Append("/")
            sb.Append("script>")
            RegisterStartupScript("stp", sb.ToString)
         End If

         LHMResult.Text = "<br><b><font color=""green"">Session Authorized:</font></b> You may
now go to the URL you originally requested: <a target=""_blank"" href=""" & req & """>" & req &
"</a>"

      'Response code 51 - Session Limit Exceeded
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "51"
         LHMResult.Text = "<br><b><font color=""red"">Session Limit Reached:</font></b> The
maximum number of guest session has been reached. Sorry for the inconvenience. Please close and
relaunch your browser to try again."

      'Response code 100 - Login Failed.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "100"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> Your
session cannot be created at this time. Sorry for the inconvenience. Please close and relaunch
your browser to try again."

      'Response code 251 - Bad HMAC.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "251"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization failed message authentication. Sorry for the inconvenience. Please
close and relaunch your browser to try again."

      'Response code 253 - Invalid SessionID.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "253"
```

```
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization failed to match a known session identity. Sorry for the
inconvenience. Please close and relaunch your browser to try again."

      'Response code 254 - Invalid CGI.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "254"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization was missing an essential parameter. Sorry for the inconvenience.
Please close and relaunch your browser to try again."

      'Response code 255 - Internal Error.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "255"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization failed due to an unspecified error. Sorry for the inconvenience.
Please close and relaunch your browser to try again."

      End If

      'Close the streams
      dataStream.Close()
      snwlReply.Close()

      'If there is some asp.net error trying to talk to the SonicWall, print it in the same
color as the background.
      Catch ex as Exception
         catchError.Text = "<font color=""9CBACE"">" & ex.ToString & "</font>"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization failed due to an unspecified error. Sorry for the inconvenience.
Please close and relaunch your browser to try again. If the problem persists, please notify an
attendant."
      End Try
End Sub


</script>

<STYLE>
body {
  font-size: 10pt;
  font-family: verdana,helvetica,arial,sans-serif;
  color:#000000;
  background-color:#9CBACE;
}

tr.heading {
  background-color:#006699;
}

.button {
   border: 1px solid #000000;
   background-color: #ffffff;
}
</STYLE>

<HTML>
<HEAD>
<TITLE>LHM Accept Script</TITLE>
</HEAD>

<BODY>
<form id="frmValidator" runat="server">

<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr class="heading">
      <td colspan=3 align="center"><font color="white"> </td>
   </tr>
   <tr class="heading">
```

```
      <td width="50%" valign="center"><font color="white"><b>Welcome <%= ip%></b></font></td>
      <td><center><img width="216" height="51" src="<%= logo%>"></center></td>
      <td width="50%" align="right" valign="center"><font color="white"><b>Powered by
SonicWall LHM</b> </font></td>
   </tr>
   <tr class="heading">
      <td colspan=3 align="center"><font color="white"> </td>
   </tr>
</table>

<table width="100%" border="0" cellpadding="2" cellspacing="0">

   <tr>
   <td><br></td>
   </tr>
   <tr>
      <td align=left>
      By clicking the <b>Accept</b> button below, you accept the following terms of
service:<br><br><b>
      1. You will not try to download bad things.<br>
      2. You will not try to upload bad things.<br>
      3. You will not try to use all the bandwidth so that others have none.<br>
      4. You will be happy when the SonicWall blocks bad things from reaching you.</b><br><br>
      </td>
      <td>
   </tr>
   <tr>
      <td><br><asp:button id="btnSubmit" class="button" text="  Accept  "
onClick="btnSubmit_Click" runat="server" /></td>
   </tr>
   <tr>
      <td><asp:Label id=LHMResult runat="server" /></td>
   </tr>
   <tr>
      <td><asp:Label id=catchError runat="server" /></td>
   </tr>
</table>
</form>
</BODY>
</HTML>
```

logout.aspx

```
<%@ Page Language="VB" Debug="true" %>
<%@ Import Namespace="System" %>
<%@ Import Namespace="System.Net" %>
<%@ Import Namespace="System.IO" %>
<%@ Import Namespace="System.Xml" %>
<%@ Import Namespace="System.Text" %>
<%@ Import Namespace="System.Security" %>
<%@ Import namespace="System.Security.Cryptography.X509Certificates"%>

<!-- #INCLUDE file="myvars.aspx" -->

<script language="VB" runat="server">

'This class allows SSL certs signed by unknown CAs to be accepted.
'This is necessary for the POST to the SonicWall authorizing the LHM session.
Public Class acceptAllCerts
   Implements System.Net.ICertificatePolicy
   Public Function CheckValidationResult(ByVal srvPoint As ServicePoint, _
   ByVal cert As X509Certificate, ByVal request As WebRequest, ByVal problem As Integer) _
      As Boolean Implements ICertificatePolicy.CheckValidationResult
         Return True
   End Function
End Class
```

```
Dim sessionId as String
Dim mgmtBaseUrl as String
Dim eventId as String = "&eventId=1"

'Grab the code and the session lifetime from the generator page
Sub Page_Load(src as Object, e as EventArgs)
   sessionId=Request.QueryString("sessId")
   mgmtBaseUrl=Request.QueryString("mgmtBaseUrl")
   sessTimer=Request.QueryString("sessTimer")


   'Use the override class in myvars.aspx to accept untrusted certificates from the SonicWall
   'This is necessary for the POST to the SonicWall authorizing the LHM session.
   System.Net.ServicePointManager.CertificatePolicy = New acceptAllCerts

   'When the page loads, make the loggedIn span visible
   loggedIn.Visible=True
   loggedOut.Visible=False

   Me.Button1.Attributes.Add("OnClick", "self.close()")

End Sub

'The Logout button
Sub btnSubmit_Click(Sender As Object, E As EventArgs)

   'Let the user know that we are setting up the session, just in case it takes more than a
second
   LHMResult.Text = "Authorizing session. Please wait."

   'The LHM cgi on the SonicWall - this does not change
   Dim loginCgi as String = "externalGuestLogoff.cgi"

   'Assemble the data to post back to the SonicWall to authorize the LHM session
   Dim loginParams as String = "sessId=" & sessionId & eventId

   'Combine mgmtBaseUrl from the original redirect with the login cgi
   Dim postToSNWL as String = mgmtBaseUrl & loginCgi

   'Convert the loginParams to a well behaved byte array
   Dim byteArray As Byte() = Encoding.UTF8.GetBytes(loginParams)

   Try
      'Make the loggedOut span visible
      loggedIn.Visible=False
      loggedOut.Visible=True

      'Create the webrequest to the SonicWall
      Dim toSNWL as WebRequest = WebRequest.Create(postToSNWL)

      'Calculate the length of the byte array
      toSNWL.ContentLength = byteArray.Length

      'Set the method for the webrequest to POST
      toSNWL.Method = "POST"

      'Set the content type
      toSNWL.ContentType = "application/x-www-form-urlencoded"

      'Open the request stream
      Dim dataStream As Stream = toSNWL.GetRequestStream()

      'Write the byte array to the request stream
      dataStream.Write(byteArray, 0, byteArray.Length)

      'Close the Stream object
```

```
      dataStream.Close()

      'Get the response
      Dim snwlReply As WebResponse = toSNWL.GetResponse()

      'Display the status - looking for 200 = OK.
      'Response.Write(CType(snwlReply, HttpWebResponse).StatusCode)

      'Grab the response and stuff it into an xml doc for possible review
      Dim snwlResponse as XmlDocument = New XmlDocument()
      snwlResponse.Load(snwlReply.GetResponseStream())

      'Set the xPath to the SNWL reply, and get the response
      Dim codePath as String = "SonicWallAccessGatewayParam/LogoffReply/ResponseCode"

      'Response.Write(snwlResponse.SelectSingleNode(codePath).InnerXml)

      'Response code 150 - Logout Succeeded
      If snwlResponse.SelectSingleNode(codePath).InnerXml = "150"
         LHMResult.Text = "<br><b><font color=""green"">Your session has been logged
out.<br><br>Thank you for using LHM Guest Services.</font></b>"

      'Response code 251 - Bad HMAC.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "251"
         LHMResult.Text = "<br><b><font color=""red"">Session logout failed:</font></b> The
request failed message authentication. Sorry for the inconvenience. Please close and relaunch
your browser to try again."

      'Response code 253 - Invalid SessionID.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "253"
         LHMResult.Text = "<br><b><font color=""red"">Session logout failed:</font></b> The
request failed to match a known session identity. Sorry for the inconvenience. Please close and
relaunch your browser to try again."

      'Response code 254 - Invalid CGI.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "254"
         LHMResult.Text = "<br><b><font color=""red"">Session logout failed:</font></b> The
request was missing an essential parameter. Sorry for the inconvenience. Please close and
relaunch your browser to try again."

      'Response code 255 - Internal Error.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "255"
         LHMResult.Text = "<br><b><font color=""red"">Session logout failed:</font></b> The
request failed due to an unspecified error. Sorry for the inconvenience. Please close and
relaunch your browser to try again."

      End If

      'Close the streams
      dataStream.Close()
      snwlReply.Close()

      'If there is some asp.net error trying to talk to the SonicWall, print it in the same
color as the background.
      Catch ex as Exception
         catchError.Text = "<font color=""9CBACE"">" & ex.ToString & "</font>"
         LHMResult.Text = "<br><b><font color=""red"">Session logout failed:</font></b> The
request failed due to an unspecified error. Sorry for the inconvenience. Please close and
relaunch your browser to try again. If the problem persists, please notify an attendant."
      End Try
End Sub

</script>
<STYLE>
body {
   font-size: 10pt;
   font-family: verdana,helvetica,arial,sans-serif;
```

```
   color:#000000;
   background-color:#9CBACE;
}

tr.heading {
   font-size: 10pt;
   background-color:#006699;
}

tr.smalltext {
   font-size: 8pt;
}

.button {
   border: 1px solid #000000;
   background-color: #ffffff;
   font-size: 8pt;
}
</STYLE>

<HTML>
<HEAD>
<TITLE>LHM Logout Page</TITLE>

<SCRIPT LANGUAGE="Javascript">

//'Javascript Seconds Countdown Timer
var SecondsToCountDown = <%= sessTimer%>;
var originalTime=" ";

function CountDown()
{
   clockStr="";

   dayStr=Math.floor(SecondsToCountDown/86400)%100000
   if(dayStr>0){
      if(dayStr>1){
         dayStr+=" days ";
      } else dayStr+=" day ";
      clockStr=dayStr;
   }
   hourStr=Math.floor(SecondsToCountDown/3600)%24
   if(hourStr>0){
      if(hourStr>1){
         hourStr+=" hours ";
      } else hourStr+=" hour ";
      clockStr+=hourStr;
   }
   minuteStr=Math.floor(SecondsToCountDown/60)%60
   if(minuteStr>0){
      if(minuteStr>1){
         minuteStr+=" minutes ";
      } else minuteStr+=" minute ";
      clockStr+=minuteStr;
   }
   secondStr=Math.floor(SecondsToCountDown/1)%60
   if(secondStr>0){
      if(secondStr>1){
         secondStr+=" seconds ";
      } else secondStr+=" second ";
      clockStr+=secondStr;
   }

   if(SecondsToCountDown > 0)
   {
      --SecondsToCountDown;
   }
```

```
   if(originalTime.length < 2)
   {
      originalTime = clockStr;
   }

   // Make sure the form is still there before trying to set a value
   if(document.frmValidator){
      document.frmValidator.originalTime.value = originalTime;
      document.frmValidator.countdown.value = clockStr;
   }

   setTimeout("CountDown()", 1000);
   if(SecondsToCountDown == 0)
   {
      document.frmValidator.countdown.value = "Session Expired";
   }
}

//'Disable right-click so that the window doesn't get refreshed since the countdown is
clientside.
document.oncontextmenu = disableRightClick;
function disableRightClick()
{
  return false;
}

//'Disable F5 key, too, on IE at least.
function noF5()
{
   var key_f5 = 116;
   if (key_f5==event.keyCode)
   {
      event.keyCode=0;
      return false;
   }
   return false;
}

document.onkeydown=noF5
document.onmousedown=disableRightClick

</SCRIPT>

</HEAD>

<BODY onload='CountDown()'>
<span id="loggedIn" runat="server">
<form id="frmValidator" runat="server">
<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr class="heading">
      <td colspan=2 align="center"> </td>
   </tr>
   <tr class="heading">
      <td colspan=2 align="center"><font color="white"><b>SonicWall LHM Logout
Window</b></font></td>
   </tr>
   <tr class="heading">
      <td colspan=2 align="center"> </td>
   </tr>
   <tr class="smalltext"><td><br></td></tr>
   <tr class="smalltext">
      <td>Original Session Time:</td>
      <td><asp:textbox width=250 id="originalTime" runat="server" /></td>
   </tr>
   <tr class="smalltext">
      <td>Remaining Session Time:</td>
```

```
      <td><asp:textbox width=250 id="countdown" runat="server" /></td>
   </tr>
   <tr class="smalltext">
      <td colspan=2><br>You may use this window to manually logout your session at any time, or
you may safely close this window if you prefer to let your session timeout
automatically.</font></td>
   </td>
   <tr>
      <td colspan=2><center><asp:button id="btnSubmit" class="button" text="  Logout  "
onClick="btnSubmit_Click" runat="server" /></center></td>
   </tr>
</table>
</form>
</span>

<span id="loggedOut" runat="server">
<form id="logout" runat="server">
<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr class="heading">
      <td colspan=2 align="center"> </td>
   </tr>
   <tr class="heading">
      <td colspan=2 align="center"><font color="white"><b>SonicWall LHM Logout
Window</b></font></td>
   </tr>
   <tr class="heading">
      <td colspan=2 align="center"> </td>
   </tr>
   <tr>
      <td><asp:Label id=LHMResult runat="server" /></td>
   </tr>
   <tr>
      <td><asp:Label id=catchError runat="server" /></td>
   </tr>
   <tr><td><br></td></tr>
   <tr>
      <td><center><asp:button id="Button1" class="button" text="  Close  " runat="server"
/></center></td>
   </tr>
</table>
</form>
</span>

</BODY>
</HTML>
```

myvars.aspx

```
<script language="VB" runat="server">

'Set the logoutPopup window flag - 0 = no popup, 1 = popup
'The use of the logoutPopup in this script is encouraged because the login event is non-
exclusive.
Dim logoutPopup as String = "1"

'Set the LHM Session Timeout
Dim sessTimer as String = "3600"

'Set the LHM Idle Timeout
Dim idleTimer as String = "300"

'Set the username to record for LHM session since this does not gather one. Set to
userName="useMAC" to use the MAC address.
Dim userName="useMAC"
'Dim userName = "LHM Guest User"
```

```
'Set the secret for use with optional HMAC auth, as configured in the Extern Guest Auth config
on the SonicWall
Dim strHmac as String = "password"

'Set the digest method for the HMAC, either MD5 or SHA1
Dim hmacType as String = "MD5"
'Dim hmacType as String = "SHA1"

'Set the logo image to use
Dim logo as String = "SonicWall.gif"

'-----------------------End of Configurable Settings-------------------------

</script>
```

ADAuth Script

Authentication Model	The Guest Client provides their username and password. These credentials are then authenticated against an Active Directory or LDAP database.	
Purpose	Classical authorization model using Active Directory through LDAP. Support for per-user session-timer and idle-timer setting provided by optionally grabbing LDAP attributes from the database during authorization.	
myvars Variables	`logoutPopup`	Controls the use of the logout popup window. Set to: • **0** to disable the popup window. • **1** to enable the popup window.
	`myLdapServer`	The IP address or resolvable FQDN of the LDAP/AD server providing authentication.
	`myLdapDomain`	The LDAP/AD domain name
	`retrAttr`	Specifies whether to retrieve session and idle timer values from the authenticating user's LDAP `attributes` (defined later). Set to: • **0** to disable retrieval. • **1** to attempt retrieval.
	`useCN`	If `reAttr`=1, then this flag sets whether to use the common name (`cn`) to retrieve attributes, or the AD default login name (`sAMAccountName`). Set to **1** to use `cn`. When authenticating against AD, this flag should be set to **0**.
	`sessAttr`	The LDAP attribute from which to retrieve the session timer (in seconds). If no value can be retrieved, or if the retrieved value is not numeric, the default session timer (`sessTimer`, defined below) are used.
	`idleAttr`	The LDAP attribute from which to retrieve the idle timer (in seconds). If no value can be retrieved, or if the retrieved value is not numeric, the default idle timer (`idleTimer`, defined below) are used.
	`sessTimer`	The default session timer in seconds.
	`idleTimer`	The default idle timer in seconds.
	`strHmac`	The shared secret for the optional HMAC function.
	`hmacType`	The digest type to use if HMAC is in use: **MD5** or **SHA1**.
	`logo`	The names of the logo (image) file to use on page headers.

Session Flow	1 The Guest Client enters their LDAP/AD username and password. 2 The provided credentials are used to bind with the configured LDAP server. 3 If the bind attempt succeeds, the user is authenticated. 4 If the `reAttr` flag is set, an attempt is made to retrieve the defined `sessAttr` and `idleAttr` attributes (such as `pager` and `mobile`) from the LDAP DB. If valid results are retrieved, they are used; otherwise the default values are used. 5 The script performs the LHM post to the SonicWall Security Appliance to authorize the session.
Additional Considerations	Requires that the LHM server be able to communicate with the configured LDAP/AD server, either by route, NAT, or VPN. If the `reAttr` option is used, it requires that the LDAP attributes be defined for user-specific values to take effect. **NOTE:** The `pager` and `mobile` attributes were selected because they are not frequently used, and because they can be set directly through Microsoft's Users and Computers MMC.)

Topics:

- default.aspx on page 843
- logout.aspx on page 850
- myvars.aspx on page 855

default.aspx

```
<%@ Page Language="VB" Debug="true" %>
<%@ Import Namespace="System" %>
<%@ Import Namespace="System.Net" %>
<%@ Import Namespace="System.IO" %>
<%@ Import Namespace="System.Xml" %>
<%@ Import Namespace="System.Text" %>
<%@ Import Namespace="System.Math" %>
<%@ Import Namespace="System.DirectoryServices" %>
<%@ Import Namespace="System.Collections" %>
<%@ Import Namespace="System.Security" %>
<%@ Import namespace="System.Security.Cryptography.X509Certificates"%>
<%@ Assembly name="System.DirectoryServices, Version=1.0.3300.0,
Culture=neutral,PublicKeyToken=b03f5f7f11d50a3a"%>

<!-- #INCLUDE file="myvars.aspx" -->

<script language="VB" runat="server">

'This class allows SSL certs signed by unknown CAs to be accepted.
'This is necessary for the POST to the SonicWall authorizing the LHM session.
Public Class acceptAllCerts
   Implements System.Net.ICertificatePolicy
   Public Function CheckValidationResult(ByVal srvPoint As ServicePoint, _
   ByVal cert As X509Certificate, ByVal request As WebRequest, ByVal problem As Integer) _
      As Boolean Implements ICertificatePolicy.CheckValidationResult
      Return True
   End Function
End Class

'Sample LHM redirect querystring:
'http://127.0.0.1/lhm/adauth/default.aspx?sessionId=0b712fd83b9f5313db5af1cea6b1004f&ip=10.50.
165.231&mac=00:0e:35:bd:c9:37&ufi=0006b11184300&mgmtBaseUrl=https://10.50.165.193:4043/&client
RedirectUrl=https://10.50.165.193:444/&req=http%3A//www.google.com/ig

Dim ip as String
Dim sessionId as String
```

```
Dim mac as String
Dim ufi as String
Dim mgmtBaseUrl as String
Dim clientRedirectUrl as String
Dim req as String
Dim hmac as String
Dim customCode as String

Sub Page_Load(src as Object, e as EventArgs)

   LHMResult.Text=""
   catchError.Text=""
   authResult.Text=""

   ip=Request.QueryString("ip")
   sessionId=Request.QueryString("sessionId")
   mac=Request.QueryString("mac")
   ufi=Request.QueryString("ufi")
   mgmtBaseUrl=Request.QueryString("mgmtBaseUrl")
   clientRedirectUrl=Request.QueryString("clientRedirectUrl")
   req=Request.QueryString("req")
   hmac=Request.QueryString("hmac")
   customCode=Request.QueryString("cc")

   'customCode grabs the "cc=" querystring value sent by the SonicWall. This allows you to use
the same
   'page (e.g. this page) for the "Session Expiration" (?cc=2), "Idle Timeout" (?cc=3) and "Max
Sessions" (?cc=4) page.
   If customCode <> "" Then
      Select Case customCode
      Case "2"
         LHMResult.Text="<br><H3><font color=""red"">Your LHM session has expired. You may try
to initiate a new session.</font></H3>"
      Case "3"
         LHMResult.Text="<br><H3><font color=""red"">You have exceeded your idle timeout.
Please log back in.</font></H3>"
      Case "4"
         LHMResult.Text="<br><H3><font color=""red"">The maximum number of sessions has been
reached. Please try again later.</font></H3>"
      End Select
   End If

   'Use the override class in myvars.aspx to accept untrusted certificates from the SonicWall
   'This is necessary for the POST to the SonicWall authorizing the LHM session.
   System.Net.ServicePointManager.CertificatePolicy = New acceptAllCerts

   'Note - the routine below for handling the hmac requires the use of the SonicSSL.dll and
libeay.dll libraries.
   'The DLL must be copied to the IIS server, and the SonicSSL dll must be registered with
"regsvr32 sonicssl.dll"
   If hmac <> "" Then

      'SonicWall URL Encode routine is different from Microsoft - this is the SonicWall method
      req=Replace(req,"%","%25")
      req=Replace(req,":","%3A")
      req=Replace(req," ","%20")
      req=Replace(req,"?","%3F")
      req=Replace(req,"+","%2B")
      req=Replace(req,"&","%26")
      req=Replace(req,"=","%3D")

      Dim strHmacText as String
      Dim objCrypto as Object
      Dim strHmacGenerated
      Dim loginError as String

      'Initialize the Crypto object
```

```
      objCrypto = Server.CreateObject("SonicSSL.Crypto")

      'The text to be encoded
      strHmacText = sessionId & ip & mac & ufi & mgmtBaseUrl & clientRedirectUrl & req

      'Calculate the hash with a key strHmac, the return value is a string converted form the
output sha1 binary.
      'The hash algorithm (MD5 or SHA1) is configured in myvars and in the Extern Guest Auth
config on the SonicWall
      If hmacType = "MD5" Then
         strHmacGenerated = objCrypto.hmac_md5(strHmacText, strHmac)
      Else
         strHmacGenerated = objCrypto.hmac_sha1(strHmacText, strHmac)
      End If

      If strHmacGenerated <> hmac Then
         Dim hmacFail as String
         hmacFail = "<font color=""red"">The HMAC failed validation. Please notify an
attendant.</font><br><br>"
         hmacFail+="<font color=""9CBACE"">Received HMAC: " & hmac & "<br>Calculated HMAC: " &
strHmacGenerated & "<br>"
         hmacFail+="Make sure the digest functions on the SonicWall and LHM server match.<br>"
         hmacFail+="Also make sure the shared secret on the SonicWall and myvars match</font>"
         catchError.Text=hmacFail
      End If

   End If

End Sub

sub OnBtnClearClicked (Sender As Object, e As EventArgs)
   txtName.Text = ""
   txtPassword.Text = ""
   authResult.Text=""
   LHMResult.Text=""
   catchError.Text=""
end sub

Sub btnSubmit_Click(Sender As Object, E As EventArgs)

   'Try to connect to LDAP with the user supplied attributes
   Try
      Dim ldapPath as String = "LDAP://" & myLdapServer
      Dim ldapUser as String = myLdapDomain & "\" & txtName.Text
      Dim validateUser as New DirectoryEntry(ldapPath,ldapUser,txtPassword.Text)

      'This is the actual authentication piece
      Dim nativeCheck as Object = validateUser.NativeObject

      'If retrAttr is set in the myvars file, attempt to retrieve the session and idle values
from LDAP
      If retrAttr = "1" Then
         Dim mySearch as New DirectorySearcher(validateUser)

         'Check the myvars for selecting either sAMAccountName or cn
         If useCN = "0" Then
            mySearch.Filter = "(sAMAccountName=" & Server.URLEncode(txtName.Text) & ")"
         Else
            mySearch.Filter = "(cn=" & Server.URLEncode(txtName.Text) & ")"
         End If
         mySearch.PageSize="1"
         mySearch.PropertiesToLoad.Add(sessAttr)
         mySearch.PropertiesToLoad.Add(idleAttr)
         Dim adResult as SearchResult

         'If we get results on the attribute query, set timer values
         adResult = mySearch.FindOne
```

```
         If Not (adResult is Nothing) Then
            If (adResult.Properties.Contains(sessAttr)) Then
                  'Check to see if the LDAP value returned is a number
                  Dim isNumber as New RegEx("^\d+$")
                  If (isNumber.IsMatch(adResult.Properties(sessAttr)(0).ToString())) Then
                     sessTimer=adResult.Properties(sessAttr)(0).ToString()
                  End If
            End If 'End If sessAttr
            If (adResult.Properties.Contains(idleAttr)) Then
                  'Check to see if the LDAP value returned is a number
                  Dim isNumber as New RegEx("^\d+$")
                  If (isNumber.IsMatch(adResult.Properties(idleAttr)(0).ToString())) Then
                     idleTimer=adResult.Properties(idleAttr)(0).ToString()
                  End If
            End If 'End if idleAttr
         End If 'End if adResult is present
      End If 'End if retrAttr is in use

      authResult.Text="<font color=""green""><b>Credentials Accepted.</b></font><br>Session
Lifetime: " & round(sessTimer/60) & " minutes.<br>Idle Timer: " & round(idleTimer/60) & "
minutes."

      'Auth succeeded - move on to LHM Auth
      LHM()

      Catch ex as Exception
         catchError.Text = "<font color=""9CBACE"">" & ex.ToString & "</font>"
         authResult.Text="<font color=""Red""><b>Credentials Rejected.</b></font><br>Please
enter a valid username and password. "
   End Try

End Sub

Sub LHM()

      'Let the user know that we are setting up the session, just in case it takes more than a
second
      LHMResult.Text = "Authorizing session. Please wait."

      'The LHM cgi on the SonicWall - this does not change
      Dim loginCgi as String = "externalGuestLogin.cgi"

      'Assemble the data to post back to the SonicWall to authorize the LHM session
      Dim loginParams as String = "sessId=" & sessionId & "&userName=" &
Server.URLEncode(txtName.Text) & "&sessionLifetime=" & sessTimer & "&idleTimeout=" & idleTimer

      'Combine mgmtBaseUrl from the original redirect with the login cgi
      Dim postToSNWL as String = mgmtBaseUrl & loginCgi

      'Convert the loginParams to a well behaved byte array
      Dim byteArray As Byte() = Encoding.UTF8.GetBytes(loginParams)

      Try
         'Create the webrequest to the SonicWall
         Dim toSNWL as WebRequest = WebRequest.Create(postToSNWL)

         'Calculate the length of the byte array
         toSNWL.ContentLength = byteArray.Length

         'Set the method for the webrequest to POST
         toSNWL.Method = "POST"

         'Set the content type
         toSNWL.ContentType = "application/x-www-form-urlencoded"

         'Open the request stream
         Dim dataStream As Stream = toSNWL.GetRequestStream()
```

```
         'Write the byte array to the request stream
         dataStream.Write(byteArray, 0, byteArray.Length)

         'Close the Stream object
         dataStream.Close()

         'Get the response
         Dim snwlReply As WebResponse = toSNWL.GetResponse()

         'Display the status - looking for 200 = OK.
         'Response.Write(CType(snwlReply, HttpWebResponse).StatusCode)

         'Grab the response and stuff it into an xml doc for possible review
         Dim snwlResponse as XmlDocument = New XmlDocument()
         snwlResponse.Load(snwlReply.GetResponseStream())

         'Set the xPath to the SNWL reply, and get the response
         Dim codePath as String =
"SonicWallAccessGatewayParam/AuthenticationReply/ResponseCode"

         'Response.Write(snwlResponse.SelectSingleNode(codePath).InnerXml)

         'Response code 50 - Login Succeeded
         If snwlResponse.SelectSingleNode(codePath).InnerXml = "50"

            'Do we want to provide a logout popup window?
            If logoutPopup = "1" Then
               'Popup hack using Javascript for logout window
               Dim sb As New System.Text.StringBuilder()
               sb.Append("<script language='javascript'>")
               sb.Append("window.open('logout.aspx?sessId=")
               sb.Append(Server.URLEncode(CStr(sessionId)))
               sb.Append("&mgmtBaseUrl=")
               sb.Append(Server.URLEncode(CStr(mgmtBaseUrl)))
               sb.Append("&sessTimer=")
               sb.Append(Server.URLEncode(CStr(sessTimer)))
               sb.Append("','logOut','toolbar=no,")
               sb.Append("addressbar=no,menubar=no,")
               sb.Append("width=400,height=250');")
               sb.Append("<")
               sb.Append("/")
               sb.Append("script>")
               RegisterStartupScript("stp", sb.ToString)
            End If

            LHMResult.Text = "<br><b><font color=""green"">Session authorized:</font></b> You
may now go to the URL you originally requested: <a target=""_blank"" href=""" & req & """>" &
req & "</a>"

         'Response code 51 - Session Limit Exceeded
         ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "51"
            LHMResult.Text = "<br><b><font color=""red"">Session Limit Reached:</font></b> The
maximum number of guest session has been reached. Sorry for the inconvenience. Please close and
relaunch your browser to try again."

         'Response code 100 - Login Failed.
         ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "100"
            LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b>
Your session cannot be created at this time. Sorry for the inconvenience. Please close and
relaunch your browser to try again."

         'Response code 251 - Bad HMAC.
         ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "251"
            LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b>
The request for authorization failed message authentication. Sorry for the inconvenience.
Please close and relaunch your browser to try again."
```

```
         'Response code 253 - Invalid SessionID.
         ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "253"
            LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b>
The request for authorization failed to match a known session identity. Sorry for the
inconvenience. Please close and relaunch your browser to try again."

         'Response code 254 - Invalid CGI.
         ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "254"
            LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b>
The request for authorization was missing an essential parameter. Sorry for the inconvenience.
Please close and relaunch your browser to try again."

         'Response code 255 - Internal Error.
         ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "255"
            LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b>
The request for authorization failed due to an unspecified error. Sorry for the inconvenience.
Please close and relaunch your browser to try again."

         End If

         'Close the streams
         dataStream.Close()
         snwlReply.Close()

         'If there is some asp.net error trying to talk to the SonicWall, print it in the same
color as the background.
         Catch ex as Exception
            catchError.Text = "<font color=""9CBACE"">" & ex.ToString & "</font>"
            LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b>
The request for authorization failed due to an unspecified error. Sorry for the inconvenience.
Please close and relaunch your browser to try again. If the problem persists, please notify an
attendant."
      End Try
End Sub
</script>

<STYLE>
body {
  font-size: 10pt;
  font-family: verdana,helvetica,arial,sans-serif;
  color:#000000;
  background-color:#9CBACE;
}

tr.heading {
  background-color:#006699;
}

.button {
   border: 1px solid #000000;
   background-color: #ffffff;
}
</STYLE>

<HTML>
<HEAD>
<TITLE>LHM ADAuth Script</TITLE>
</HEAD>

<BODY>
<form id="frmValidator" runat="server">

<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr class="heading">
      <td colspan=3 align="center"><font color="white"> </td>
   </tr>
```

```
   <tr class="heading">
      <td width="50%" valign="center"><font color="white"><b>LDAP/AD LHM
Authentication</b></font></td>
      <td><center><img width="216" height="51" src="<%= logo %>"></center></td>
      <td width="50%" align="right" valign="center"><font color="white"><b>Powered by
SonicWall LHM</b> </font></td>
   </tr>
   <tr class="heading">
      <td colspan=3 align="center"><font color="white"> </td>
   </tr>
</table>

<table width="90%" border="0" cellpadding="2" cellspacing="0">
   <tr>
      <td><b>Welcome <%= ip%> to SonicWall's LHM AD/LDAP Authenticator.</b><br><br>Enter your
LDAP or Active Directory username and password to obtain secure guest Internet
access.<br><br>If your domain account specifies session timeout values, those values will be
applied to your account, otherwise you will receive the default one hour (60 minutes) of access
with a five minute idle timeout.<br>
      </td>
   </tr>
</table>
<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr class="heading">
      <td colspan=3 align="center"><font color="white">Authentication domain:
<%=myLdapDomain%></td>
   </tr>
</table>
<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr>
      <td width="30%"><br>Enter your login name:</td>
      <td width="30%"><asp:TextBox id="txtName" runat="server" /></td>
      <td width="40%"><asp:RequiredFieldValidator id="valTxtName" ControlToValidate="txtName"
ErrorMessage="Please enter your name." runat="server" /></td>
   </tr>
   <tr>
      <td width="30%"><br>Enter your password:</td>
      <td width="30%"><asp:TextBox id="txtPassword" textmode="password" runat="server" /></td>
      <td width="40%"><asp:RequiredFieldValidator id="valTxtPassword"
ControlToValidate="txtPassword" ErrorMessage="Please enter your password." runat="server"
/></td>
   </tr>
   <tr>
      <td></td><td><asp:Label id=authResult runat="server" /> </td>
   </tr>
   <tr>
      <td></td>
      <td><asp:button id="btnSubmit" class="button" text="  Submit  "
onClick="btnSubmit_Click" runat="server" />
        
      <asp:button id="btnClear" class="button" text=" Clear All " CausesValidation="False"
onClick="OnBtnClearClicked" runat="server" />
      </td>
   </tr>
   <tr>
      <td colspan=2><asp:Label id=LHMResult runat="server" /></td>
   </tr>
   <tr>
      <td colspan=2><asp:Label id=catchError runat="server" /></td>
   </tr>
</table>
</form>
</BODY>
</HTML>
```

logout.aspx

```
<%@ Page Language="VB" Debug="true" %>
<%@ Import Namespace="System" %>
<%@ Import Namespace="System.Net" %>
<%@ Import Namespace="System.IO" %>
<%@ Import Namespace="System.Xml" %>
<%@ Import Namespace="System.Text" %>
<%@ Import Namespace="System.Security" %>
<%@ Import namespace="System.Security.Cryptography.X509Certificates"%>

<!-- #INCLUDE file="myvars.aspx" -->

<script language="VB" runat="server">

'This class allows SSL certs signed by unknown CAs to be accepted.
'This is necessary for the POST to the SonicWall authorizing the LHM session.
Public Class acceptAllCerts
   Implements System.Net.ICertificatePolicy
   Public Function CheckValidationResult(ByVal srvPoint As ServicePoint, _
   ByVal cert As X509Certificate, ByVal request As WebRequest, ByVal problem As Integer) _
      As Boolean Implements ICertificatePolicy.CheckValidationResult
      Return True
   End Function
End Class

Dim sessionId as String
Dim mgmtBaseUrl as String
Dim eventId as String = "&eventId=1"

'Grab the code and the session lifetime from the generator page
Sub Page_Load(src as Object, e as EventArgs)
   sessionId=Request.QueryString("sessId")
   mgmtBaseUrl=Request.QueryString("mgmtBaseUrl")
   sessTimer=Request.QueryString("sessTimer")


   'Use the override class in myvars.aspx to accept untrusted certificates from the SonicWall
   'This is necessary for the POST to the SonicWall authorizing the LHM session.
   System.Net.ServicePointManager.CertificatePolicy = New acceptAllCerts

   'When the page loads, make the loggedIn span visible
   loggedIn.Visible=True
   loggedOut.Visible=False

   Me.Button1.Attributes.Add("OnClick", "self.close()")

End Sub

'The Logout button
Sub btnSubmit_Click(Sender As Object, E As EventArgs)

   'Let the user know that we are setting up the session, just in case it takes more than a
second
   LHMResult.Text = "Authorizing session. Please wait."

   'The LHM cgi on the SonicWall - this does not change
   Dim loginCgi as String = "externalGuestLogoff.cgi"

   'Assemble the data to post back to the SonicWall to authorize the LHM session
   Dim loginParams as String = "sessId=" & sessionId & eventId

   'Combine mgmtBaseUrl from the original redirect with the login cgi
   Dim postToSNWL as String = mgmtBaseUrl & loginCgi

   'Convert the loginParams to a well behaved byte array
   Dim byteArray As Byte() = Encoding.UTF8.GetBytes(loginParams)
```

```
   Try
      'Make the loggedOut span visible
      loggedIn.Visible=False
      loggedOut.Visible=True

      'Create the webrequest to the SonicWall
      Dim toSNWL as WebRequest = WebRequest.Create(postToSNWL)

      'Calculate the length of the byte array
      toSNWL.ContentLength = byteArray.Length

      'Set the method for the webrequest to POST
      toSNWL.Method = "POST"

      'Set the content type
      toSNWL.ContentType = "application/x-www-form-urlencoded"

      'Open the request stream
      Dim dataStream As Stream = toSNWL.GetRequestStream()

      'Write the byte array to the request stream
      dataStream.Write(byteArray, 0, byteArray.Length)

      'Close the Stream object
      dataStream.Close()

      'Get the response
      Dim snwlReply As WebResponse = toSNWL.GetResponse()

      'Display the status - looking for 200 = OK.
      'Response.Write(CType(snwlReply, HttpWebResponse).StatusCode)

      'Grab the response and stuff it into an xml doc for possible review
      Dim snwlResponse as XmlDocument = New XmlDocument()
      snwlResponse.Load(snwlReply.GetResponseStream())

      'Set the xPath to the SNWL reply, and get the response
      Dim codePath as String = "SonicWallAccessGatewayParam/LogoffReply/ResponseCode"

      'Response.Write(snwlResponse.SelectSingleNode(codePath).InnerXml)

      'Response code 150 - Logout Succeeded
      If snwlResponse.SelectSingleNode(codePath).InnerXml = "150"
         LHMResult.Text = "<br><b><font color=""green"">Your session has been logged
out.<br><br>Thank you for using LHM Guest Services.</font></b>"

      'Response code 251 - Bad HMAC.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "251"
         LHMResult.Text = "<br><b><font color=""red"">Session logout failed:</font></b> The
request failed message authentication. Sorry for the inconvenience. Please close and relaunch
your browser to try again."

      'Response code 253 - Invalid SessionID.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "253"
         LHMResult.Text = "<br><b><font color=""red"">Session logout failed:</font></b> The
request failed to match a known session identity. Sorry for the inconvenience. Please close and
relaunch your browser to try again."

      'Response code 254 - Invalid CGI.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "254"
         LHMResult.Text = "<br><b><font color=""red"">Session logout failed:</font></b> The
request was missing an essential parameter. Sorry for the inconvenience. Please close and
relaunch your browser to try again."

      'Response code 255 - Internal Error.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "255"
```

```
         LHMResult.Text = "<br><b><font color=""red"">Session logout failed:</font></b> The
request failed due to an unspecified error. Sorry for the inconvenience. Please close and
relaunch your browser to try again."

      End If

      'Close the streams
      dataStream.Close()
      snwlReply.Close()

      'If there is some asp.net error trying to talk to the SonicWall, print it in the same
color as the background.
      Catch ex as Exception
         catchError.Text = "<font color=""9CBACE"">" & ex.ToString & "</font>"
         LHMResult.Text = "<br><b><font color=""red"">Session logout failed:</font></b> The
request failed due to an unspecified error. Sorry for the inconvenience. Please close and
relaunch your browser to try again. If the problem persists, please notify an attendant."
      End Try
End Sub

</script>
<STYLE>
body {
   font-size: 10pt;
   font-family: verdana,helvetica,arial,sans-serif;
   color:#000000;
   background-color:#9CBACE;
}

tr.heading {
   font-size: 10pt;
   background-color:#006699;
}

tr.smalltext {
   font-size: 8pt;
}

.button {
   border: 1px solid #000000;
   background-color: #ffffff;
   font-size: 8pt;
}
</STYLE>

<HTML>
<HEAD>
<TITLE>LHM Logout Page</TITLE>

<SCRIPT LANGUAGE="Javascript">

//'Javascript Seconds Countdown Timer
var SecondsToCountDown = <%= sessTimer%>;
var originalTime=" ";

function CountDown()
{
   clockStr="";

   dayStr=Math.floor(SecondsToCountDown/86400)%100000
   if(dayStr>0){
      if(dayStr>1){
         dayStr+=" days ";
      } else dayStr+=" day ";
      clockStr=dayStr;
   }
   hourStr=Math.floor(SecondsToCountDown/3600)%24
```

```
   if(hourStr>0){
      if(hourStr>1){
         hourStr+=" hours ";
      } else hourStr+=" hour ";
      clockStr+=hourStr;
   }
   minuteStr=Math.floor(SecondsToCountDown/60)%60
   if(minuteStr>0){
      if(minuteStr>1){
         minuteStr+=" minutes ";
      } else minuteStr+=" minute ";
      clockStr+=minuteStr;
   }
   secondStr=Math.floor(SecondsToCountDown/1)%60
   if(secondStr>0){
      if(secondStr>1){
         secondStr+=" seconds ";
      } else secondStr+=" second ";
      clockStr+=secondStr;
   }

   if(SecondsToCountDown > 0)
   {
      --SecondsToCountDown;
   }

   if(originalTime.length < 2)
   {
      originalTime = clockStr;
   }

   // Make sure the form is still there before trying to set a value
   if(document.frmValidator){
      document.frmValidator.originalTime.value = originalTime;
      document.frmValidator.countdown.value = clockStr;
   }

   setTimeout("CountDown()", 1000);
   if(SecondsToCountDown == 0)
   {
      document.frmValidator.countdown.value = "Session Expired";
   }
}

//'Disable right-click so that the window doesn't get refreshed since the countdown is
clientside.
document.oncontextmenu = disableRightClick;
function disableRightClick()
{
  return false;
}

//'Disable F5 key, too, on IE at least.
function noF5()
{
   var key_f5 = 116;
   if (key_f5==event.keyCode)
   {
      event.keyCode=0;
      return false;
   }
   return false;
}

document.onkeydown=noF5
document.onmousedown=disableRightClick
```

```
</SCRIPT>

</HEAD>

<BODY onload='CountDown()'>
<span id="loggedIn" runat="server">
<form id="frmValidator" runat="server">
<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr class="heading">
      <td colspan=2 align="center"> </td>
   </tr>
   <tr class="heading">
      <td colspan=2 align="center"><font color="white"><b>SonicWall LHM Logout
Window</b></font></td>
   </tr>
   <tr class="heading">
      <td colspan=2 align="center"> </td>
   </tr>
   <tr class="smalltext"><td><br></td></tr>
   <tr class="smalltext">
      <td>Original Session Time:</td>
      <td><asp:textbox width=250 id="originalTime" runat="server" /></td>
   </tr>
   <tr class="smalltext">
      <td>Remaining Session Time:</td>
      <td><asp:textbox width=250 id="countdown" runat="server" /></td>
   </tr>
   <tr class="smalltext">
      <td colspan=2><br>You may use this window to manually logout your session at any time, or
you may safely close this window if you prefer to let your session timeout
automatically.</font></td>
   </td>
   <tr>
      <td colspan=2><center><asp:button id="btnSubmit" class="button" text="  Logout  "
onClick="btnSubmit_Click" runat="server" /></center></td>
   </tr>
</table>
</form>
</span>

<span id="loggedOut" runat="server">
<form id="logout" runat="server">
<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr class="heading">
      <td colspan=2 align="center"> </td>
   </tr>
   <tr class="heading">
      <td colspan=2 align="center"><font color="white"><b>SonicWall LHM Logout
Window</b></font></td>
   </tr>
   <tr class="heading">
      <td colspan=2 align="center"> </td>
   </tr>
   <tr>
      <td><asp:Label id=LHMResult runat="server" /></td>
   </tr>
   <tr>
      <td><asp:Label id=catchError runat="server" /></td>
   </tr>
   <tr><td><br></td></tr>
   <tr>
      <td><center><asp:button id="Button1" class="button" text="  Close  " runat="server"
/></center></td>
   </tr>
</table>
</form>
</span>
```

```
</BODY>
</HTML>
```

myvars.aspx

```
<script language="VB" runat="server">

'Set the logoutPopup window flag - 0 = no popup, 1 = popup
'The use of the logoutPopup in this script is encouraged because the login event is non-
exclusive.
Dim logoutPopup as String = "1"

'Set the LDAP server IP or Name
Dim myLdapServer as String = "10.50.128.40"

'Set the LDAP domain
Dim myLdapDomain as String = "sv.us.SonicWall.com"

'Set the retrAttr to 0 to use default session and idle timeouts
'Set the retrAttr to 1 to try to retrieve the session and idle timeouts from LDAP attributes.
Dim retrAttr as String ="1"

'Set useCN=1 to use common name (e.g. "joe levy", non-Active Directory LDAP) for attribute
retrieval (retrAttr).
'Set useCN=0 to use saMACcountName (e.g. "jlevy", Active Directory / Windows) for attribute
retrieval.
Dim useCN as String = "0"

'If using retrAttr=1, you must define the ldap attributes from which to retrieve the values
   'Set the ldap attribute from which to retrieve the session timeout value (use is optional)
   Dim sessAttr as String = "pager"

   'Set the ldap attribute from which to retrieve the idle timeout value (use is optional)
   Dim idleAttr as String = "mobile"

'If retrAttr=0, of if no attributes value can be retrieved, use the following timeout values
   'Set the default LHM Session Timeout (for when no attributes is retrieved)
   Dim sessTimer as String = "3600"

   'Set the default LHM Idle Timeout (for when no attributes is retrieved)
   Dim idleTimer as String = "300"

'Set the secret for use with optional HMAC auth, as configured in the Extern Guest Auth config
on the SonicWall
Dim strHmac as String = "password"

'Set the digest method for the HMAC, either MD5 or SHA1
Dim hmacType as String = "MD5"
'Dim hmacType as String = "SHA1"

'Set the logo image to use
Dim logo as String = "SonicWall.gif"

'------------------------End of Configurable Settings-------------------------
</script>
```

Guestbook Script

Authentication Model	The Guest Client provides their name, address, phone, email, URL (optional), and comment (optional) information.
Purpose	Gather market information; write the information to a database for later use.

myvars Variables	`logoutPopup`	Controls the use of the logout popup window. Set to: • **0** to disable the popup window. • **1** to enable the popup window.
	`sessTimer`	The session timer in seconds.
	`idleTimer`	The idle timer in seconds.
	`strHmac`	The shared secret for the optional HMAC function.
	`hmacType`	The digest type to use if HMAC is in use: **MD5** or **SHA1**.
	`logo`	The names of the logo (image) file to use on page headers.
Session Flow	1 The Guest Client enters their personal information and clicks Submit. 2 The entered information is written to a local `.mdb` database file for later use. 3 The LHM post string is assembled with the `sessionID`, the username (as provided in the web-form), the default session lifetime and idle lifetime. 4 The script performs the LHM post to the SonicWall Security Appliance to authorize the session.	
Additional Considerations	Because the script is writing to the database, it is necessary to configure write privileges for the **IUSR_MACHINENAME** and **IWAM_MACHINENAME** (or **ASPNET**) accounts, as described in I want to use the sample scripts SonicWall provided. What do I need to do to use them? on page 824.	

Topics:

- default.aspx on page 856
- logout.aspx on page 863
- myvars.aspx on page 869

default.aspx

```
<%@ Page Language="VB" Debug="true" %>
<%@ Import Namespace="System" %>
<%@ Import Namespace="System.Net" %>
<%@ Import Namespace="System.IO" %>
<%@ Import Namespace="System.Xml" %>
<%@ Import Namespace="System.Text" %>
<%@ Import Namespace="System.Data" %>
<%@ Import Namespace="System.Data.OleDB" %>
<%@ Import Namespace="System.Security" %>
<%@ Import namespace="System.Security.Cryptography.X509Certificates"%>

<!-- #INCLUDE file="myvars.aspx" -->

<script language="VB" runat="server">

'This class allows SSL certs signed by unknown CAs to be accepted.
'This is necessary for the POST to the SonicWall authorizing the LHM session.
Public Class acceptAllCerts
   Implements System.Net.ICertificatePolicy
   Public Function CheckValidationResult(ByVal srvPoint As ServicePoint, _
   ByVal cert As X509Certificate, ByVal request As WebRequest, ByVal problem As Integer) _
      As Boolean Implements ICertificatePolicy.CheckValidationResult
      Return True
   End Function
End Class

'Sample LHM redirect querystring:
'http://127.0.0.1/lhm/guestbook/default.aspx?sessionId=0b712fd83b9f5313db5af1cea6b1004f&ip=10.
50.165.231&mac=00:0e:35:bd:c9:37&ufi=0006b11184300&mgmtBaseUrl=https://10.50.165.193:4043/&cli
entRedirectUrl=https://10.50.165.193:444/&req=http%3A//www.google.com/ig
```

```
Dim ip as String
Dim sessionId as String
Dim mac as String
Dim ufi as String
Dim mgmtBaseUrl as String
Dim clientRedirectUrl as String
Dim req as String
Dim hmac as String
Dim customCode as String

Sub Page_Load(src as Object, e as EventArgs)

   LHMResult.Text=""
   catchError.Text=""

   ip=Request.QueryString("ip")
   sessionId=Request.QueryString("sessionId")
   mac=Request.QueryString("mac")
   ufi=Request.QueryString("ufi")
   mgmtBaseUrl=Request.QueryString("mgmtBaseUrl")
   clientRedirectUrl=Request.QueryString("clientRedirectUrl")
   req=Request.QueryString("req")
   hmac=Request.QueryString("hmac")
   customCode=Request.QueryString("cc")

   'customCode grabs the "cc=" querystring value sent by the SonicWall. This allows you to use
the same
   'page (e.g. this page) for the "Session Expiration" (?cc=2), "Idle Timeout" (?cc=3) and "Max
Sessions" (?cc=4) page.
   If customCode <> "" Then
      Select Case customCode
      Case "2"
         LHMResult.Text="<br><H3><font color=""red"">Your LHM session has expired. You may try
to initiate a new session.</font></H3>"
      Case "3"
         LHMResult.Text="<br><H3><font color=""red"">You have exceeded your idle timeout.
Please log back in.</font></H3>"
      Case "4"
         LHMResult.Text="<br><H3><font color=""red"">The maximum number of sessions has been
reached. Please try again later.</font></H3>"
      End Select
   End If

   'Use the override class in myvars.aspx to accept untrusted certificates from the SonicWall
   'This is necessary for the POST to the SonicWall authorizing the LHM session.
   System.Net.ServicePointManager.CertificatePolicy = New acceptAllCerts

   'Note - the routine below for handling the hmac requires the use of the SonicSSL.dll and
libeay.dll libraries.
   'The DLL must be copied to the IIS server, and the SonicSSL dll must be registered with
"regsvr32 sonicssl.dll"
   If hmac <> "" Then

      'SonicWall URL Encode routine is different from Microsoft - this is the SonicWall method
      req=Replace(req,"%","%25")
      req=Replace(req,":","%3A")
      req=Replace(req," ","%20")
      req=Replace(req,"?","%3F")
      req=Replace(req,"+","%2B")
      req=Replace(req,"&","%26")
      req=Replace(req,"=","%3D")

      Dim strHmacText as String
      Dim objCrypto as Object
      Dim strHmacGenerated
      Dim loginError as String
```

```
      'Initialize the Crypto object
      objCrypto = Server.CreateObject("SonicSSL.Crypto")

      'The text to be encoded
      strHmacText = sessionId & ip & mac & ufi & mgmtBaseUrl & clientRedirectUrl & req

      'Calculate the hash with a key strHmac, the return value is a string converted form the
output sha1 binary.
      'The hash algorithm (MD5 or SHA1) is configured in myvars and in the Extern Guest Auth
config on the SonicWall
      If hmacType = "MD5" Then
         strHmacGenerated = objCrypto.hmac_md5(strHmacText, strHmac)
      Else
         strHmacGenerated = objCrypto.hmac_sha1(strHmacText, strHmac)
      End If

      If strHmacGenerated <> hmac Then
         Dim hmacFail as String
         hmacFail = "<font color=""red"">The HMAC failed validation. Please notify an
attendant.</font><br><br>"
         hmacFail+="<font color=""9CBACE"">Received HMAC: " & hmac & "<br>Calculated HMAC: " &
strHmacGenerated & "<br>"
         hmacFail+="Make sure the digest functions on the SonicWall and LHM server match.<br>"
         hmacFail+="Also make sure the shared secret on the SonicWall and myvars match</font>"
         catchError.Text=hmacFail
      End If

   End If

End Sub

sub OnBtnClearClicked (Sender As Object, e As EventArgs)
   txtName.Text = ""
   txtAddress.Text = ""
   txtCity.Text = ""
   txtState.Text = ""
   txtZip.Text = ""
   txtPhone.Text = ""
   txtEMail.Text = ""
   txtURL.Text = ""
   txtComment.Text = ""
   LHMResult.Text=""
   catchError.Text=""
end sub

Sub btnSubmit_Click(Sender As Object, E As EventArgs)

   Try
      'Try to write the submitted info to the database file
      Dim strConn as string = "Provider=Microsoft.Jet.OLEDB.4.0;Data Source=" &
server.mappath("guestbook.mdb") & ";"

      Dim MySQL as string = "INSERT INTO Guestbook (Name, Address, City, State, Zip, Phone,
EMail, URL, Comment) VALUES ('" & txtName.Text & "','" & txtAddress.Text & "','" & txtCity.Text
& "','" & txtState.Text & "','" & txtZip.Text & "','" & txtPhone.Text & "','" & txtEMail.Text &
"','" & txtURL.Text & "','" & txtComment.Text & "')"
      Dim MyConn as New OleDBConnection (strConn)
      Dim cmd as New OleDBCommand (MySQL, MyConn)
      MyConn.Open ()
      cmd.ExecuteNonQuery ()
      MyConn.Close ()

      Catch ex as Exception
         catchError.Text = "<font color=""9CBACE"">" & ex.ToString & "</font>"
   End Try
```

```
   'Let the user know that we are setting up the session, just in case it takes more than a
second
   LHMResult.Text = "Authorizing session. Please wait."

   'The LHM cgi on the SonicWall - this does not change
   Dim loginCgi as String = "externalGuestLogin.cgi"

   'Assemble the data to post back to the SonicWall to authorize the LHM session
   Dim loginParams as String = "sessId=" & sessionId & "&userName=" &
Server.URLEncode(txtName.Text) & "&sessionLifetime=" & sessTimer & "&idleTimeout=" & idleTimer

   'Combine mgmtBaseUrl from the original redirect with the login cgi
   Dim postToSNWL as String = mgmtBaseUrl & loginCgi

   'Convert the loginParams to a well behaved byte array
   Dim byteArray As Byte() = Encoding.UTF8.GetBytes(loginParams)

   Try
      'Create the webrequest to the SonicWall
      Dim toSNWL as WebRequest = WebRequest.Create(postToSNWL)

      'Calculate the length of the byte array
      toSNWL.ContentLength = byteArray.Length

      'Set the method for the webrequest to POST
      toSNWL.Method = "POST"

      'Set the content type
      toSNWL.ContentType = "application/x-www-form-urlencoded"

      'Open the request stream
      Dim dataStream As Stream = toSNWL.GetRequestStream()

      'Write the byte array to the request stream
      dataStream.Write(byteArray, 0, byteArray.Length)

      'Close the Stream object
      dataStream.Close()

      'Get the response
      Dim snwlReply As WebResponse = toSNWL.GetResponse()

      'Display the status - looking for 200 = OK.
      'Response.Write(CType(snwlReply, HttpWebResponse).StatusCode)

      'Grab the response and stuff it into an xml doc for possible review
      Dim snwlResponse as XmlDocument = New XmlDocument()
      snwlResponse.Load(snwlReply.GetResponseStream())

      'Set the xPath to the SNWL reply, and get the response
      Dim codePath as String = "SonicWallAccessGatewayParam/AuthenticationReply/ResponseCode"

      'Response.Write(snwlResponse.SelectSingleNode(codePath).InnerXml)

      'Response code 50 - Login Succeeded
      If snwlResponse.SelectSingleNode(codePath).InnerXml = "50"

         'Do we want to provide a logout popup window?
         If logoutPopup = "1" Then
            'Popup hack using Javascript for logout window
            Dim sb As New System.Text.StringBuilder()
            sb.Append("<script language='javascript'>")
            sb.Append("window.open('logout.aspx?sessId=")
            sb.Append(Server.URLEncode(CStr(sessionId)))
            sb.Append("&mgmtBaseUrl=")
            sb.Append(Server.URLEncode(CStr(mgmtBaseUrl)))
            sb.Append("&sessTimer=")
```

```
            sb.Append(Server.URLEncode(CStr(sessTimer)))
            sb.Append("','logOut','toolbar=no,")
            sb.Append("addressbar=no,menubar=no,")
            sb.Append("width=400,height=250');")
            sb.Append("<")
            sb.Append("/")
            sb.Append("script>")
            RegisterStartupScript("stp", sb.ToString)
         End If

         LHMResult.Text = "<br><b><font color=""green"">Session authorized:</font></b> You may
now go to the URL you originally requested: <a target=""_blank"" href=""" & req & """>" & req &
"</a>"

      'Response code 51 - Session Limit Exceeded
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "51"
         LHMResult.Text = "<br><b><font color=""red"">Session Limit Reached:</font></b> The
maximum number of guest session has been reached. Sorry for the inconvenience. Please close and
relaunch your browser to try again."

      'Response code 100 - Login Failed.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "100"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> Your
session cannot be created at this time. Sorry for the inconvenience. Please close and relaunch
your browser to try again."

      'Response code 251 - Bad HMAC.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "251"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization failed message authentication. Sorry for the inconvenience. Please
close and relaunch your browser to try again."

      'Response code 253 - Invalid SessionID.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "253"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization failed to match a known session identity. Sorry for the
inconvenience. Please close and relaunch your browser to try again."

      'Response code 254 - Invalid CGI.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "254"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization was missing an essential parameter. Sorry for the inconvenience.
Please close and relaunch your browser to try again."

      'Response code 255 - Internal Error.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "255"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization failed due to an unspecified error. Sorry for the inconvenience.
Please close and relaunch your browser to try again."

      End If

      'Close the streams
      dataStream.Close()
      snwlReply.Close()

      'If there is some asp.net error trying to talk to the SonicWall, print it in the same
color as the background.
      Catch ex as Exception
         catchError.Text = "<font color=""9CBACE"">" & ex.ToString & "</font>"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization failed due to an unspecified error. Sorry for the inconvenience.
Please close and relaunch your browser to try again. If the problem persists, please notify an
attendant."
      End Try
End Sub
</script>
```

```
<STYLE>
body {
  font-size: 10pt;
  font-family: verdana,helvetica,arial,sans-serif;
  color:#000000;
  background-color:#9CBACE;
}

tr.heading {
  background-color:#006699;
}

.button {
   border: 1px solid #000000;
   background-color: #ffffff;
}
</STYLE>

<HTML>
<HEAD>
<TITLE>LHM Guestbook Script</TITLE>
</HEAD>

<BODY>
<form id="frmValidator" runat="server">

<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr class="heading">
      <td colspan=3 align="center"><font color="white"> </td>
   </tr>
   <tr class="heading">
      <td width="50%" valign="center"><font color="white"><b>LHM Guestbook</b></font></td>
      <td><center><img width="216" height="51" src="<%= logo %>"></center></td>
      <td width="50%" align="right" valign="center"><font color="white"><b>Powered by
SonicWall LHM</b> </font></td>
   </tr>
   <tr class="heading">
      <td colspan=3 align="center"><font color="white"> </td>
   </tr>
</table>

<table width="90%" border="0" cellpadding="2" cellspacing="0">
   <tr>
      <td>Welcome <%= ip%> to SonicWall's LHM Guestbook. In exchange for providing us with your
contact information,
      along with your permission to occasionally contact you while you are in the middle of
dinner, we will
      provide you with <b>one complimentary hour of secure Internet access.</b><br>
      </td>
   </tr>
</table>
<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr class="heading">
      <td colspan=3><font color="white"><center><b>Thank you for your
participation.</b></center></td>
   </tr>
</table>
<table width="100%" border="0" cellpadding="0" cellspacing="0">
   <tr>
      <td width="30%"><br>Enter your full name:</td>
      <td width="30%"><asp:TextBox id="txtName" runat="server" /></td>
      <td width="40%"><asp:RequiredFieldValidator id="valTxtName" ControlToValidate="txtName"
ErrorMessage="Please enter your name." runat="server" /></td>
   </tr>
   <tr>
      <td width="30%"><br>Enter your address:</td>
```

```
      <td width="30%"><asp:TextBox id="txtAddress" runat="server" /></td>
      <td width="40%"><asp:RequiredFieldValidator id="valTxtAddress"
ControlToValidate="txtAddress" ErrorMessage="Please enter your address." runat="server"
/></td>
   </tr>
   <tr>
      <td width="30%"><br>Enter your city:</td>
      <td width="30%"><asp:TextBox id="txtCity" runat="server" /></td>
      <td width="40%"><asp:RequiredFieldValidator id="valTxtCity" ControlToValidate="txtCity"
ErrorMessage="Please enter your city." runat="server" /></td>
   </tr>
   <tr>
      <td width="30%"><br>Enter your State:</td>
      <td width="30%"><asp:TextBox id="txtState" runat="server" /></td>
      <td width="40%"><asp:RequiredFieldValidator id="valTxtState"
ControlToValidate="txtState" ErrorMessage="Please enter your State." runat="server" /></td>
   </tr>
   <tr>
      <td width="30%"><br>Enter your zip code:</td>
      <td width="30%"><asp:TextBox id="txtZip" runat="server" /></td>
      <td width="40%"><asp:RequiredFieldValidator id="valTxtZip" ControlToValidate="txtZip"
ErrorMessage="Please enter your zip code." Display="Dynamic" runat="server" />
      <asp:RegularExpressionValidator id=regEx1 runat="server" Display="Dynamic"
ControlToValidate="txtZip" ErrorMessage="Please enter in the format #####"
ValidationExpression="^\d{5}"></asp:RegularExpressionValidator>
      </td>
   </tr>
   <tr>
      <td width="30%"><br>Enter your phone number:</td>
      <td width="30%"><asp:TextBox id="txtPhone" runat="server" /></td>
      <td width="40%"><asp:RequiredFieldValidator id="valTxtPhone"
ControlToValidate="txtPhone" ErrorMessage="Please enter your phone number." Display="Dynamic"
runat="server" />
      <asp:RegularExpressionValidator id=regEx2 runat="server" Display="Dynamic"
ControlToValidate="txtPhone" ErrorMessage="Please enter in the format ###-###-####"
ValidationExpression="((\(\d{3}\)?)|(\d{3}-))?\d{3}-\d{4}"></asp:RegularExpressionValidator>
      </td>
   </tr>
   <tr>
      <td width="30%"><br>Enter your email address:</td>
      <td width="30%"><asp:TextBox id="txtEmail" runat="server" /></td>
      <td width="40%"><asp:RegularExpressionValidator id=regEx3 runat="server"
ControlToValidate="txtEmail" ValidationExpression=".*@.*\..*" ErrorMessage="Please enter a
valid email address." Display="Dynamic" />
      </asp:RegularExpressionValidator>
      <asp:RequiredFieldValidator id="valTxtEnail" runat="server" ControlToValidate="txtEmail"
ErrorMessage="Please enter you email address." Display="Dynamic" />
      </asp:RequiredFieldValidator>
      </td>
   </tr>
   <tr>
      <td width="30%"><br>Enter your web site URL (optional):</td>
      <td width="30%"><asp:TextBox id="txtURL" runat="server" /></td>
   </tr>
   <tr>
      <td valign="top" width="30%"><br>Enter a comment (optional):</td>
      <td width="30%"><asp:TextBox id="txtComment" width="300px" height="100px"
textmode="MultiLine" tabindex="4" runat="server" /></td>
   </tr>
   <tr>
      <td></td>
      <td><br><asp:button id="btnSubmit" class="button" text=" Submit "
onClick="btnSubmit_Click" runat="server" />
        
      <asp:button id="btnClear" class="button" text=" Clear All " CausesValidation="False"
onClick="OnBtnClearClicked" runat="server" />
      </td>
```

```
   </tr>
   <tr>
      <td colspan=2><asp:Label id=LHMResult runat="server" /></td>
   </tr>
   <tr>
      <td colspan=2><asp:Label id=catchError runat="server" /></td>
   </tr>
</table>
</form>
</BODY>
</HTML>
```

logout.aspx

```
<%@ Page Language="VB" Debug="true" %>
<%@ Import Namespace="System" %>
<%@ Import Namespace="System.Net" %>
<%@ Import Namespace="System.IO" %>
<%@ Import Namespace="System.Xml" %>
<%@ Import Namespace="System.Text" %>
<%@ Import Namespace="System.Data" %>
<%@ Import Namespace="System.Data.OleDB" %>
<%@ Import Namespace="System.Security" %>
<%@ Import namespace="System.Security.Cryptography.X509Certificates"%>

<!-- #INCLUDE file="myvars.aspx" -->

<script language="VB" runat="server">

'This class allows SSL certs signed by unknown CAs to be accepted.
'This is necessary for the POST to the SonicWall authorizing the LHM session.
Public Class acceptAllCerts
   Implements System.Net.ICertificatePolicy
   Public Function CheckValidationResult(ByVal srvPoint As ServicePoint, _
   ByVal cert As X509Certificate, ByVal request As WebRequest, ByVal problem As Integer) _
      As Boolean Implements ICertificatePolicy.CheckValidationResult
      Return True
   End Function
End Class

'Sample LHM redirect querystring:
'http://127.0.0.1/lhm/guestbook/default.aspx?sessionId=0b712fd83b9f5313db5af1cea6b1004f&ip=10.
50.165.231&mac=00:0e:35:bd:c9:37&ufi=0006b11184300&mgmtBaseUrl=https://10.50.165.193:4043/&cli
entRedirectUrl=https://10.50.165.193:444/&req=http%3A//www.google.com/ig

Dim ip as String
Dim sessionId as String
Dim mac as String
Dim ufi as String
Dim mgmtBaseUrl as String
Dim clientRedirectUrl as String
Dim req as String
Dim hmac as String
Dim customCode as String

Sub Page_Load(src as Object, e as EventArgs)

   LHMResult.Text=""
   catchError.Text=""

   ip=Request.QueryString("ip")
   sessionId=Request.QueryString("sessionId")
   mac=Request.QueryString("mac")
   ufi=Request.QueryString("ufi")
   mgmtBaseUrl=Request.QueryString("mgmtBaseUrl")
   clientRedirectUrl=Request.QueryString("clientRedirectUrl")
```

```
   req=Request.QueryString("req")
   hmac=Request.QueryString("hmac")
   customCode=Request.QueryString("cc")

   'customCode grabs the "cc=" querystring value sent by the SonicWall. This allows you to use
the same
   'page (e.g. this page) for the "Session Expiration" (?cc=2), "Idle Timeout" (?cc=3) and "Max
Sessions" (?cc=4) page.
   If customCode <> "" Then
      Select Case customCode
      Case "2"
         LHMResult.Text="<br><H3><font color=""red"">Your LHM session has expired. You may try
to initiate a new session.</font></H3>"
      Case "3"
         LHMResult.Text="<br><H3><font color=""red"">You have exceeded your idle timeout.
Please log back in.</font></H3>"
      Case "4"
         LHMResult.Text="<br><H3><font color=""red"">The maximum number of sessions has been
reached. Please try again later.</font></H3>"
      End Select
   End If

   'Use the override class in myvars.aspx to accept untrusted certificates from the SonicWall
   'This is necessary for the POST to the SonicWall authorizing the LHM session.
   System.Net.ServicePointManager.CertificatePolicy = New acceptAllCerts

   'Note - the routine below for handling the hmac requires the use of the SonicSSL.dll and
libeay.dll libraries.
   'The DLL must be copied to the IIS server, and the SonicSSL dll must be registered with
"regsvr32 sonicssl.dll"
   If hmac <> "" Then

      'SonicWall URL Encode routine is different from Microsoft - this is the SonicWall method
      req=Replace(req,"%","%25")
      req=Replace(req,":","%3A")
      req=Replace(req," ","%20")
      req=Replace(req,"?","%3F")
      req=Replace(req,"+","%2B")
      req=Replace(req,"&","%26")
      req=Replace(req,"=","%3D")

      Dim strHmacText as String
      Dim objCrypto as Object
      Dim strHmacGenerated
      Dim loginError as String

      'Initialize the Crypto object
      objCrypto = Server.CreateObject("SonicSSL.Crypto")

      'The text to be encoded
      strHmacText = sessionId & ip & mac & ufi & mgmtBaseUrl & clientRedirectUrl & req

      'Calculate the hash with a key strHmac, the return value is a string converted form the
output sha1 binary.
      'The hash algorithm (MD5 or SHA1) is configured in myvars and in the Extern Guest Auth
config on the SonicWall
      If hmacType = "MD5" Then
         strHmacGenerated = objCrypto.hmac_md5(strHmacText, strHmac)
      Else
         strHmacGenerated = objCrypto.hmac_sha1(strHmacText, strHmac)
      End If

      If strHmacGenerated <> hmac Then
         Dim hmacFail as String
         hmacFail = "<font color=""red"">The HMAC failed validation. Please notify an
attendant.</font><br><br>"
```

```
         hmacFail+="<font color=""9CBACE"">Received HMAC: " & hmac & "<br>Calculated HMAC: " &
strHmacGenerated & "<br>"
         hmacFail+="Make sure the digest functions on the SonicWall and LHM server match.<br>"
         hmacFail+="Also make sure the shared secret on the SonicWall and myvars match</font>"
         catchError.Text=hmacFail
      End If

   End If

End Sub

sub OnBtnClearClicked (Sender As Object, e As EventArgs)
   txtName.Text = ""
   txtAddress.Text = ""
   txtCity.Text = ""
   txtState.Text = ""
   txtZip.Text = ""
   txtPhone.Text = ""
   txtEMail.Text = ""
   txtURL.Text = ""
   txtComment.Text = ""
   LHMResult.Text=""
   catchError.Text=""
end sub

Sub btnSubmit_Click(Sender As Object, E As EventArgs)

   Try
      'Try to write the submitted info to the database file
      Dim strConn as string = "Provider=Microsoft.Jet.OLEDB.4.0;Data Source=" &
server.mappath("guestbook.mdb") & ";"

      Dim MySQL as string = "INSERT INTO Guestbook (Name, Address, City, State, Zip, Phone,
EMail, URL, Comment) VALUES ('" & txtName.Text & "','" & txtAddress.Text & "','" & txtCity.Text
& "','" & txtState.Text & "','" & txtZip.Text & "','" & txtPhone.Text & "','" & txtEMail.Text &
"','" & txtURL.Text & "','" & txtComment.Text & "')"
      Dim MyConn as New OleDBConnection (strConn)
      Dim cmd as New OleDBCommand (MySQL, MyConn)
      MyConn.Open ()
      cmd.ExecuteNonQuery ()
      MyConn.Close ()

      Catch ex as Exception
         catchError.Text = "<font color=""9CBACE"">" & ex.ToString & "</font>"
   End Try

   'Let the user know that we are setting up the session, just in case it takes more than a
second
   LHMResult.Text = "Authorizing session. Please wait."

   'The LHM cgi on the SonicWall - this does not change
   Dim loginCgi as String = "externalGuestLogin.cgi"

   'Assemble the data to post back to the SonicWall to authorize the LHM session
   Dim loginParams as String = "sessId=" & sessionId & "&userName=" &
Server.URLEncode(txtName.Text) & "&sessionLifetime=" & sessTimer & "&idleTimeout=" & idleTimer

   'Combine mgmtBaseUrl from the original redirect with the login cgi
   Dim postToSNWL as String = mgmtBaseUrl & loginCgi

   'Convert the loginParams to a well behaved byte array
   Dim byteArray As Byte() = Encoding.UTF8.GetBytes(loginParams)

   Try
      'Create the webrequest to the SonicWall
      Dim toSNWL as WebRequest = WebRequest.Create(postToSNWL)
```

```
      'Calculate the length of the byte array
      toSNWL.ContentLength = byteArray.Length

      'Set the method for the webrequest to POST
      toSNWL.Method = "POST"

      'Set the content type
      toSNWL.ContentType = "application/x-www-form-urlencoded"

      'Open the request stream
      Dim dataStream As Stream = toSNWL.GetRequestStream()

      'Write the byte array to the request stream
      dataStream.Write(byteArray, 0, byteArray.Length)

      'Close the Stream object
      dataStream.Close()

      'Get the response
      Dim snwlReply As WebResponse = toSNWL.GetResponse()

      'Display the status - looking for 200 = OK.
      'Response.Write(CType(snwlReply, HttpWebResponse).StatusCode)

      'Grab the response and stuff it into an xml doc for possible review
      Dim snwlResponse as XmlDocument = New XmlDocument()
      snwlResponse.Load(snwlReply.GetResponseStream())

      'Set the xPath to the SNWL reply, and get the response
      Dim codePath as String = "SonicWallAccessGatewayParam/AuthenticationReply/ResponseCode"

      'Response.Write(snwlResponse.SelectSingleNode(codePath).InnerXml)

      'Response code 50 - Login Succeeded
      If snwlResponse.SelectSingleNode(codePath).InnerXml = "50"

         'Do we want to provide a logout popup window?
         If logoutPopup = "1" Then
            'Popup hack using Javascript for logout window
            Dim sb As New System.Text.StringBuilder()
            sb.Append("<script language='javascript'>")
            sb.Append("window.open('logout.aspx?sessId=")
            sb.Append(Server.URLEncode(CStr(sessionId)))
            sb.Append("&mgmtBaseUrl=")
            sb.Append(Server.URLEncode(CStr(mgmtBaseUrl)))
            sb.Append("&sessTimer=")
            sb.Append(Server.URLEncode(CStr(sessTimer)))
            sb.Append("','logOut','toolbar=no,")
            sb.Append("addressbar=no,menubar=no,")
            sb.Append("width=400,height=250');")
            sb.Append("<")
            sb.Append("/")
            sb.Append("script>")
            RegisterStartupScript("stp", sb.ToString)
         End If

         LHMResult.Text = "<br><b><font color=""green"">Session authorized:</font></b> You may
now go to the URL you originally requested: <a target=""_blank"" href=""" & req & """>" & req &
"</a>"

      'Response code 51 - Session Limit Exceeded
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "51"
         LHMResult.Text = "<br><b><font color=""red"">Session Limit Reached:</font></b> The
maximum number of guest session has been reached. Sorry for the inconvenience. Please close and
relaunch your browser to try again."

      'Response code 100 - Login Failed.
```

```
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "100"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> Your
session cannot be created at this time. Sorry for the inconvenience. Please close and relaunch
your browser to try again."

      'Response code 251 - Bad HMAC.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "251"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization failed message authentication. Sorry for the inconvenience. Please
close and relaunch your browser to try again."

      'Response code 253 - Invalid SessionID.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "253"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization failed to match a known session identity. Sorry for the
inconvenience. Please close and relaunch your browser to try again."

      'Response code 254 - Invalid CGI.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "254"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization was missing an essential parameter. Sorry for the inconvenience.
Please close and relaunch your browser to try again."

      'Response code 255 - Internal Error.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "255"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization failed due to an unspecified error. Sorry for the inconvenience.
Please close and relaunch your browser to try again."

      End If

      'Close the streams
      dataStream.Close()
      snwlReply.Close()

      'If there is some asp.net error trying to talk to the SonicWall, print it in the same
color as the background.
      Catch ex as Exception
         catchError.Text = "<font color=""9CBACE"">" & ex.ToString & "</font>"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization failed due to an unspecified error. Sorry for the inconvenience.
Please close and relaunch your browser to try again. If the problem persists, please notify an
attendant."
      End Try
End Sub
</script>

<STYLE>
body {
  font-size: 10pt;
  font-family: verdana,helvetica,arial,sans-serif;
  color:#000000;
  background-color:#9CBACE;
}

tr.heading {
  background-color:#006699;
}

.button {
   border: 1px solid #000000;
   background-color: #ffffff;
}
</STYLE>

<HTML>
<HEAD>
```

```
<TITLE>LHM Guestbook Script</TITLE>
</HEAD>

<BODY>
<form id="frmValidator" runat="server">

<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr class="heading">
      <td colspan=3 align="center"><font color="white"> </td>
   </tr>
   <tr class="heading">
      <td width="50%" valign="center"><font color="white"><b>LHM Guestbook</b></font></td>
      <td><center><img width="216" height="51" src="<%= logo %>"></center></td>
      <td width="50%" align="right" valign="center"><font color="white"><b>Powered by
SonicWall LHM</b> </font></td>
   </tr>
   <tr class="heading">
      <td colspan=3 align="center"><font color="white"> </td>
   </tr>
</table>

<table width="90%" border="0" cellpadding="2" cellspacing="0">
   <tr>
      <td>Welcome <%= ip%> to SonicWall's LHM Guestbook. In exchange for providing us with your
contact information,
      along with your permission to occasionally contact you while you are in the middle of
dinner, we will
      provide you with <b>one complimentary hour of secure Internet access.</b><br>
      </td>
   </tr>
</table>
<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr class="heading">
      <td colspan=3><font color="white"><center><b>Thank you for your
participation.</b></center></td>
   </tr>
</table>
<table width="100%" border="0" cellpadding="0" cellspacing="0">
   <tr>
      <td width="30%"><br>Enter your full name:</td>
      <td width="30%"><asp:TextBox id="txtName" runat="server" /></td>
      <td width="40%"><asp:RequiredFieldValidator id="valTxtName" ControlToValidate="txtName"
ErrorMessage="Please enter your name." runat="server" /></td>
   </tr>
   <tr>
      <td width="30%"><br>Enter your address:</td>
      <td width="30%"><asp:TextBox id="txtAddress" runat="server" /></td>
      <td width="40%"><asp:RequiredFieldValidator id="valTxtAddress"
ControlToValidate="txtAddress" ErrorMessage="Please enter your address." runat="server"
/></td>
   </tr>
   <tr>
      <td width="30%"><br>Enter your city:</td>
      <td width="30%"><asp:TextBox id="txtCity" runat="server" /></td>
      <td width="40%"><asp:RequiredFieldValidator id="valTxtCity" ControlToValidate="txtCity"
ErrorMessage="Please enter your city." runat="server" /></td>
   </tr>
   <tr>
      <td width="30%"><br>Enter your State:</td>
      <td width="30%"><asp:TextBox id="txtState" runat="server" /></td>
      <td width="40%"><asp:RequiredFieldValidator id="valTxtState"
ControlToValidate="txtState" ErrorMessage="Please enter your State." runat="server" /></td>
   </tr>
   <tr>
      <td width="30%"><br>Enter your zip code:</td>
      <td width="30%"><asp:TextBox id="txtZip" runat="server" /></td>
```

```
      <td width="40%"><asp:RequiredFieldValidator id="valTxtZip" ControlToValidate="txtZip"
ErrorMessage="Please enter your zip code." Display="Dynamic" runat="server" />
      <asp:RegularExpressionValidator id=regEx1 runat="server" Display="Dynamic"
ControlToValidate="txtZip" ErrorMessage="Please enter in the format #####"
ValidationExpression="^\d{5}"></asp:RegularExpressionValidator>
      </td>
   </tr>
   <tr>
      <td width="30%"><br>Enter your phone number:</td>
      <td width="30%"><asp:TextBox id="txtPhone" runat="server" /></td>
      <td width="40%"><asp:RequiredFieldValidator id="valTxtPhone"
ControlToValidate="txtPhone" ErrorMessage="Please enter your phone number." Display="Dynamic"
runat="server" />
      <asp:RegularExpressionValidator id=regEx2 runat="server" Display="Dynamic"
ControlToValidate="txtPhone" ErrorMessage="Please enter in the format ###-###-####"
ValidationExpression="((\(\d{3}\)?)|(\d{3}-))?\d{3}-\d{4}"></asp:RegularExpressionValidator>
      </td>
   </tr>
   <tr>
      <td width="30%"><br>Enter your email address:</td>
      <td width="30%"><asp:TextBox id="txtEmail" runat="server" /></td>
      <td width="40%"><asp:RegularExpressionValidator id=regEx3 runat="server"
ControlToValidate="txtEmail" ValidationExpression=".*@.*\..*" ErrorMessage="Please enter a
valid email address." Display="Dynamic" />
      </asp:RegularExpressionValidator>
      <asp:RequiredFieldValidator id="valTxtEnail" runat="server" ControlToValidate="txtEmail"
ErrorMessage="Please enter you email address." Display="Dynamic" />
      </asp:RequiredFieldValidator>
      </td>
   </tr>
   <tr>
      <td width="30%"><br>Enter your web site URL (optional):</td>
      <td width="30%"><asp:TextBox id="txtURL" runat="server" /></td>
   </tr>
   <tr>
      <td valign="top" width="30%"><br>Enter a comment (optional):</td>
      <td width="30%"><asp:TextBox id="txtComment" width="300px" height="100px"
textmode="MultiLine" tabindex="4" runat="server" /></td>
   </tr>
   <tr>
      <td></td>
      <td><br><asp:button id="btnSubmit" class="button" text="  Submit  "
onClick="btnSubmit_Click" runat="server" />
        
      <asp:button id="btnClear" class="button" text=" Clear All " CausesValidation="False"
onClick="OnBtnClearClicked" runat="server" />
      </td>
   </tr>
   <tr>
      <td colspan=2><asp:Label id=LHMResult runat="server" /></td>
   </tr>
   <tr>
      <td colspan=2><asp:Label id=catchError runat="server" /></td>
   </tr>
</table>
</form>
</BODY>
</HTML>
```

myvars.aspx

```
<script language="VB" runat="server">

'Set the logoutPopup window flag - 0 = no popup, 1 = popup
'The use of the logoutPopup in this script is discouraged because although the login event
'is non-exclusive, the login event produces data where redundancy is undesirable.
Dim logoutPopup as String = "0"
```

```
'Set the LHM Session Timeout
Dim sessTimer as String = "3600"

'Set the LHM Idle Timeout
Dim idleTimer as String = "300"

'Set the secret for use with optional HMAC auth, as configured in the Extern Guest Auth config
on the SonicWall
Dim strHmac as String = "password"

'Set the digest method for the HMAC, either MD5 or SHA1
Dim hmacType as String = "MD5"
'Dim hmacType as String = "SHA1"

'Set the logo image to use
Dim logo as String = "SonicWall.gif"

'-----------------------End of Configurable Settings-------------------------

</script>
```

LHMQuiz Script

Authentication Model	The Guest Client takes a quiz. A passing score serves as the authentication credentials	
Purpose	It is common for network access to be provided in a classroom environment. By using a passing score on a test of the material being taught as the method for authentication, an instructor can ensure that the course material has been mastered before the irresistible temptation of the Internet diverts attention. The script also emails the completed passing test to the test-taker, and mails failing tests to the proctor/instructor.	
myvars Variables	`logoutPopup`	Controls the use of the logout popup window. Set to: • **0** to disable the popup window. • **1** to enable the popup window.
	`passingScore`	The score (an integer representing a percentage) required to pass the quiz.
	`quizFile`	The filename for the XML source for the quiz (such as `quiz.xml`, `shortquiz.xml`).
	`quizName`	The name of the quiz, used throughout the script.
	`quizFrom`	The `From: email` address used when emailing the quiz.
	`quizTo`	The `To: email` address where failing quizzes are to be sent (such as the test proctor or instructor).
	`imagePath`	The email includes an attachment for the correct and incorrect answers. This sets the path for those image files. This is generally set to the same path of the script files themselves.
	`smtpServer`	The IP address or resolvable FQDN of the SMTP server to be used for quiz result delivery. This can be set to `127.0.0.1` if the local IIS SMTP server instances is to be used.
	`sessTimer`	The session timer in seconds.
	`idleTimer`	The idle timer in seconds.
	`strHmac`	The shared secret for the optional HMAC function.
	`hmacType`	The digest type to use if HMAC is in use: **MD5** or **SHA1**.

`logo`	The names of the logo (image) file to use on page headers.

Session Flow

1. The Guest Client is prompted to enter their full name and email address. A correct/valid email address is required for delivery of the completed passing quiz.
2. After entering name and email, the Guest Client is redirected to the `quiz.aspx` page. This is where the multiple choice test is administered.
3. The test questions themselves are contained in the `quiz.xml` file, defined by the `quiz.xsd` (XML Schema Definition) file. The `quiz.xml` file can and should be edited to customize the quiz, but the `quiz.xsd` document should not be edited unless absolutely necessary.

   Two versions of the quiz are included: `quiz.xml` (containing 10 questions) and `shortquiz.xml` (containing 2 questions, for testing that the script works). The quiz supports any number of questions, and each question supports any number of answers, one of which must be marked the correct answer, with `correct=yes`. It should be fairly straightforward to modify the provided `quiz.xml` file as needed.
4. At the end of the quiz, the results are shown. If it is a:
   - Failing score, the test results are emailed to the instructor (email address defined in `myvars`), and the Guest Client is prompted to take the test again. The LHM session is not authorized.
   - Passing score, the test results are emailed to the test-taker, and the LHM session is authorized.

   The emailed test is sent in an HTML format, and includes the `checkmark.gif` and `block.gif` (right and wrong) graphics as an attachment so that they can be displayed in the email.
5. If the test was passed, the LHM post string is assembled with the `sessionID`, the username (as provided in the web-form), the default session lifetime and idle lifetime.
6. The script performs the LHM post to the SonicWall Security Appliance to authorize the session.

Additional Considerations

Access to an SMTP server is required to deliver the test results. Because the script is relaying the mail through the server, the SMTP server needs to be configured to allow relaying from the LHM server. This is best accomplished by configuring the SMTP server to allow relaying from the IP address of the LHM server.

Most IIS installations include a local SMTP server, so it is convenient to use this local SMTP server for mail delivery by configuring the `smtpServer` variable in `myvars` as `127.0.0.1`.

Even when using the local SMTP server for mail delivery, it is necessary to allow relaying. In most configurations, this is performed by:

1. Going into the IIS MMC configurator.
2. Right clicking on **Default SMTP Virtual Server**.
3. Selecting **Properties**.
4. Selecting **Access**.
5. Clicking **Relay** .
6. Adding `127.0.0.1` to the access granted list.

When using a non-local SMTP server, that SMTP server should be configured to allow the LHM server to relay by its actual IP address.

Topics:

- default.aspx on page 872
- logout.aspx on page 875
- myvars.aspx on page 880
- quiz.aspx on page 881

default.aspx

```
<%@ Page Language="VB" Debug="true" %>
<%@ Import Namespace="System" %>

<!-- #INCLUDE file="myvars.aspx" -->

<script runat="server">

'Sample LHM redirect querystring:
'http://10.50.165.231/xmlquiz/default.aspx?sessionId=0b712fd83b9f5313db5af1cea6b1004f&ip=10.50
.165.231&mac=00:0e:35:bd:c9:37&ufi=0006b11184300&mgmtBaseUrl=https://10.50.165.193:4043/&clien
tRedirectUrl=https://10.50.165.193:444/&req=http%3A//www.google.com/ig

Dim ip as String
Dim sessionId as String
Dim mac as String
Dim ufi as String
Dim mgmtBaseUrl as String
Dim clientRedirectUrl as String
Dim req as String
Dim hmac as String
Dim emailAddr as String
Dim userName as String
Dim customCode as String

Sub Page_Load(src as Object, e as EventArgs)
   ip=Request.QueryString("ip")
   sessionId=Request.QueryString("sessionId")
   mac=Request.QueryString("mac")
   ufi=Request.QueryString("ufi")
   mgmtBaseUrl=Request.QueryString("mgmtBaseUrl")
   clientRedirectUrl=Request.QueryString("clientRedirectUrl")
   req=Request.QueryString("req")
   hmac=Request.QueryString("hmac")
   customCode=Request.QueryString("cc")

   'customCode grabs the "cc=" querystring value sent by the SonicWall. This allows you to use
the same
   'page (e.g. this page) for the "Session Expiration" (?cc=2), "Idle Timeout" (?cc=3) and "Max
Sessions" (?cc=4) page.
   If customCode <> "" Then
      Select Case customCode
      Case "2"
         LHMResult.Text="<br><H3><font color=""red"">Your LHM session has expired. You may try
to initiate a new session.</font></H3>"
      Case "3"
         LHMResult.Text="<br><H3><font color=""red"">You have exceeded your idle timeout.
Please log back in.</font></H3>"
      Case "4"
         LHMResult.Text="<br><H3><font color=""red"">The maximum number of sessions has been
reached. Please try again later.</font></H3>"
      End Select
   End If

   'Note - the routine below for handling the hmac requires the use of the SonicSSL.dll and
libeay.dll libraries.
   'The DLL must be copied to the IIS server, and the SonicSSL dll must be registered with
"regsvr32 sonicssl.dll"
   If hmac <> "" Then

      'SonicWall URL Encode routine is different from Microsoft - this is the SonicWall method
      req=Replace(req,"%","%25")
      req=Replace(req,":","%3A")
      req=Replace(req," ","%20")
      req=Replace(req,"?","%3F")
      req=Replace(req,"+","%2B")
```

```
      req=Replace(req,"&","%26")
      req=Replace(req,"=","%3D")

      Dim strHmacText as String
      Dim objCrypto as Object
      Dim strHmacGenerated
      Dim loginError as String

      'Initialize the Crypto object
      objCrypto = Server.CreateObject("SonicSSL.Crypto")

      'The text to be encoded
      strHmacText = sessionId & ip & mac & ufi & mgmtBaseUrl & clientRedirectUrl & req

      'Calculate the hash with a key strHmac, the return value is a string converted form the
output sha1 binary.
      'The hash algorithm (MD5 or SHA1) is configured in myvars and in the Extern Guest Auth
config on the SonicWall
      If hmacType = "MD5" Then
         strHmacGenerated = objCrypto.hmac_md5(strHmacText, strHmac)
      Else
         strHmacGenerated = objCrypto.hmac_sha1(strHmacText, strHmac)
      End If

      If strHmacGenerated <> hmac Then
         Dim hmacFail as String
         hmacFail = "<font color=""red"">The HMAC failed validation. Please notify an
attendant.</font><br><br>"
         hmacFail+="<font color=""9CBACE"">Received HMAC: " & hmac & "<br>Calculated HMAC: " &
strHmacGenerated & "<br>"
         hmacFail+="Make sure the digest functions on the SonicWall and LHM server match.<br>"
         hmacFail+="Also make sure the shared secret on the SonicWall and myvars match</font>"
         catchError.Text=hmacFail
      End If

   End If

End Sub

'When the submit button is clicked, pass the variables we need and load the quiz
Sub btnSubmit_Click(Sender As Object, E As EventArgs)

   Context.Items.Add("req", req)
   Context.Items.Add("sessionId", sessionId)
   Context.Items.Add("emailAddr",clientEmail.Text)
   Context.Items.Add("userName",clientName.Text)
   Context.Items.Add("mgmtBaseUrl",mgmtBaseUrl)
   Server.Transfer("quiz.aspx",true)

End Sub

</script>

<STYLE>
body {
  font-size: 10pt;
  font-family: verdana,helvetica,arial,sans-serif;
  color:#000000;
  background-color:#9CBACE;
}

tr.heading {
  background-color:#006699;
}

.button {
   border: 1px solid #000000;
```

```
   background-color: #ffffff;
}
</STYLE>

<HTML>
<HEAD>
<TITLE>LHM Quiz Script</TITLE>
</HEAD>

<BODY>
<form id="frmValidator" runat="server">

<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr class="heading">
      <td colspan=3 align="center"><font color="white"> </td>
   </tr>
   <tr class="heading">
      <td width="50%" valign="center"><font color="white"><b>LHM Quiz
Authorization</b></font></td>
      <td><center><img width="216" height="51" src="SonicWall.gif"></center></td>
      <td width="50%" align="right" valign="center"><font color="white"><b>Powered by
SonicWall LHM</b> </font></td>
   </tr>
   <tr class="heading">
      <td colspan=3 align="center"><font color="white"> </td>
   </tr>
</table>
<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr>
      <td width="30%"><br>Enter your full name:</td>
      <td width="20%"><asp:TextBox id="clientName" runat="server" /></td>
      <td ><asp:RequiredFieldValidator id="valTxtName" ControlToValidate="clientName"
ErrorMessage="Please enter your name." Display="Dynamic" runat="server" /></td>
   </tr>
   <tr>
      <td width="30%"><br>Enter your real email address:</td>
      <td width="20%"><asp:TextBox id="clientEmail" runat="server" /></td>
      <td ><asp:RegularExpressionValidator id="fromEmail" runat="server"
ControlToValidate="clientEmail" ValidationExpression=".*@.*\..*" ErrorMessage="Please enter a
valid email address." Display="Dynamic" />
      </asp:RegularExpressionValidator>
      <asp:RequiredFieldValidator id="fromRequired" runat="server"
ControlToValidate="clientEmail" ErrorMessage="Please enter your email address."
Display="Dynamic" />
      </asp:RequiredFieldValidator>
      </td>
   </tr>
   <tr>
      <td></td>
      <td><asp:button id="btnSubmit" class="button" text="  Submit  "
onClick="btnSubmit_Click" runat="server" /><br></td>
   </tr>

   <tr class="heading">
      <td colspan=3 align="left"><font color="white"><b>Welcome Quiztaker <%=
ip%></b></font></td>
   </tr>
</table>
<table width="70%" border="0" cellpadding="2" cellspacing="0">
   <tr>
      <td>
      <br>You have been redirected here by Lightweight Hotspot Messaging.
      This environment has been setup to demonstrate the flexibility of LHM, including
      support for both wired and wireless clients, and also the ability for LHM to use
      more than just username and password authentication for providing access.<br><br>
      The page that you are about to continue on to is a <%= quizName %> written in ASP.net.
```

```
      A passing score of <%= passingScore%>% will serve as the authentication for LHM, and will
grant
      you network access. You must pass the test to continue, and will be prompted to retake
      the entire quiz if you you do not pass. <br><br>
      When you are done, the completed test will be emailed to you at the address you
      specify above.<br><br>
      So it's not just a good way to prove your understanding of some
      key SonicOS concepts, but also a practical example of the versatility of LHM.
      </td>
   </tr>
   <tr>
      <td><asp:Label id=LHMResult runat="server" /></td>
   </tr>
   <tr>
      <td colspan=2><asp:Label id=catchError runat="server" /></td>
   </tr>
</table>
</form>
</BODY>
</HTML>
```

logout.aspx

```
<%@ Page Language="VB" Debug="true" %>
<%@ Import Namespace="System" %>
<%@ Import Namespace="System.Net" %>
<%@ Import Namespace="System.IO" %>
<%@ Import Namespace="System.Xml" %>
<%@ Import Namespace="System.Text" %>
<%@ Import Namespace="System.Security" %>
<%@ Import namespace="System.Security.Cryptography.X509Certificates"%>

<!-- #INCLUDE file="myvars.aspx" -->

<script language="VB" runat="server">

'This class allows SSL certs signed by unknown CAs to be accepted.
'This is necessary for the POST to the SonicWall authorizing the LHM session.
Public Class acceptAllCerts
   Implements System.Net.ICertificatePolicy
   Public Function CheckValidationResult(ByVal srvPoint As ServicePoint, _
   ByVal cert As X509Certificate, ByVal request As WebRequest, ByVal problem As Integer) _
      As Boolean Implements ICertificatePolicy.CheckValidationResult
      Return True
   End Function
End Class

Dim sessionId as String
Dim mgmtBaseUrl as String
Dim eventId as String = "&eventId=1"

'Grab the code and the session lifetime from the generator page
Sub Page_Load(src as Object, e as EventArgs)
   sessionId=Request.QueryString("sessId")
   mgmtBaseUrl=Request.QueryString("mgmtBaseUrl")
   sessTimer=Request.QueryString("sessTimer")


   'Use the override class in myvars.aspx to accept untrusted certificates from the SonicWall
   'This is necessary for the POST to the SonicWall authorizing the LHM session.
   System.Net.ServicePointManager.CertificatePolicy = New acceptAllCerts

   'When the page loads, make the loggedIn span visible
   loggedIn.Visible=True
   loggedOut.Visible=False
```

```
   Me.Button1.Attributes.Add("OnClick", "self.close()")

End Sub

'The Logout button
Sub btnSubmit_Click(Sender As Object, E As EventArgs)

   'Let the user know that we are setting up the session, just in case it takes more than a
second
   LHMResult.Text = "Authorizing session. Please wait."

   'The LHM cgi on the SonicWall - this does not change
   Dim loginCgi as String = "externalGuestLogoff.cgi"

   'Assemble the data to post back to the SonicWall to authorize the LHM session
   Dim loginParams as String = "sessId=" & sessionId & eventId

   'Combine mgmtBaseUrl from the original redirect with the login cgi
   Dim postToSNWL as String = mgmtBaseUrl & loginCgi

   'Convert the loginParams to a well behaved byte array
   Dim byteArray As Byte() = Encoding.UTF8.GetBytes(loginParams)

   Try
      'Make the loggedOut span visible
      loggedIn.Visible=False
      loggedOut.Visible=True

      'Create the webrequest to the SonicWall
      Dim toSNWL as WebRequest = WebRequest.Create(postToSNWL)

      'Calculate the length of the byte array
      toSNWL.ContentLength = byteArray.Length

      'Set the method for the webrequest to POST
      toSNWL.Method = "POST"

      'Set the content type
      toSNWL.ContentType = "application/x-www-form-urlencoded"

      'Open the request stream
      Dim dataStream As Stream = toSNWL.GetRequestStream()

      'Write the byte array to the request stream
      dataStream.Write(byteArray, 0, byteArray.Length)

      'Close the Stream object
      dataStream.Close()

      'Get the response
      Dim snwlReply As WebResponse = toSNWL.GetResponse()

      'Display the status - looking for 200 = OK.
      'Response.Write(CType(snwlReply, HttpWebResponse).StatusCode)

      'Grab the response and stuff it into an xml doc for possible review
      Dim snwlResponse as XmlDocument = New XmlDocument()
      snwlResponse.Load(snwlReply.GetResponseStream())

      'Set the xPath to the SNWL reply, and get the response
      Dim codePath as String = "SonicWallAccessGatewayParam/LogoffReply/ResponseCode"

      'Response.Write(snwlResponse.SelectSingleNode(codePath).InnerXml)

      'Response code 150 - Logout Succeeded
      If snwlResponse.SelectSingleNode(codePath).InnerXml = "150"
```

```
         LHMResult.Text = "<br><b><font color=""green"">Your session has been logged
out.<br><br>Thank you for using LHM Guest Services.</font></b>"

      'Response code 251 - Bad HMAC.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "251"
         LHMResult.Text = "<br><b><font color=""red"">Session logout failed:</font></b> The
request failed message authentication. Sorry for the inconvenience. Please close and relaunch
your browser to try again."

      'Response code 253 - Invalid SessionID.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "253"
         LHMResult.Text = "<br><b><font color=""red"">Session logout failed:</font></b> The
request failed to match a known session identity. Sorry for the inconvenience. Please close and
relaunch your browser to try again."

      'Response code 254 - Invalid CGI.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "254"
         LHMResult.Text = "<br><b><font color=""red"">Session logout failed:</font></b> The
request was missing an essential parameter. Sorry for the inconvenience. Please close and
relaunch your browser to try again."

      'Response code 255 - Internal Error.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "255"
         LHMResult.Text = "<br><b><font color=""red"">Session logout failed:</font></b> The
request failed due to an unspecified error. Sorry for the inconvenience. Please close and
relaunch your browser to try again."

      End If

      'Close the streams
      dataStream.Close()
      snwlReply.Close()

      'If there is some asp.net error trying to talk to the SonicWall, print it in the same
color as the background.
      Catch ex as Exception
         catchError.Text = "<font color=""9CBACE"">" & ex.ToString & "</font>"
         LHMResult.Text = "<br><b><font color=""red"">Session logout failed:</font></b> The
request failed due to an unspecified error. Sorry for the inconvenience. Please close and
relaunch your browser to try again. If the problem persists, please notify an attendant."
      End Try
End Sub

</script>
<STYLE>
body {
   font-size: 10pt;
   font-family: verdana,helvetica,arial,sans-serif;
   color:#000000;
   background-color:#9CBACE;
}

tr.heading {
   font-size: 10pt;
   background-color:#006699;
}

tr.smalltext {
   font-size: 8pt;
}

.button {
   border: 1px solid #000000;
   background-color: #ffffff;
   font-size: 8pt;
}
</STYLE>
```

```
<HTML>
<HEAD>
<TITLE>LHM Logout Page</TITLE>

<SCRIPT LANGUAGE="Javascript">

//'Javascript Seconds Countdown Timer
var SecondsToCountDown = <%= sessTimer%>;
var originalTime=" ";

function CountDown()
{
   clockStr="";

   dayStr=Math.floor(SecondsToCountDown/86400)%100000
   if(dayStr>0){
      if(dayStr>1){
         dayStr+=" days ";
      } else dayStr+=" day ";
      clockStr=dayStr;
   }
   hourStr=Math.floor(SecondsToCountDown/3600)%24
   if(hourStr>0){
      if(hourStr>1){
         hourStr+=" hours ";
      } else hourStr+=" hour ";
      clockStr+=hourStr;
   }
   minuteStr=Math.floor(SecondsToCountDown/60)%60
   if(minuteStr>0){
      if(minuteStr>1){
         minuteStr+=" minutes ";
      } else minuteStr+=" minute ";
      clockStr+=minuteStr;
   }
   secondStr=Math.floor(SecondsToCountDown/1)%60
   if(secondStr>0){
      if(secondStr>1){
         secondStr+=" seconds ";
      } else secondStr+=" second ";
      clockStr+=secondStr;
   }

   if(SecondsToCountDown > 0)
   {
      --SecondsToCountDown;
   }

   if(originalTime.length < 2)
   {
      originalTime = clockStr;
   }

   // Make sure the form is still there before trying to set a value
   if(document.frmValidator){
      document.frmValidator.originalTime.value = originalTime;
      document.frmValidator.countdown.value = clockStr;
   }

   setTimeout("CountDown()", 1000);
   if(SecondsToCountDown == 0)
   {
      document.frmValidator.countdown.value = "Session Expired";
   }
}
```

```
//'Disable right-click so that the window doesn't get refreshed since the countdown is
clientside.
document.oncontextmenu = disableRightClick;
function disableRightClick()
{
  return false;
}

//'Disable F5 key, too, on IE at least.
function noF5()
{
   var key_f5 = 116;
   if (key_f5==event.keyCode)
   {
      event.keyCode=0;
      return false;
   }
   return false;
}

document.onkeydown=noF5
document.onmousedown=disableRightClick

</SCRIPT>

</HEAD>

<BODY onload='CountDown()'>
<span id="loggedIn" runat="server">
<form id="frmValidator" runat="server">
<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr class="heading">
      <td colspan=2 align="center"> </td>
   </tr>
   <tr class="heading">
      <td colspan=2 align="center"><font color="white"><b>SonicWall LHM Logout
Window</b></font></td>
   </tr>
   <tr class="heading">
      <td colspan=2 align="center"> </td>
   </tr>
   <tr class="smalltext"><td><br></td></tr>
   <tr class="smalltext">
      <td>Original Session Time:</td>
      <td><asp:textbox width=250 id="originalTime" runat="server" /></td>
   </tr>
   <tr class="smalltext">
      <td>Remaining Session Time:</td>
      <td><asp:textbox width=250 id="countdown" runat="server" /></td>
   </tr>
   <tr class="smalltext">
      <td colspan=2><br>You may use this window to manually logout your session at any time, or
you may safely close this window if you prefer to let your session timeout
automatically.</font></td>
   </td>
   <tr>
      <td colspan=2><center><asp:button id="btnSubmit" class="button" text="  Logout  "
onClick="btnSubmit_Click" runat="server" /></center></td>
   </tr>
</table>
</form>
</span>

<span id="loggedOut" runat="server">
<form id="logout" runat="server">
<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr class="heading">
```

```
      <td colspan=2 align="center"> </td>
   </tr>
   <tr class="heading">
      <td colspan=2 align="center"><font color="white"><b>SonicWall LHM Logout
Window</b></font></td>
   </tr>
   <tr class="heading">
      <td colspan=2 align="center"> </td>
   </tr>
   <tr>
      <td><asp:Label id=LHMResult runat="server" /></td>
   </tr>
   <tr>
      <td><asp:Label id=catchError runat="server" /></td>
   </tr>
   <tr><td><br></td></tr>
   <tr>
      <td><center><asp:button id="Button1" class="button" text="  Close  " runat="server"
/></center></td>
   </tr>
</table>
</form>
</span>

</BODY>
</HTML>
```

myvars.aspx

```
<script language="VB" runat="server">

'Set the logoutPopup window flag - 0 = no popup, 1 = popup
'The use of the logoutPopup in this script is discouraged because although the login event
'is non-exclusive, the login event produces data where redundancy is undesirable.
Dim logoutPopup as String = "0"

'Set the passing score
Dim passingScore as Integer = 80

'Set the filename of the quiz XML source
Dim quizFile as String = "quiz.xml"
'Dim quizFile as String = "shortquiz.xml"

'Set the name of the Quiz
Dim quizName as String = "SonicOS Quiz"

'Set the emailed quiz results "from" email address
Dim quizFrom as String = "joelevy@SonicWall.com"

'Set the email address to send failed test results to (the proctor/instructor)
Dim quizTo as String = "joelevy@SonicWall.com"

'Set the path for check and block embedded images - usually the same path as the quiz
Dim imagePath as String = "C:\inetpub\wwwroot\lhm\lhmquiz\"

'Set the IP or resolvable FQDN for the SMTP Server
'Make sure the server is configured to relay from the IP address of this server
'If setting to 127.0.0.1 (local IIS SMTP), you need to allow IIS SMTP to relay from 127.0.0.1
Dim smtpServer as String = "127.0.0.1"

'Set the LHM Session Timeout
Dim sessTimer as String = "86400"

'Set the LHM Idle Timeout
Dim idleTimer as String = "3600"
```

```
'Set the secret for use with optional HMAC auth, as configured in the Extern Guest Auth config
on the SonicWall
Dim strHmac as String = "password"

'Set the digest method for the HMAC, either MD5 or SHA1
Dim hmacType as String = "MD5"
'Dim hmacType as String = "SHA1"

'Set the logo image to use
Dim logo as String = "SonicWall.gif"

'------------------------End of Configurable Settings-------------------------

</script>
```

quiz.aspx

```
<%@ Page Language="VB" Debug="true" %>
<%@ Import Namespace="System" %>
<%@ Import Namespace="System.Xml" %>
<%@ Import Namespace="System.Net" %>
<%@ Import Namespace="System.IO" %>
<%@ Import Namespace="System.Text" %>
<%@ Import Namespace="System.Security" %>
<%@ Import namespace="System.Security.Cryptography.X509Certificates"%>
<%@ Import Namespace="System.Web" %>
<%@ Import Namespace="System.Web.Mail" %>

<!-- Original quiz code from www.codeproject.com -->

<!-- #INCLUDE file="myvars.aspx" -->

<script language="VB" runat="server">

'This class allows SSL certs signed by unknown CAs to be accepted.
'This is necessary for the POST to the SonicWall authorizing the LHM session.
Public Class acceptAllCerts
   Implements System.Net.ICertificatePolicy
   Public Function CheckValidationResult(ByVal srvPoint As ServicePoint, _
   ByVal cert As X509Certificate, ByVal request As WebRequest, ByVal problem As Integer) _
      As Boolean Implements ICertificatePolicy.CheckValidationResult
      Return True
   End Function
End Class

'Set the path to the XML quiz data
Dim strXmlFilePath as String = Server.MapPath(quizFile)

'Setup our variables
Dim emailAddr as String
Dim userName as String
Dim req as String
Dim sessionId as String
Dim mgmtBaseUrl as String
Dim xDoc as XmlDocument = New XmlDocument()
Dim intTotalQuestion as Integer
Dim intQuestionNo as Integer = 1
Dim intScore as Integer = 0
Dim arrAnswerHistory as new ArrayList()
Dim arrRightOrWrong as new ArrayList()
Dim arrCorrect as new ArrayList()

Sub Page_Load(src as Object, e as EventArgs)

   LHMResult.Text=""
```

```
   catchError.Text=""

   'Grab context items set in default.aspx
   emailAddr = Context.Items("emailAddr")
   userName = Context.Items("userName")
   req = Context.Items("req")
   sessionId = Context.Items("sessionId")
   mgmtBaseUrl = Context.Items("mgmtBaseUrl")

   'Use the override class in myvars.aspx to accept untrusted certificates from the SonicWall
   'This is necessary for the POST to the SonicWall authorizing the LHM session.
   System.Net.ServicePointManager.CertificatePolicy = New acceptAllCerts

   'Load xml data
   xDoc.Load(strXmlFilePath)

   'Start a new quiz?
   If Not Page.IsPostBack Then

      'Yes. Count total question
      intTotalQuestion = xDoc.SelectNodes("/quiz/mchoice").Count

      'Record start time
      ViewState("StartTime") = DateTime.Now

      ShowQuestion(intQuestionNo)
   End If
End Sub

Sub btnSubmit_Click(src as Object, e as EventArgs)

   'Retrieve variables from ViewState
   intTotalQuestion = ViewState("TotalQuestion")
   intQuestionNo = ViewState("QuestionNo")
   intScore = ViewState("Score")
   arrAnswerHistory = ViewState("AnswerHistory")
   arrRightOrWrong = ViewState("RightOrWrong")
   arrCorrect = ViewState("AnswerList")
   req = ViewState("origReq")
   userName = ViewState("origUserName")
   emailAddr = ViewState("origEmailAddr")
   mgmtBaseUrl = ViewState("mgmtUrl")
   sessionId = ViewState("sessID")

   'Correct answer?
   If rblAnswer.SelectedItem.Value = ViewState("CorrectAnswer") Then
      intScore += 1
      arrRightOrWrong.Add(0)
   Else
      arrRightOrWrong.Add(rblAnswer.SelectedItem.Value)
   End If

   'Remember all selected answers
   arrAnswerHistory.Add(rblAnswer.SelectedItem.Value)
   arrCorrect.Add(ViewState("CorrectAnswer"))

   'End of quiz?
   If intQuestionNo=intTotalQuestion Then

      'Yes. Show the result.
      QuizScreen.Visible = False
      ResultScreen.Visible = True

      'Render result screen
      ShowResult()

   Else
```

```
      'Not yet. Show another question.
      QuizScreen.Visible = True
      ResultScreen.Visible = False
      intQuestionNo += 1

      'Render next question
      ShowQuestion(intQuestionNo)
   End If
End Sub

Sub ShowQuestion(intQuestionNo as Integer)
   Dim xNodeList as XmlNodeList
   Dim xNodeAttr as Object
   Dim strXPath as String
   Dim i as Integer
   Dim tsTimeSpent as TimeSpan

   strXPath = "/quiz/mchoice[" & intQuestionNo.ToString() & "]"

   'Extract question
   lblQuestion.Text = intQuestionNo.ToString() & ". " & xDoc.SelectSingleNode(strXPath &
"/question").InnerXml

   'Extract answers
   xNodeList = xDoc.SelectNodes(strXPath & "/answer")

   'Clear previous listitems
   rblAnswer.Items.Clear

   For i = 0 to xNodeList.Count-1

      'Add item to radiobuttonlist
      rblAnswer.Items.Add(new ListItem(xNodeList.Item(i).InnerText, i+1))

      'Extract correct answer
      xNodeAttr = xNodeList.Item(i).Attributes.ItemOf("correct")
      If not xNodeAttr is Nothing Then
         If xNodeAttr.Value = "yes" Then
            ViewState("CorrectAnswer") = i+1
         End If
      End If
   Next

   'Output Total Question and passing score
   lblTotalQuestion.Text = intTotalQuestion
   lblPassingScore.Text = passingScore

   'Output Time Spent
   tsTimeSpent = DateTime.Now.Subtract(ViewState("StartTime"))
   lblTimeSpent.Text = tsTimeSpent.Minutes.ToString() & ":" & tsTimeSpent.Seconds.ToString()

   'Store data to viewstate
   ViewState("TotalQuestion") = intTotalQuestion
   ViewState("Score") = intScore
   ViewState("QuestionNo") = intQuestionNo
   ViewState("AnswerHistory") = arrAnswerHistory
   ViewState("RightOrWrong") = arrRightOrWrong
   ViewState("AnswerList") = arrCorrect
   ViewState("origReq")=req
   ViewState("origUserName")=userName
   ViewState("origEmailAddr")=emailAddr
   ViewState("mgmtUrl")=mgmtBaseUrl
   ViewState("sessID")=sessionID

End Sub
```

```
Sub ShowResult()
   Dim strResult as String
   Dim intCompetency as Integer
   Dim i as Integer
   Dim strXPath as String
   Dim tsTimeSpent as TimeSpan

   tsTimeSpent = DateTime.Now.Subtract(ViewState("StartTime"))

   strResult  = "<center>"

   if passingScore <= Int(intScore/intTotalQuestion*100).ToString()
      strResult += "<h2><font color=""green"">You Passed!</h3></font>"
   else
      strResult += "<h2><font color=""red"">You Failed!</h3><b>Please review the answers and
retake the test.</b><br></font>"
   End If

   strResult += "User Name: " & userName & "<br>"
   strResult += "Elapsed Time: " & tsTimeSpent.Minutes.ToString() & ":" &
tsTimeSpent.Seconds.ToString() & "<br>"
   strResult += "Correct Answers: " & intScore.ToString() & " out of " &
intTotalQuestion.ToString() & "<br>"
   strResult += "Your Percentage: " & Int(intScore/intTotalQuestion*100).ToString() & "%<br>"
   strResult += "Required Percentage:" & passingScore.ToString() & "%<br>"
   strResult += "</center>"

   strResult += "<h3>Quiz Results</h3>"
   For i = 1 to intTotalQuestion
      strXPath = "/quiz/mchoice[" & i.ToString() & "]"
      strResult += "<b>" & i.ToString() & ". " & xDoc.SelectNodes(strXPath &
"/question").Item(0).InnerXml & "</b><br>"
      If arrRightOrWrong.Item(i-1)=0 Then
         strResult += "<img src = ""checkMark.gif""><font color=""green""> "
         strResult += "<b>You answered:</b> " & xDoc.SelectNodes(strXPath & "/answer[" &
arrAnswerHistory.Item(i-1).ToString() & "]").Item(0).InnerXml & "</font><br><br>"
      Else
         strResult += "<img src = ""Block.gif""><font color=""red""> "
         strResult += "<b>You answered:</b> " & xDoc.SelectNodes(strXPath & "/answer[" &
arrAnswerHistory.Item(i-1).ToString() & "]").Item(0).InnerXml & "<br>"
         strResult += "The correct anwer is: " & xDoc.SelectNodes(strXPath & "/answer[" &
arrCorrect.Item(i-1).ToString() & "]").Item(0).InnerXml & "</font><br><br>"
      End If
   Next

   'Setup the common Mail settings
   Dim objMail As MailMessage
   objMail = New MailMessage()
   objMail.From = quizFrom
   objMail.Body = strResult
   objMail.BodyFormat = MailFormat.Html

   'Path to the attachments for the Check and X images - update these in myvars.aspx
   objMail.Attachments.Add(New MailAttachment(imagePath & "block.gif"))
   objMail.Attachments.Add(New MailAttachment(imagePath & "checkMark.gif"))

   'Address of the SMTP server - can be localhost if SMTP is running on IIS - in myvars.aspx
   SmtpMail.SmtpServer = smtpServer

   'Determine pass/fail
   If passingScore <= Int(intScore/intTotalQuestion*100).ToString()

      'Mail the passing test result to the test-taker
      'Be sure to update the mail fields in myvars.aspx
      objMail.To =emailAddr
      objMail.Subject = quizName & " Results for " & emailAddr
```

```
      'Send the mail
      SmtpMail.Send(objMail)
      strResult +="Your test is being emailed to you at " & emailAddr

      'Send the session Auth message to LHM
      postLHM()

   else
      'Mail failing test results to the instuctor
      objMail.To =quizTo
      objMail.Subject = "Failing " & quizName & " Test Results for " & emailAddr

      'Send the mail
      SmtpMail.Send(objMail)
      strResult += "<a href=""quiz.aspx"">Click here to retake the quiz</a>"
   End If

   'Write it
   lblResult.Text = strResult

End Sub

Sub postLHM()

   'The LHM cgi on the SonicWall - this does not change
   Dim loginCgi as String = "externalGuestLogin.cgi"

   'Assemble the data to post back to the SonicWall to authorize the LHM session
   Dim loginParams as String = "sessId=" & sessionId & "&userName=" &
Server.URLEncode(userName) & "&sessionLifetime=" & sessTimer & "&idleTimeout=" & idleTimer

   'Combine mgmtBaseUrl from the original redirect with the login cgi
   Dim postToSNWL as String = mgmtBaseUrl & loginCgi

   'Convert the loginParams to a well behaved byte array
   Dim byteArray As Byte() = Encoding.UTF8.GetBytes(loginParams)

   Try
      'Let the user know that we are setting up the session, just in case it takes more than a
second
      LHMResult.Text = "Authorizing session. Please wait."

      'Create the webrequest to the SonicWall
      Dim toSNWL as WebRequest = WebRequest.Create(postToSNWL)

      'Calculate the length of the byte array
      toSNWL.ContentLength = byteArray.Length

      'Set the method for the webrequest to POST
      toSNWL.Method = "POST"

      'Set the content type
      toSNWL.ContentType = "application/x-www-form-urlencoded"

      'Open the request stream
      Dim dataStream As Stream = toSNWL.GetRequestStream()

      'Write the byte array to the request stream
      dataStream.Write(byteArray, 0, byteArray.Length)

      'Close the Stream object
      dataStream.Close()

      'Get the response
      Dim snwlReply As WebResponse = toSNWL.GetResponse()

      'Display the status - looking for 200 = OK.
```

```
      'Response.Write(CType(snwlReply, HttpWebResponse).StatusCode)

      'Grab the response and stuff it into an xml doc for possible review
      Dim snwlResponse as XmlDocument = New XmlDocument()
      snwlResponse.Load(snwlReply.GetResponseStream())

      'Set the xPath to the SNWL reply, and get the response
      Dim codePath as String = "SonicWallAccessGatewayParam/AuthenticationReply/ResponseCode"

      'Response code 50 - Login Succeeded

      If snwlResponse.SelectSingleNode(codePath).InnerXml = "50"

         'Do we want to provide a logout popup window?
         If logoutPopup = "1" Then
            'Popup hack using Javascript for logout window
            Dim sb As New System.Text.StringBuilder()
            sb.Append("<script language='javascript'>")
            sb.Append("window.open('logout.aspx?sessId=")
            sb.Append(Server.URLEncode(CStr(sessionId)))
            sb.Append("&mgmtBaseUrl=")
            sb.Append(Server.URLEncode(CStr(mgmtBaseUrl)))
            sb.Append("&sessTimer=")
            sb.Append(Server.URLEncode(CStr(sessTimer)))
            sb.Append("','logOut','toolbar=no,")
            sb.Append("addressbar=no,menubar=no,")
            sb.Append("width=400,height=250');")
            sb.Append("<")
            sb.Append("/")
            sb.Append("script>")
            RegisterStartupScript("stp", sb.ToString)
         End If

         LHMResult.Text = "<br><b><font color=""green"">Session authorized:</font></b> You may
now go to the URL you originally requested: <a target=""_blank"" href=""" & req & """>" & req &
"</a>"

      'Response code 51 - Session Limit Exceeded
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "51"
         LHMResult.Text = "<br><b><font color=""red"">Session Limit Reached:</font></b> The
maximum number of guest session has been reached. Sorry for the inconvenience. Please close and
relaunch your browser to try again."

      'Response code 100 - Login Failed.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "100"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> Your
session cannot be created at this time. Sorry for the inconvenience. Please close and relaunch
your browser to try again."

      'Response code 251 - Bad HMAC.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "251"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization failed message authentication. Sorry for the inconvenience. Please
close and relaunch your browser to try again."

      'Response code 253 - Invalid SessionID.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "253"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization failed to match a known session identity. Sorry for the
inconvenience. Please close and relaunch your browser to try again."

      'Response code 254 - Invalid CGI.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "254"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization was missing an essential parameter. Sorry for the inconvenience.
Please close and relaunch your browser to try again."
```

```
      'Response code 255 - Internal Error.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "255"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization failed due to an unspecified error. Sorry for the inconvenience.
Please close and relaunch your browser to try again."

      End If

      'Close the streams
      dataStream.Close()
      snwlReply.Close()

      'If there is some asp.net error trying to talk to the SonicWall, print it
      'in the same color as the background, but still show the quiz results.
      Catch ex as Exception
         catchError.Text = "<font color=""9CBACE"">" & ex.ToString & "</font>"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization failed due to an unspecified error. Sorry for the inconvenience.
Please close and relaunch your browser to try again. If the problem persists, please notify an
attendant."
   End Try
End Sub

</script>
<html>
<head>
<title><%= quizName %> </title>
</head>
<style>
body {
  font-size: 10pt;
  font-family: verdana,helvetica,arial,sans-serif;
  color:#000000;
  background-color:#9CBACE;
}

tr.heading {
  background-color:#006699;
}

.button {
   border: 1px solid #000000;
   background-color: #ffffff;
}
</style>

<HTML>
<HEAD>
<TITLE>LHM Quiz Script</TITLE>
</HEAD>

<body>
<span id="QuizScreen" runat="server">
<form runat="server">
<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr class="heading">
      <td colspan=3 align="center"><font color="white"> </td>
   </tr>
   <tr class="heading">
      <td width="50%" valign="center"><font color="white"><b><%= quizName %>  - <%=
userName%></b></font></td>
      <td><center><img width="216" height="51" src="<%= logo %>"></center></td>
      <td width="50%" align="right" valign="center"><font color="white"><b>This quiz has
<asp:label id="lblTotalQuestion" runat="server" /> questions</b></font></td>
   </tr>
   <tr class="heading">
      <td colspan=3 align="center"><font color="white"> </td>
```

```
   </tr>
</table>
<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr>
      <td colspan="2">
         <b><asp:label id="lblQuestion" runat="server" /></b><br>
         <asp:radiobuttonlist id="rblAnswer" RepeatDirection="vertical" TextAlign="right"
RepeatLayout="table" runat="server" /><br>
         <asp:button id="btnSubmit" class="button" text="  Submit  " onClick="btnSubmit_Click"
runat="server" />
         <asp:requiredfieldvalidator ControlToValidate="rblAnswer" ErrorMessage="Please select
an answer" runat="server" />
      </td>
   </tr>
   <tr class="heading">
      <td width="70%"><font color="white"><b>Score required to pass <asp:label
id="lblPassingScore" runat="server" />%</b></font></td>
      <td width="30%" align="right"><font color="white"><b>Time spent <asp:label
id="lblTimeSpent" runat="server" /></b></font></td>
   </tr>
</table>
</form>
</span>

<span id="ResultScreen" runat="server"> <asp:label id="lblResult" runat="server" /> <br>
<asp:Label id=LHMResult runat="server" />
<asp:Label id=catchError runat="server" />
</span>

</body>
</html>
```

PayPal Script

Authentication Model	The Guest Client buys 1-hour or 24-hour access with a **Buy Now** button using their PayPal account. Payment is made through PayPal to the hotspot provider's PayPal merchant account.
Purpose	Nearly everyone who buys or sells on the Internet uses PayPal. It is very easy to setup a buyer account, and to link it to any form of payment (such as credit card, bank card, checking account). It is almost equally easy to upgrade a buyer-only account to a merchant account. Having a merchant account allows PayPal users to accept payment from other PayPal users for goods or services. The funds transfer is run through PayPal, providing merchants a way to do business online, accepting any form of payment, without having to setup any sort of complicated payment processing. This eliminates what is perhaps the single biggest obstacle to being a fee-based hotspot provider. Paypal provides a feature called **Buy Now**, which allows for one-click transactions. The buttons are forms, generated with the assistance of PayPal, that contain information about the item or service being purchased. When the buyer clicks on **Buy Now**, the session is redirected to the PayPal site with a `querystring` containing all the details of the transaction (such as the seller, the item, the price). Rather than using **Buy Now** (which is client-side rather than server-side code), the PayPal script uses a custom, server-side Buy Now routine.

Also included in the Buy Now redirect is the path for the auto-return. Auto-return is a PayPal feature that sends the buyer back to the merchant's site after the PayPal transaction. Auto-return is required when using PDT (`pdtPath`, described below).

The custom Buy Now redirect also embeds the LHM `sessionID` and the `mgmtBaseUrl` into a custom string in the Buy Now redirect to PayPal. This allows us to track the session even though it leaves the LHM server, goes to PayPal, and then comes back (through auto-return for PDT).

The basic PayPal payment system provides notification of payment to merchants by email. This is acceptable for physical goods because the purchase/ship transaction does not have to occur in real-time; the merchant can wait hours or days for the notification before shipping the product. For transactions that require instantaneous delivery, such as buying hotspot access, a more real-time method of payment is required.

PayPal offers two methods of payment notification:

- Instant Payment Notification (IPN), which works by PayPal making a web-services call to the merchant's site indicating that payment for a particular transaction has cleared. Unfortunately, this does not always occur in real-time (it can take up to 20 minutes for this asynchronous notification to arrive), so it was not employed in this script. (More can be read about IPN at https://www.paypal.com/cgi-bin/webscr?cmd=p/xcl/rec/ipn-intro-outside)
- Payment Data Transfer (PDT: see http://www.paypal.com/cgi-bin/webscr?cmd=p/xcl/rec/pdt-intro-outside). This method occurs in absolute real-time using PayPal's auto-return method. PDT provides instant notification to the merchant of the state of a transaction (either SUCCESS or FAIL), as well as of the `payment_status` (Completed, Pending, Denied, Failed, Refunded, Reversed, or Cancelled_Reversal). By instantly knowing the status of the transaction and the payment, it is possible to immediately provide service without the risk of losing payment.

myvars Variables

`logoutPopup` — Controls the use of the logout popup window. Set to:

- **0** to disable the popup window.
- **1** to enable the popup window.

`debugFlag` — Sets the debug output for the PayPal PDT transfer:

- **0** = Off
- **1** = On

`pdtPath` — The path to which the Guest Client is redirected by the PDT auto-return (described above in the Purpose section).

`paypalCGI` — The URL for the PayPal CGI serving as the gateway for the PayPal transaction. The URL itself should not be changed, but there are two options; either the:

- Live (real) PayPal site.
- Paypal sandbox (part of the PayPal developer network), which can be used for testing.

`myBusiness` — The email address (how PayPal recognizes the business) of the hotspot provider. This must match the email address of the merchant account that is receiving payment for the transactions.

`token` — The Payment Data Transfer option generates a unique token for each merchant. This is where you specify your PayPal-provided unique token. The token must be correct, or the PDT transaction (not the actual PayPal transaction) fails.

`itemName1` `itemName2`	The names of the two access options, such as `1 Hour Secure Internet Access` and `24 Hours Secure Internet Access`.
`itemNumber1` `itemNumber2`	The item number (a mostly arbitrary internal PayPal reference) for the two access options, such as `1hour` and `24hour`.
`itemTimer1` `itemTimer2`	The session timer, in seconds, for the two access options, such as `3600` for 1 hour and `86400` for 24 hours.
`itemAmount1` `itemAmount2`	The price in US dollars for the two access options, such as `0.01` (one cent) and `0.02` (two cents). Limited time promotional bargain pricing.
`itemButton1` `itemButton2`	The button text for the two access options, such as `1 Hour Access - $0.01` and `24 Hours Access - $0.02`.
`strHmac`	The shared secret for the optional HMAC function.
`hmacType`	The digest type to use if HMAC is in use: **MD5** or **SHA1**.
`logo`	The names of the logo (image) file to use on page headers.

Session Flow

1 The Guest Client launches their web-browser, and is redirected by LHM to `http://<lhmserver>/paypal/default.aspx`, where `<lhmserver>` is your LHM server.

2 Guest client (buyer) clicks on one of the **Buy Now** buttons, such as **1 Hour Access - $0.01**.

3 The client is redirected to the PayPal site with a `querystring` containing all the information about the merchant, the item, the LHM session (in the custom variable), and the auto-return URL (defined in `myvars` as pdtPath).

The `pdtPath` resides on the LHM server. The path should be the same as the `default.aspx` path (as configured on the SonicWall Security Appliance), but should point to the `pdt.aspx` file. This way, when the PayPal transaction is completed and PayPal redirects the client back to the merchant site, the client is redirected back to the `http://<lhmserver>/paypal/pdt.aspx` page.

HTTP can be used on the LHM Server because no sensitive information is entered on the LHM server itself; the PayPal transaction occurs through HTTPS directly between the Guest Client and PayPal.

Sample Buy Now redirect string:

```
https://www.sandbox.paypal.com/cgi-
bin/webscr?cmd=_xclick&business=demo@SonicWall.co
m&item_name=1%20Hour%20Access&item_number=1hour&a
mount=0.01&currency_code=USD&lc=US&bn=PP-
BuyNowBF&no_note=1&no_shipping=1&cancel_return=ht
tp://lhmserverpaypal/default.aspx&return=http://l
hmserver/lhm/paypal/pdt.aspx&custom=35378e67833fa
a3de83aa3b771https%3a%2f%2f172.16.17.1%3a4043%2f
```

4 The Guest Client logs into PayPal (or creates a new account, as needed) and completes the transaction with PayPal. After the transaction is completed, the client is redirected back to `http://<lhmserver>/paypal/pdt.aspx`. Included in the redirect is a `querystring` containing the transaction id (`tx`), the status (`st`), the amount (`amt`), the currency type (`cc`), the custom value (`cm`), and an encrypted signature (`sig`).

Sample redirect string:

```
http://lhmserver/lhm/paypal/pdt.aspx?tx=4LN76482J
F4605045&st=Completed&amt=0.01&cc=USD&cm=35378e67
833faa3b771https%3a%2f%2f172%2e16%2e17%2e1%3a4043
%2f&sig=qdsNC4flKwtPviggoGAXCpeV9gS%2f2E%2bGGVbTZ
3STrUV1Ci9K3c2zTdJMuuKCmRiif1SybsZtUqDYqzzfMg64AF
3PKCk85rrPubYT4K4aC
```

5 The Guest Client accessing the `pdt.aspx` script at the URL above starts the PDT process on the LHM server. The script builds a `querystring` consisting on `cmd=_notify-synch` (indicating that it is a PDT transacation) along with the `tx` (transaction ID) and the `at` variable set to the merchant's token (defined in `myvars`). This is then POSTed to the `paypalCGI` URL (as defined in `myvars`).

6 PayPal responds to the POST with a SUCCESS of a FAIL code.

- FAIL – the script indicates to the client that the PayPal transaction fails, and they are prompted to seek assistance.
- SUCCESS – provides details about the transaction:

```
SUCCESS
txn_type=web_accept
payment_date=00%3A39%3A48+Oct+30%2C+2005+PDT
last_name=Niqual
item_name=1+Hour+Secure+Internet+Access
payment_gross=0.01
mc_currency=USD
business=lhmdemo%40SonicWall.com
payment_type=instant
payer_status=verified
tax=0.00
payer_email=lhmClient%40SonicWall.com
txn_id=84K306380G150640T
quantity=1
receiver_email=lhmdemo%40SonicWall.com
first_name=Sah
payer_id=XWRZGABD6UV2W
receiver_id=REW4W5WANU294
item_number=1hour
payment_status=Completed
payment_fee=0.01
mc_fee=0.01
shipping=0.00
mc_gross=0.01
custom=35378e67833faa3de833755d3aa3b771https%3A//172
.16.17.1%3A4043/
charset=windows-1252
```

7 The script checks the `payment_status` to make sure the payment is completed. If it is not completed, an incomplete-payment message is provided to the user.

8 If `payment_status` is completed, the script also obtains the client name, item name, amount, transaction ID, business, and custom variables for generating the client's receipt, a userName for the LHM session, and identifying the LHM `sessionID` and `mgmtBaseUrl`.

9 The script presents the PayPal transaction receipt to the Guest Client.

10 The script performs the LHM post to the SonicWall Security Appliance to authorize the session.

Additional Considerations

Requires a PayPal merchant account.

Requires that the PayPal account be setup for auto-return and for PDT (see http://www.paypal.com/cgi-bin/webscr?cmd=p/xcl/rec/pdt-techview-outside)

For testing, it is strongly suggested that a (free) PayPal sandbox account be setup through the PayPal Developer's Network (https://developer.paypal.com) and (https://www.sandbox.paypal.com)

IMPORTANT: Because the Guest Client is redirected directly to the PayPal site, ALL PayPal site IP addresses must be setup on the SonicWall Security Appliance as Allowed Networks on the Guest Services configuration. These include the following:

www.paypal.com

```
64.4.241.32
64.4.241.33
216.113.188.32
216.113.188.35
216.113.188.66
216.113.188.67
```

www.paypalobjects.com

```
216.113.188.25
64.4.241.62
216.113.188.9
```

www.sandbox.paypal.com

```
66.135.197.160
```

developer.paypal.com

```
66.135.197.163
```

Topics:

- default.aspx on page 892
- logout.aspx on page 897
- myvars.aspx on page 902
- pdt.aspx on page 903

default.aspx

```
<%@ Page Language="VB" Debug="true" %>
<%@ Import Namespace="System" %>
<%@ Import Namespace="System.Net" %>
<%@ Import Namespace="System.IO" %>
<%@ Import Namespace="System.Xml" %>
<%@ Import Namespace="System.Text" %>
<%@ Import Namespace="System.Security" %>
<%@ Import namespace="System.Security.Cryptography.X509Certificates"%>
```

```
<!-- #INCLUDE file="myvars.aspx" -->

<script language="VB" runat="server">

'This class allows SSL certs signed by unknown CAs to be accepted.
'This is necessary for the POST to the SonicWall authorizing the LHM session.
Public Class acceptAllCerts
   Implements System.Net.ICertificatePolicy
   Public Function CheckValidationResult(ByVal srvPoint As ServicePoint, _
   ByVal cert As X509Certificate, ByVal request As WebRequest, ByVal problem As Integer) _
      As Boolean Implements ICertificatePolicy.CheckValidationResult
      Return True
   End Function
End Class

'Note: For PayPal authorization to work, it is necessary to set up the PayPal sites
(www.paypal.com, www.paypalobjects.com, and www.sandbox.paypal.com) as a bypass network on
WGS. This is so that WGS/LHM users can access PayPal directly to complete the payment
transactions. This list currently includes the following addresses:  [64.4.241.32, 64.4.241.33,
216.113.188.32, 216.113.188.35, 216.113.188.66, 216.113.188.67], [216.113.188.25, 64.4.241.62,
216.113.188.9] and [66.135.197.160].

'Sample LHM redirect querystring:
'http://127.0.0.1/lhm/paypal/default.aspx?sessionId=0b712fd83b9f5313db5af1cea6b1004f&ip=10.50.
165.231&mac=00:0e:35:bd:c9:37&ufi=0006b11184300&mgmtBaseUrl=https://10.50.165.193:4043/&client
RedirectUrl=https://10.50.165.193:444/&req=http%3A//www.google.com/ig

Dim ip as String
Dim sessionId as String
Dim mac as String
Dim ufi as String
Dim mgmtBaseUrl as String
Dim clientRedirectUrl as String
Dim req as String
Dim hmac as String
Dim customCode as String

Sub Page_Load(src as Object, e as EventArgs)

   ip=Request.QueryString("ip")
   sessionId=Request.QueryString("sessionId")
   mac=Request.QueryString("mac")
   ufi=Request.QueryString("ufi")
   mgmtBaseUrl=Request.QueryString("mgmtBaseUrl")
   clientRedirectUrl=Request.QueryString("clientRedirectUrl")
   req=Request.QueryString("req")
   hmac=Request.QueryString("hmac")
   customCode=Request.QueryString("cc")

   'customCode grabs the "cc=" querystring value sent by the SonicWall. This allows you to use
the same
   'page (e.g. this page) for the "Session Expiration" (?cc=2), "Idle Timeout" (?cc=3) and "Max
Sessions" (?cc=4) page.
   If customCode <> "" Then
      Select Case customCode
      Case "2"
         LHMResult.Text="<br><H3><font color=""red"">Your LHM session has expired. You may try
to initiate a new session.</font></H3>"
      Case "3"
         LHMResult.Text="<br><H3><font color=""red"">You have exceeded your idle timeout.
Please log back in.</font></H3>"
      Case "4"
         LHMResult.Text="<br><H3><font color=""red"">The maximum number of sessions has been
reached. Please try again later.</font></H3>"
      End Select
   End If
```

```
   'Set the button Text for the two buttons with the variable configured in myvars
   btnBuyNow1.Text=itemButton1
   btnBuyNow2.Text=itemButton2

   'Use the override class in myvars.aspx to accept untrusted certificates from the SonicWall
   'This is necessary for the POST to the SonicWall authorizing the LHM session.
   System.Net.ServicePointManager.CertificatePolicy = New acceptAllCerts

   'Note - the routine below for handling the hmac requires the use of the SonicSSL.dll and
libeay.dll libraries.
   'The DLL must be copied to the IIS server, and the SonicSSL dll must be registered with
"regsvr32 sonicssl.dll"
   If hmac <> "" Then

      'SonicWall URL Encode routine is different from Microsoft - this is the SonicWall method
      req=Replace(req,"%","%25")
      req=Replace(req,":","%3A")
      req=Replace(req," ","%20")
      req=Replace(req,"?","%3F")
      req=Replace(req,"+","%2B")
      req=Replace(req,"&","%26")
      req=Replace(req,"=","%3D")

      Dim strHmacText as String
      Dim objCrypto as Object
      Dim strHmacGenerated
      Dim loginError as String

      'Initialize the Crypto object
      objCrypto = Server.CreateObject("SonicSSL.Crypto")

      'The text to be encoded
      strHmacText = sessionId & ip & mac & ufi & mgmtBaseUrl & clientRedirectUrl & req

      'Calculate the hash with a key strHmac, the return value is a string converted form the
output sha1 binary.
      'The hash algorithm (MD5 or SHA1) is configured in myvars and in the Extern Guest Auth
config on the SonicWall
      If hmacType = "MD5" Then
         strHmacGenerated = objCrypto.hmac_md5(strHmacText, strHmac)
      Else
         strHmacGenerated = objCrypto.hmac_sha1(strHmacText, strHmac)
      End If

      If strHmacGenerated <> hmac Then
         Dim hmacFail as String
         hmacFail = "<font color=""red"">The HMAC failed validation. Please notify an
attendant.</font><br><br>"
         hmacFail+="<font color=""9CBACE"">Received HMAC: " & hmac & "<br>Calculated HMAC: " &
strHmacGenerated & "<br>"
         hmacFail+="Make sure the digest functions on the SonicWall and LHM server match.<br>"
         hmacFail+="Also make sure the shared secret on the SonicWall and myvars match</font>"
         catchError.Text=hmacFail
      End If

   End If

End Sub

Sub btnBuyNow_Click(Sender As Object, E As EventArgs)

'sample redirect generated by this routine:
'https://www.paypal.com/cgi-
bin/webscr?cmd=_xclick&business=jlevy@SonicWall.com&item_name=24%20Hour%20Secure%20Internet%20
Access&item_number=24hour&amount=0.02&currency_code=USD&lc=US&bn=PP-
BuyNowBF&no_note=1&no_shipping=1&cancel_return=http://127.0.0.1/lhm/paypal/default.aspx&return
=http://www.moosifer.com/pdt.aspx
```

```
'sample redirect from the paypal server back the LHM server on transaction completion
(modified).
'http://127.0.0.1/lhm/paypal/pdt.aspx?tx=4PG453F7LS133715V&st=Completed&amt=0.02&cc=USD&cm=&si
g=EZhZtJygi7RTXulJt4SEhVBRi%2bJwLaC9z9kRLsrsXk4gQKnzvI5vjGy0vdhKPXAVyhbh%2bwBxWon2cieEQDJ9P6R9
qqjuKnzvI5vjGy0vdhKPXAVyJ3GtOq5Jd3%2fvTY3s7FrRcKdKnzvI5vjGy0vdhKPXAVyyEKNxY3d

   Dim str, itemName, itemNumber, itemAmount As String
   Dim sb As New StringBuilder()

   'Determine which button was pressed, and set item attributes appropriately
   Select Case Sender.Text
      Case itemButton1
         itemName = itemName1
         itemNumber = itemNumber1
         itemAmount = itemAmount1
      Case itemButton2
         itemName = itemName2
         itemNumber = itemNumber2
         itemAmount = itemAmount2
   End Select

   'The paypal CGI URL - You can select either the real CGI or the sandbox CGI in myvars
   sb.Append(paypalCGI & "?")
   'The cmd passed to PayPal - do not change!
   sb.Append("cmd=_xclick")
   'The email address of the paypal merchant receiving payment. Replace in myvars with your
paypal email address.
   sb.Append("&business=" & myBusiness)
   'The name of the item being purchased. This is the first item option (e.g. 1 hour). Set in
myvars
   sb.Append("&item_name=" & itemName)
   'The optional item id
   sb.Append("&item_number=" & itemNumber)
   'The price being charged for the item (access)
   sb.Append("&amount=" & itemAmount)
   'The currency
   sb.Append("&currency_code=USD")
   'The country
   sb.Append("&lc=US")
   'The banana nullifier
   sb.Append("&bn=PP-BuyNowBF")
   'Disables the note option on the transaction
   sb.Append("&no_note=1")
   'Disables the shipping option on the transaction
   sb.Append("&no_shipping=1")
   'Build the path to return the client to (the LHM server address) on a cancelled transaction
   sb.Append("&cancel_return=http://" & Request.ServerVariables("SERVER_NAME") &
Request.ServerVariables("URL"))
   'The return (success page) path to return the buyer to after the transaction. This is the
PDT receiver/processor page.
   sb.Append("&return=" & pdtPath)
   'The LHM sessionID - append this so that it can be returned to us later by the PDT
transaction - do not change!
   sb.Append("&custom=" & sessionId & Server.URLEncode(mgmtBaseUrl))
   'Optional notify_url that paypal will asynchronously send IPN confirmation to. Not used
since it's not real-time.
   'sb.Append("&notify_url=http://www.moosifer.com/ipn.aspx")
   str = sb.ToString
   Response.Redirect(str)

End Sub

</script>

<STYLE>
body {
```

```
  font-size: 10pt;
  font-family: verdana,helvetica,arial,sans-serif;
  color:#000000;
  background-color:#9CBACE;
}

tr.heading {
  background-color:#006699;
}

.button {
   border: 1px solid #000000;
   background-color: #ffffff;
}
</STYLE>

<HTML>
<HEAD>
<TITLE>LHM PayPal Script</TITLE>
</HEAD>

<BODY>
<form id="frmValidator" runat="server">

<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr class="heading">
      <td colspan=3 align="center"><font color="white"> </td>
   </tr>
   <tr class="heading">
      <td width="50%" valign="center"><font color="white"><b>LHM Access with PayPal Buy
Now</b></font></td>
      <td><center><img width="216" height="51" src="<%= logo%>"></center></td>
      <td width="50%" align="right" valign="center"><font color="white"><b>Powered by
SonicWall LHM</b> </font></td>
   </tr>
   <tr class="heading">
      <td colspan=3 align="center"><font color="white"> </td>
   </tr>
</table>

<table width="100%" border="0" cellpadding="2" cellspacing="0">

   <tr>
      <td colspan=3><br></td>
   </tr>
   <tr>
      <td colspan=3 align="left">Purchase Secure Internet Access through SonicWall's LHM and
PayPal's Buy Now feature.
      <br><br>The two Buy Now buttons below will send you to PayPal's website where you can use
your PayPal account to pay <b>$<%= itemAmount1 %> for <%= itemName1 %></b>, or <b>$<%=
itemAmount2 %> for <%= itemName2 %></b>.
      <br><br>
      PayPal will then redirect you to this site to initiate the Payment Data Transfer (PDT)
exchange. The PDT exchange begins with the LHM server posting a paypal constructed querystring
back to PayPal. The response to the post will then be parsed by the LHM server to determine if
the PayPal transaction was successful. After all data are exchanged and verified, LHM will
authorize access on the SonicWall for the period of time purchased.
      <br><br>
      The clock for access will start immediately upon successful session authorization, and
can be used on the local SonicWall appliance by the client (as tracked by IP and MAC address)
so long as session time remains. The idle timeout will effectively be disabled by setting the
idle timer to the same value as the session timer.
      <br><br>
      Please select "<%= itemName1 %>" or "<%= itemName2 %>" below. You will be redirected to
the PayPal site, and will be returned to this site on transaction completion.

      <br><br>
```

```
      </td>
   </tr>
   <tr class="heading">
      <td colspan=3 align="center"><font color="white"> </td>
   </tr>
   <tr class="heading">
      <td align="center"><asp:Button ID="btnBuyNow1" Class="button" OnClick="btnBuyNow_Click"
runat="server" />
        <asp:Button ID="btnBuyNow2" Class="button" OnClick="btnBuyNow_Click"
runat="server" /></td>
   </tr>
   <tr class="heading">
      <td colspan=3 align="center"><font color="white"> </td>
   </tr>
   <tr>
      <td colspan=3><asp:Label id=LHMResult runat="server" /></td>
   </tr>
   <tr>
      <td colspan=3><asp:Label id=catchError runat="server"/></td>
   </tr>
</table>

</form>
</BODY>
</HTML>
```

logout.aspx

```
<%@ Page Language="VB" Debug="true" %>
<%@ Import Namespace="System" %>
<%@ Import Namespace="System.Net" %>
<%@ Import Namespace="System.IO" %>
<%@ Import Namespace="System.Xml" %>
<%@ Import Namespace="System.Text" %>
<%@ Import Namespace="System.Security" %>
<%@ Import namespace="System.Security.Cryptography.X509Certificates"%>

<!-- #INCLUDE file="myvars.aspx" -->

<script language="VB" runat="server">

'This class allows SSL certs signed by unknown CAs to be accepted.
'This is necessary for the POST to the SonicWall authorizing the LHM session.
Public Class acceptAllCerts
   Implements System.Net.ICertificatePolicy
   Public Function CheckValidationResult(ByVal srvPoint As ServicePoint, _
   ByVal cert As X509Certificate, ByVal request As WebRequest, ByVal problem As Integer) _
      As Boolean Implements ICertificatePolicy.CheckValidationResult
      Return True
   End Function
End Class

Dim sessionId as String
Dim mgmtBaseUrl as String
Dim eventId as String = "&eventId=1"

'Grab the code and the session lifetime from the generator page
Sub Page_Load(src as Object, e as EventArgs)
   sessionId=Request.QueryString("sessId")
   mgmtBaseUrl=Request.QueryString("mgmtBaseUrl")
   sessTimer=Request.QueryString("sessTimer")


   'Use the override class in myvars.aspx to accept untrusted certificates from the SonicWall
   'This is necessary for the POST to the SonicWall authorizing the LHM session.
   System.Net.ServicePointManager.CertificatePolicy = New acceptAllCerts
```

```
   'When the page loads, make the loggedIn span visible
   loggedIn.Visible=True
   loggedOut.Visible=False

   Me.Button1.Attributes.Add("OnClick", "self.close()")

End Sub

'The Logout button
Sub btnSubmit_Click(Sender As Object, E As EventArgs)

   'Let the user know that we are setting up the session, just in case it takes more than a
second
   LHMResult.Text = "Authorizing session. Please wait."

   'The LHM cgi on the SonicWall - this does not change
   Dim loginCgi as String = "externalGuestLogoff.cgi"

   'Assemble the data to post back to the SonicWall to authorize the LHM session
   Dim loginParams as String = "sessId=" & sessionId & eventId

   'Combine mgmtBaseUrl from the original redirect with the login cgi
   Dim postToSNWL as String = mgmtBaseUrl & loginCgi

   'Convert the loginParams to a well behaved byte array
   Dim byteArray As Byte() = Encoding.UTF8.GetBytes(loginParams)

   Try
      'Make the loggedOut span visible
      loggedIn.Visible=False
      loggedOut.Visible=True

      'Create the webrequest to the SonicWall
      Dim toSNWL as WebRequest = WebRequest.Create(postToSNWL)

      'Calculate the length of the byte array
      toSNWL.ContentLength = byteArray.Length

      'Set the method for the webrequest to POST
      toSNWL.Method = "POST"

      'Set the content type
      toSNWL.ContentType = "application/x-www-form-urlencoded"

      'Open the request stream
      Dim dataStream As Stream = toSNWL.GetRequestStream()

      'Write the byte array to the request stream
      dataStream.Write(byteArray, 0, byteArray.Length)

      'Close the Stream object
      dataStream.Close()

      'Get the response
      Dim snwlReply As WebResponse = toSNWL.GetResponse()

      'Display the status - looking for 200 = OK.
      'Response.Write(CType(snwlReply, HttpWebResponse).StatusCode)

      'Grab the response and stuff it into an xml doc for possible review
      Dim snwlResponse as XmlDocument = New XmlDocument()
      snwlResponse.Load(snwlReply.GetResponseStream())

      'Set the xPath to the SNWL reply, and get the response
      Dim codePath as String = "SonicWallAccessGatewayParam/LogoffReply/ResponseCode"
```

```
      'Response.Write(snwlResponse.SelectSingleNode(codePath).InnerXml)

      'Response code 150 - Logout Succeeded
      If snwlResponse.SelectSingleNode(codePath).InnerXml = "150"
         LHMResult.Text = "<br><b><font color=""green"">Your session has been logged
out.<br><br>Thank you for using LHM Guest Services.</font></b>"

      'Response code 251 - Bad HMAC.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "251"
         LHMResult.Text = "<br><b><font color=""red"">Session logout failed:</font></b> The
request failed message authentication. Sorry for the inconvenience. Please close and relaunch
your browser to try again."

      'Response code 253 - Invalid SessionID.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "253"
         LHMResult.Text = "<br><b><font color=""red"">Session logout failed:</font></b> The
request failed to match a known session identity. Sorry for the inconvenience. Please close and
relaunch your browser to try again."

      'Response code 254 - Invalid CGI.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "254"
         LHMResult.Text = "<br><b><font color=""red"">Session logout failed:</font></b> The
request was missing an essential parameter. Sorry for the inconvenience. Please close and
relaunch your browser to try again."

      'Response code 255 - Internal Error.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "255"
         LHMResult.Text = "<br><b><font color=""red"">Session logout failed:</font></b> The
request failed due to an unspecified error. Sorry for the inconvenience. Please close and
relaunch your browser to try again."

      End If

      'Close the streams
      dataStream.Close()
      snwlReply.Close()

      'If there is some asp.net error trying to talk to the SonicWall, print it in the same
color as the background.
      Catch ex as Exception
         catchError.Text = "<font color=""9CBACE"">" & ex.ToString & "</font>"
         LHMResult.Text = "<br><b><font color=""red"">Session logout failed:</font></b> The
request failed due to an unspecified error. Sorry for the inconvenience. Please close and
relaunch your browser to try again. If the problem persists, please notify an attendant."
      End Try
End Sub

</script>
<STYLE>
body {
   font-size: 10pt;
   font-family: verdana,helvetica,arial,sans-serif;
   color:#000000;
   background-color:#9CBACE;
}

tr.heading {
   font-size: 10pt;
   background-color:#006699;
}

tr.smalltext {
   font-size: 8pt;
}

.button {
   border: 1px solid #000000;
```

```
   background-color: #ffffff;
   font-size: 8pt;
}
</STYLE>

<HTML>
<HEAD>
<TITLE>LHM Logout Page</TITLE>

<SCRIPT LANGUAGE="Javascript">

//'Javascript Seconds Countdown Timer
var SecondsToCountDown = <%= sessTimer%>;
var originalTime=" ";

function CountDown()
{
   clockStr="";

   dayStr=Math.floor(SecondsToCountDown/86400)%100000
   if(dayStr>0){
      if(dayStr>1){
         dayStr+=" days ";
      } else dayStr+=" day ";
      clockStr=dayStr;
   }
   hourStr=Math.floor(SecondsToCountDown/3600)%24
   if(hourStr>0){
      if(hourStr>1){
         hourStr+=" hours ";
      } else hourStr+=" hour ";
      clockStr+=hourStr;
   }
   minuteStr=Math.floor(SecondsToCountDown/60)%60
   if(minuteStr>0){
      if(minuteStr>1){
         minuteStr+=" minutes ";
      } else minuteStr+=" minute ";
      clockStr+=minuteStr;
   }
   secondStr=Math.floor(SecondsToCountDown/1)%60
   if(secondStr>0){
      if(secondStr>1){
         secondStr+=" seconds ";
      } else secondStr+=" second ";
      clockStr+=secondStr;
   }

   if(SecondsToCountDown > 0)
   {
      --SecondsToCountDown;
   }

   if(originalTime.length < 2)
   {
      originalTime = clockStr;
   }

   // Make sure the form is still there before trying to set a value
   if(document.frmValidator){
      document.frmValidator.originalTime.value = originalTime;
      document.frmValidator.countdown.value = clockStr;
   }

   setTimeout("CountDown()", 1000);
   if(SecondsToCountDown == 0)
   {
```

```
      document.frmValidator.countdown.value = "Session Expired";
   }
}

//'Disable right-click so that the window doesn't get refreshed since the countdown is
clientside.
document.oncontextmenu = disableRightClick;
function disableRightClick()
{
  return false;
}

//'Disable F5 key, too, on IE at least.
function noF5()
{
   var key_f5 = 116;
   if (key_f5==event.keyCode)
   {
      event.keyCode=0;
      return false;
   }
   return false;
}

document.onkeydown=noF5
document.onmousedown=disableRightClick

</SCRIPT>

</HEAD>

<BODY onload='CountDown()'>
<span id="loggedIn" runat="server">
<form id="frmValidator" runat="server">
<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr class="heading">
      <td colspan=2 align="center"> </td>
   </tr>
   <tr class="heading">
      <td colspan=2 align="center"><font color="white"><b>SonicWall LHM Logout
Window</b></font></td>
   </tr>
   <tr class="heading">
      <td colspan=2 align="center"> </td>
   </tr>
   <tr class="smalltext"><td><br></td></tr>
   <tr class="smalltext">
      <td>Original Session Time:</td>
      <td><asp:textbox width=250 id="originalTime" runat="server" /></td>
   </tr>
   <tr class="smalltext">
      <td>Remaining Session Time:</td>
      <td><asp:textbox width=250 id="countdown" runat="server" /></td>
   </tr>
   <tr class="smalltext">
      <td colspan=2><br>You may use this window to manually logout your session at any time, or
you may safely close this window if you prefer to let your session timeout
automatically.</font></td>
   </td>
   <tr>
      <td colspan=2><center><asp:button id="btnSubmit" class="button" text="  Logout  "
onClick="btnSubmit_Click" runat="server" /></center></td>
   </tr>
</table>
</form>
</span>
```

```
<span id="loggedOut" runat="server">
<form id="logout" runat="server">
<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr class="heading">
      <td colspan=2 align="center"> </td>
   </tr>
   <tr class="heading">
      <td colspan=2 align="center"><font color="white"><b>SonicWall LHM Logout
Window</b></font></td>
   </tr>
   <tr class="heading">
      <td colspan=2 align="center"> </td>
   </tr>
   <tr>
      <td><asp:Label id=LHMResult runat="server" /></td>
   </tr>
   <tr>
      <td><asp:Label id=catchError runat="server" /></td>
   </tr>
   <tr><td><br></td></tr>
   <tr>
      <td><center><asp:button id="Button1" class="button" text="  Close  " runat="server"
/></center></td>
   </tr>
</table>
</form>
</span>

</BODY>
</HTML>
```

myvars.aspx

```
<script language="VB" runat="server">

'Set the logoutPopup window flag - 0 = no popup, 1 = popup
'The use of the logoutPopup in this script is discouraged because the login event is exclusive.
Dim logoutPopup as String = "0"

'Set the debug flag (0 = off, 1 = on)
Dim debugFlag as String = "0"

'Set the path and file for the PDT responder script - this should be the same path as the LHM
settings
'configured on the SonicWall "External Web Server Settings" page, but pointing to the PDT
handler script.
'Refer to http://www.paypal.com/cgi-bin/webscr?cmd=p/xcl/rec/pdt-techview-outside for
information on PDT
Dim pdtPath as String = "http://10.50.165.2/lhm/paypal/pdt.aspx"

'Set the path the PayPal processing CGI. Use the sandbox (https://developer.paypal.com) and
(https://www.sandbox.paypal.com) for testing
'Using the sandbox requires a developer network account and login.
Dim paypalCGI as String = "https://www.sandbox.paypal.com/cgi-bin/webscr"
'Dim paypalCGI as String = "https://www.paypal.com/cgi-bin/webscr"

'Set the email address of the PayPal merchant account to which payment will be made
'The following is a valid sandbox account, but requires authentication by the parent (real)
account.
'You must replace this with you own (real or sandbox account) for use.
Dim myBusiness as String = "lhmdemo@SonicWall.com"

'Set this to token from PayPal account. It must be your actual, valid token.
'Refer to http://paypaltech.com/PDTGen/PDTtokenhelp.htm for information on the identity token
'The following is a valid sandbox token, but requires authentication by the parent (real)
account.
```

```
'You must replace this with you own (real or sandbox token) for use.
Dim token as String = "ucistq6vmKGWPxwJbrTJFDhFq889RxYt_6Mkz_3viraSzjiQJ5iPYCZ5Mdq"

'Set the names for the purchase item options (e.g. 1 hour Access, 3 hours access, etc.)
Dim itemName1 as String = "1 Hour Secure Internet Access"
Dim itemName2 as String = "24 Hours Secure Internet Access"

'Set the paypal querystring number for purchase item options (e.g. 1hour, 60mins, itemone,
etc.)
Dim itemNumber1 as String = "1hour"
Dim itemNumber2 as String = "24hour"

'Set the purchase item options session and idle timers (timers use the same value since we do
not want sessions idling out)
Dim itemTimer1 as String = "3600"'One hour, in minutes
Dim itemTimer2 as String = "86400"'24 hours

'Set the costs in dollars for purchase item options (e.g. one penny = 0.01, one dollar = 1.00,
etc.)
Dim itemAmount1 as String = "0.01"
Dim itemAmount2 as String = "0.02"

'Set the button names and descriptions for purchase item options
Dim itemButton1 as String = "1 Hour Access - $0.01"
Dim itemButton2 as String = "24 Hours Access - $0.02"

'Set the secret for use with optional HMAC auth, as configured in the Extern Guest Auth config
on the SonicWall
Dim strHmac as String = "password"

'Set the digest method for the HMAC, either MD5 or SHA1
Dim hmacType as String = "MD5"
'Dim hmacType as String = "SHA1"

'Set the logo image to use
Dim logo as String = "SonicWall.gif"
'------------------------End of Configurable Settings--------------------------

</script>
```

pdt.aspx

```
<%@ Page Language="VB" Debug="true" %>
<%@ Import Namespace="System" %>
<%@ Import Namespace="System.Net" %>
<%@ Import Namespace="System.IO" %>
<%@ Import Namespace="System.Xml" %>
<%@ Import Namespace="System.Text" %>
<%@ Import Namespace="System.Security" %>
<%@ Import namespace="System.Security.Cryptography.X509Certificates"%>

<!-- #INCLUDE file="myvars.aspx" -->

<script language="VB" runat="server">

'This class allows SSL certs signed by unknown CAs to be accepted.
'This is necessary for the POST to the SonicWall authorizing the LHM session.
Public Class acceptAllCerts
   Implements System.Net.ICertificatePolicy
   Public Function CheckValidationResult(ByVal srvPoint As ServicePoint, _
   ByVal cert As X509Certificate, ByVal request As WebRequest, ByVal problem As Integer) _
      As Boolean Implements ICertificatePolicy.CheckValidationResult
      Return True
   End Function
End Class
```

```
'Sample LHM redirect querystring:
'http://127.0.0.1/lhm/paypal/default.aspx?sessionId=0b712fd83b9f5313db5af1cea6b1004f&ip=10.50.
165.231&mac=00:0e:35:bd:c9:37&ufi=0006b11184300&mgmtBaseUrl=https://10.50.165.193:4043/&client
RedirectUrl=https://10.50.165.193:444/&req=http%3A//www.google.com/ig

Dim ip as String
Dim sessionId as String
Dim mac as String
Dim ufi as String
Dim mgmtBaseUrl as String
Dim clientRedirectUrl as String
Dim req as String
Dim sessTimer as String
Dim idleTimer as String
Dim userName as String
Dim hmac as String
Dim firstname, lastName, itemName, mcGross, mcCurrency, itemNumber, business, txn, payStatus As
String

Sub Page_Load(ByVal sender As System.Object, ByVal e As System.EventArgs) Handles MyBase.Load

   'Use the override class to accept untrusted certificates from the SonicWall
   'This is necessary for the POST to the SonicWall authorizing the LHM session.
   System.Net.ServicePointManager.CertificatePolicy = New acceptAllCerts

   Dim tx, PDTvalidateQuery As String
   Dim strResponse As HttpWebResponse
   Dim temp As String
   Dim PDTArray() As String
   Dim iParts, sResults(0, 0), aParts(), sParts(), sKey, sValue, snwlCustom As String
   Dim i As Integer

   'Set tx to value of tx passed in via Querystring from PayPal
   tx = Request.QueryString("tx")

   'Set string = to the cmd value, tx and at that needs to be
   'POSTed back to PayPal to validate the PDT
   PDTvalidateQuery = "cmd=_notify-synch&tx=" & tx & "&at=" & token

   'Now we need to POST this info back to PayPal for validation of the PDT
   'Create the request back
   Dim req As HttpWebRequest = CType(WebRequest.Create(paypalCGI), HttpWebRequest)

   'Set values for the request back
   'set method
    req.Method = "POST"
    'set content type
    req.ContentType = "application/x-www-form-urlencoded"
    'set length
    req.ContentLength = PDTvalidateQuery.Length

   'Write the request back to PayPal
   Dim stOut As StreamWriter = New StreamWriter(req.GetRequestStream(), Encoding.ASCII)
   stOut.Write(PDTvalidateQuery)
   stOut.Close()

   Try
      strResponse = CType(req.GetResponse(), HttpWebResponse)
   Catch ex As SystemException
      catchError.Text = "<font color=""9CBACE"">" & ex.ToString & "</font>"
   End Try

        'Once we write the stream back to PayPal, we need to read the response.

        Dim IPNResponseStream As Stream = strResponse.GetResponseStream
        Dim encode As Encoding = System.Text.Encoding.GetEncoding("utf-8")
        Dim readStream As New StreamReader(IPNResponseStream, encode)
```

```
  'Read the response in String variable "temp"
  temp = readStream.ReadToEnd

'Debug flag, set in myvars - prints the whole output from the POST reply
If debugFlag = "1" Then
   OutputEntirePDTString(temp)
End If

  'Check to see if the 1st line of the response was "SUCCESS"
  If Mid(temp, 1, 7) = "SUCCESS" Then

      'if it is SUCCESS, the code below puts the response in a nice array
      temp = Mid(temp, 9)
      sParts = Split(temp, vbLf)
      iParts = UBound(sParts) - 1
      ReDim sResults(iParts, 1)

      For i = 0 To iParts

          aParts = Split(sParts(i), "=")
          sKey = aParts(0)
          sValue = aParts(1)
          sResults(i, 0) = sKey
          sResults(i, 1) = sValue

          'You can add more case statements here for other returned variables

      Try
          Select Case sKey
              Case "first_name"
                  firstname = Server.URLDecode(sValue)
              Case "last_name"
                  lastName = Server.URLDecode(sValue)
              Case "item_name"
                  itemName = Server.URLDecode(sValue)
              Case "mc_gross"
                  mcGross = sValue
              Case "mc_currency"
                  mcCurrency = sValue
         Case "item_number"
            itemNumber = Server.URLDecode(sValue)
         Case "business"
            business = Server.URLDecode(sValue)
         Case "txn_id"
            txn = sValue
         Case "payment_status"
            payStatus = sValue
              Case "custom"
            snwlCustom = sValue
            sessionID = snwlCustom.SubString(0,32)
            mgmtBaseUrl=(Server.URLDecode(Mid(snwlCustom,33)))
          End Select
      Catch ex as Exception
         catchError.Text = "<font color=""9CBACE"">" & ex.ToString & "</font>"
      End Try

         Next

      If payStatus = "Completed" Then
         'Transaction Succeeded - Give the Guest a receipt
         Dim receipt as String

         receipt =  "<h3>Transaction Succeeded. Thank you for selecting SonicWall
LHM.</h3><br>"
         receipt + = "<b>Transaction Invoice:</b><br><br>"
         receipt + = "Name: " & firstname & " " & lastName & "<br>"
```

```
            receipt + = "Description: " & itemName & "<br>"
            receipt + = "Amount: " & mcCurrency & " " & mcGross &"<br>"
            receipt + = "Paid to: " & business &"<br>"
            receipt + = "Transaction ID: " & txn &"<br>"
            receipt + = "<br><br>"

            paypalResult.Text = receipt

            LHMResult.Text = "Authorizing your LHM session."

            'Setup the LHM session variables and call LHM Routine
            'Set the session and idle timers to match the variables set in myvars
            If itemNumber = itemNumber1 Then
               sessTimer=itemTimer1
               idleTimer=itemTimer1
            Else
               sessTimer=itemTimer2
               idleTimer=itemTimer2
            End If

            userName = firstname & " " & lastName

            LHM()
         Else
            'The transaction itself was a success, but the payment status was not Completed.
            paypalResult.Text = "The transaction succeeded, but the payment was not completed.
The session cannot be authorized at this time."
         End If

        Else
            ' If PDT response is not "SUCCESS"
          paypalResult.Text = "The PayPal transaction did not succeed. The returned status is:
<b>" &  temp & "</b>"
        End If

        'Close the streams
        readStream.Close()
        strResponse.Close()

    End Sub

   'This is the parser for the debug function to print the entire resonse to the PDT POST
    Private Function OutputEntirePDTString(ByVal myPDTString As String) As String
        Dim tempString() As String = Split(myPDTString, vbLf)
        Dim x As Integer
        For x = 0 To tempString.GetUpperBound(0)
            Response.Write(tempString(x) & "<br>")
        Next
    End Function

Sub LHM()

   'Let the user know that we are setting up the session, just in case it takes more than a
second
   LHMResult.Text = "Authorizing session. Please wait."

   'The LHM cgi on the SonicWall - this does not change
   Dim loginCgi as String = "externalGuestLogin.cgi"

   'Assemble the data to post back to the SonicWall to authorize the LHM session
   Dim loginParams as String = "sessId=" & sessionId & "&userName=" &
Server.URLEncode(userName) & "&sessionLifetime=" & sessTimer & "&idleTimeout=" & idleTimer

   'Combine mgmtBaseUrl from the original redirect with the login cgi
   Dim postToSNWL as String = mgmtBaseUrl & loginCgi

   'Convert the loginParams to a well behaved byte array
```

```
   Dim byteArray As Byte() = Encoding.UTF8.GetBytes(loginParams)

   Try
      'Create the webrequest to the SonicWall
      Dim toSNWL as WebRequest = WebRequest.Create(postToSNWL)

      'Calculate the length of the byte array
      toSNWL.ContentLength = byteArray.Length

      'Set the method for the webrequest to POST
      toSNWL.Method = "POST"

      'Set the content type
      toSNWL.ContentType = "application/x-www-form-urlencoded"

      'Open the request stream
      Dim dataStream As Stream = toSNWL.GetRequestStream()

      'Write the byte array to the request stream
      dataStream.Write(byteArray, 0, byteArray.Length)

      'Close the Stream object
      dataStream.Close()

      'Get the response
      Dim snwlReply As WebResponse = toSNWL.GetResponse()

      'Display the status - looking for 200 = OK.
      'Response.Write(CType(snwlReply, HttpWebResponse).StatusCode)

      'Grab the response and stuff it into an xml doc for possible review
      Dim snwlResponse as XmlDocument = New XmlDocument()
      snwlResponse.Load(snwlReply.GetResponseStream())

      'Set the xPath to the SNWL reply, and get the response
      Dim codePath as String = "SonicWallAccessGatewayParam/AuthenticationReply/ResponseCode"

      'Response.Write(snwlResponse.SelectSingleNode(codePath).InnerXml)

      'Response code 50 - Login Succeeded

      If snwlResponse.SelectSingleNode(codePath).InnerXml = "50"

         'Do we want to provide a logout popup window?
         If logoutPopup = "1" Then
            'Popup hack using Javascript for logout window
            Dim sb As New System.Text.StringBuilder()
            sb.Append("<script language='javascript'>")
            sb.Append("window.open('logout.aspx?sessId=")
            sb.Append(Server.URLEncode(CStr(sessionId)))
            sb.Append("&mgmtBaseUrl=")
            sb.Append(Server.URLEncode(CStr(mgmtBaseUrl)))
            sb.Append("&sessTimer=")
            sb.Append(Server.URLEncode(CStr(sessTimer)))
            sb.Append("','logOut','toolbar=no,")
            sb.Append("addressbar=no,menubar=no,")
            sb.Append("width=400,height=250');")
            sb.Append("<")
            sb.Append("/")
            sb.Append("script>")
            RegisterStartupScript("stp", sb.ToString)
         End If

         LHMResult.Text = "<br><b><font color=""green"">Session authorized:</font></b> You may
now begin your secure Internet access session."

      'Response code 51 - Session Limit Exceeded
```

```
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "51"
         LHMResult.Text = "<br><b><font color=""red"">Session Limit Reached:</font></b> The
maximum number of guest session has been reached. Sorry for the inconvenience. Please close and
relaunch your browser to try again."

      'Response code 100 - Login Failed.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "100"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> Your
session cannot be created at this time. Sorry for the inconvenience. Please close and relaunch
your browser to try again."

      'Response code 251 - Bad HMAC.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "251"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization failed message authentication. Sorry for the inconvenience. Please
close and relaunch your browser to try again."

      'Response code 253 - Invalid SessionID.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "253"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization failed to match a known session identity. Sorry for the
inconvenience. Please close and relaunch your browser to try again."

      'Response code 254 - Invalid CGI.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "254"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization was missing an essential parameter. Sorry for the inconvenience.
Please close and relaunch your browser to try again."

      'Response code 255 - Internal Error.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "255"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization failed due to an unspecified error. Sorry for the inconvenience.
Please close and relaunch your browser to try again."

      End If

      'Close the streams
      dataStream.Close()
      snwlReply.Close()

      'If there is some asp.net error trying to talk to the SonicWall, print it in the same
color as the background.
      Catch ex as Exception
         catchError.Text = "<font color=""9CBACE"">" & ex.ToString & "</font>"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization failed due to an unspecified error. Sorry for the inconvenience.
Please close and relaunch your browser to try again. If the problem persists, please notify an
attendant."
      End Try
End Sub

</script>

<STYLE>
body {
  font-size: 10pt;
  font-family: verdana,helvetica,arial,sans-serif;
  color:#000000;
  background-color:#9CBACE;
}

tr.heading {
  background-color:#006699;
}

.button {
```

```
   border: 1px solid #000000;
   background-color: #ffffff;
}
</STYLE>

<HTML>
<HEAD>
<TITLE>LHM PayPal Script</TITLE>
</HEAD>

<BODY>

<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr class="heading">
      <td colspan=3 align="center"><font color="white"> </td>
   </tr>
   <tr class="heading">
      <td width="50%" valign="center"><font color="white"><b>LHM Access with PayPal Buy
Now</b></font></td>
      <td><center><img width="216" height="51" src="<%= logo%>"></center></td>
      <td width="50%" align="right" valign="center"><font color="white"><b>Powered by
SonicWall LHM</b> </font></td>
   </tr>
   <tr class="heading">
      <td colspan=3 align="center"><font color="white"> </td>
   </tr>
</table>

<table width="100%" border="0" cellpadding="2" cellspacing="0">

   <tr>
   <td><br></td>
   </tr>
   <tr>
      <td><asp:Label id=paypalResult runat="server" /></td>
   </tr>
   <tr>
      <td><asp:Label id=LHMResult runat="server" /></td>
   </tr>
   <tr>
      <td><asp:Label id=catchError runat="server" /></td>
   </tr>
</table>
</BODY>
</HTML>
```

Random Script

Authentication Model	The Guest Client enters an algorithmically validated, randomly generated passcode.
Purpose	Traditional passcode authentication requires that a passcode be generated prior to use and stored on the authenticating platform. For example, Wireless Guest Services requires that accounts be generated on the particular SonicWall Security Appliance on which they are used. The Random script eliminates this dependency by using a salted algorithm to generate and validate passcodes. This means that passcodes never have to be stored anywhere, and as long as the salt is the same, passcodes are completely migratory (that is, they can be used at any site, even against different LHM servers).

The practical implication of this is that guest account passcodes can be generated in bulk, distributed, and used at any time in the future. For example, passcodes could be generated (using a particular salt), printed (for example, on certificates, business cards, scratch cards) distributed, and used at any site whose LHM server employs the same algorithmic salt. The passcodes could be given an absolute (rather than relative) expiration date, at which time the salt can be changed to invalidate the expired passcodes.

The same way that a common salt can be used to validate a set of passcodes across multiple sites, unique salts can ensure that passcodes generated at one site cannot be used at another with a dissimilar salt; so although a common algorithm is used to generate and validate all passcodes, the addition of the salt to the hash function provides uniqueness as needed.

In addition to the `default.aspx` script is a `generator.aspx` script, which is where passcodes are generated. Anywhere from 1 to 999 passcodes may be generated at one time. After generation, individual passcodes can be printed or the entire list can be exported to a `.csv` file.

Support was included for two classes of passcodes: 1 hour and 24 hour. Either type of passcode can be generated by the generator script.

How the generation algorithm works:

1 Generate a random code (root-passcode) of `randChars` (integer with a default value of six) characters, as defined in `myvars`. The character set for the random code generator can be modified within the `default.aspx` file.

2 The salt (defined in `myvars` as the salt string) is prefixed to the root-passcode.

3 A SHA1 hash is then calculated on the resulting string. Three pairs of characters are then obtained from the hash; for a:
- 1-hour passcode, the 408 pair are obtained (characters 4,5 + 0,1 + 8,9).
- 24-hour passcode, the 752 pair are obtained (characters 7,8 + 5,6 + 2,3).

4 The six characters chosen from the hash are then concatenated to root-passcode.

5 The result is the distributable passcode.

The validation algorithm works in reverse:

1 Guest client enters their passcode (call this `enteredCode`).

2 The script grabs the first `randChars` characters of the entered code (call this root-passcode).

3 The salt is prefixed to the root-passcode, and a SHA1 hash is calculated. The 408 pair of characters are obtained and attached to the root-passcode. The 408 pair is then matched to the `enteredCode`:
- If the 408 pair matches, then it is validated as a 1-hour passcode.
- If the 408 pair did not match, then the 752 pair is tried. If this matches the `enteredCode`, then it is validated as a 24-hour passcode.
- If neither matches, then the code is not valid.

After the `enteredCode` has been validated, the `usedcodes.mdb` database is queried to see if the code has already been used. If the `enteredCode` is not found in the database, the LHM session authorization sequence commences, using the MAC address as the userName. After the LHM session is authorized and an acknowledgment has been received by the LHM server, the root-passcode from the `enteredCode` is written to the `usedcodes.mdb` database so that it cannot be re-used. When (if) the salt is changed, it is advisable to flush the database.

myvars Variables

`logoutPopup` Controls the use of the logout popup window. Set to:
- **0** to disable the popup window.
- **1** to enable the popup window.

	`useDB`	Controls the use of the used passcode database. If `useDB` is: • **0**, then the database is not read from or written to, allowing passcodes to be used repeatedly. • **1**, then used passcodes are written to the database, and new authentication processes check the database to determine whether the passcodes have already been used.
	`randChars`	The number of random characters to include in the root-passcode. The default is six. This results in 12-character passcodes because the hash component always adds an additional six characters.
	`salt`	The salt to use in computing the hash. Be sure to use a good salt to prevent unwanted passcode migration/collisions.
	`sessTimer`	The session timer in seconds.
	`idleTimer`	The idle timer in seconds.
	`strHmac`	The shared secret for the optional HMAC function.
	`hmacType`	The digest type to use if HMAC is in use: **MD5** or **SHA1**.
	`logo`	The name of the logo (image) file to use on page headers.
Session Flow	1 The Guest Client enters their passcode. 2 The passcode is validated using algorithmic validation, described in the **Purpose** section above. 3 If the code is validated, it is checked for previous use in the `usedcodes.mdb` database. 4 If it is not present, the LHM session (either 1-hour or 24-hours) is initiated, using the MAC address as the username. 5 After the LHM session is initiated, the script writes the root-passcode to the `usedcodes.mdb` database so that it cannot be reused. 6 The script performs the LHM post to the SonicWall Security Appliance to authorize the session.	
Additional Considerations	Because the script is writing to the database, it is necessary to configure write privileges for the **IUSR_MACHINENAME** and **IWAM_MACHINENAME** (or **ASPNET**) accounts, as described in the I want to use the sample scripts SonicWall provided. What do I need to do to use them? on page 824 The `generator.aspx` script should be located in a secure (publicly inaccessible) area on the web-server.	

Topics:

- default.aspx on page 911
- generator.aspx on page 919
- logout.aspx on page 923
- myvars.aspx on page 928
- print.aspx on page 928

default.aspx

```
<%@ Page Language="VB" Debug="true" %>
<%@ Import Namespace="System" %>
<%@ Import Namespace="System.Net" %>
<%@ Import Namespace="System.IO" %>
<%@ Import Namespace="System.Xml" %>
<%@ Import Namespace="System.Text" %>
<%@ Import Namespace="System.Data" %>
```

```
<%@ Import Namespace="System.Data.OleDB" %>
<%@ Import Namespace="System.Security" %>
<%@ Import Namespace="System.Security.Cryptography" %>
<%@ Import namespace="System.Security.Cryptography.X509Certificates"%>

<!-- #INCLUDE file="myvars.aspx" -->

<script language="VB" runat="server">

'This class allows SSL certs signed by unknown CAs to be accepted.
'This is necessary for the POST to the SonicWall authorizing the LHM session.
Public Class acceptAllCerts
   Implements System.Net.ICertificatePolicy
   Public Function CheckValidationResult(ByVal srvPoint As ServicePoint, _
   ByVal cert As X509Certificate, ByVal request As WebRequest, ByVal problem As Integer) _
      As Boolean Implements ICertificatePolicy.CheckValidationResult
      Return True
   End Function
End Class

'Sample LHM redirect querystring:
'http://127.0.0.1/lhm/random/default.aspx?sessionId=0b712fd83b9f5313db5af1cea6b1004f&ip=10.50.
165.231&mac=00:0e:35:bd:c9:37&ufi=0006b11184300&mgmtBaseUrl=https://10.50.165.193:4043/&client
RedirectUrl=https://10.50.165.193:444/&req=http%3A//www.google.com/ig

Dim ip as String
Dim sessionId as String
Dim mac as String
Dim ufi as String
Dim mgmtBaseUrl as String
Dim clientRedirectUrl as String
Dim req as String
Dim hmac as String
Dim customCode as String

Dim passCode as String
Dim grabCode as String

Sub Page_Load(Source as Object, E as EventArgs)

   LHMResult.Text=""
   catchError.Text=""
   authResult.Text=""

   ip=Request.QueryString("ip")
   sessionId=Request.QueryString("sessionId")
   mac=Request.QueryString("mac")
   ufi=Request.QueryString("ufi")
   mgmtBaseUrl=Request.QueryString("mgmtBaseUrl")
   clientRedirectUrl=Request.QueryString("clientRedirectUrl")
   req=Request.QueryString("req")
   hmac=Request.QueryString("hmac")
   customCode=Request.QueryString("cc")

   'customCode grabs the "cc=" querystring value sent by the SonicWall. This allows you to use
the same
   'page (e.g. this page) for the "Session Expiration" (?cc=2), "Idle Timeout" (?cc=3) and "Max
Sessions" (?cc=4) page.
   If customCode <> "" Then
      Select Case customCode
      Case "2"
         LHMResult.Text="<br><H3><font color=""red"">Your LHM session has expired. You may try
to initiate a new session.</font></H3>"
      Case "3"
         LHMResult.Text="<br><H3><font color=""red"">You have exceeded your idle timeout.
Please log back in.</font></H3>"
      Case "4"
```

```
         LHMResult.Text="<br><H3><font color=""red"">The maximum number of sessions has been
reached. Please try again later.</font></H3>"
      End Select
   End If

   'Use the override class in myvars.aspx to accept untrusted certificates from the SonicWall
   'This is necessary for the POST to the SonicWall authorizing the LHM session.
   System.Net.ServicePointManager.CertificatePolicy = New acceptAllCerts

   'Note - the routine below for handling the hmac requires the use of the SonicSSL.dll and
libeay.dll libraries.
   'The DLL must be copied to the IIS server, and the SonicSSL dll must be registered with
"regsvr32 sonicssl.dll"
   If hmac <> "" Then

      'SonicWall URL Encode routine is different from Microsoft - this is the SonicWall method
      req=Replace(req,"%","%25")
      req=Replace(req,":","%3A")
      req=Replace(req," ","%20")
      req=Replace(req,"?","%3F")
      req=Replace(req,"+","%2B")
      req=Replace(req,"&","%26")
      req=Replace(req,"=","%3D")

      Dim strHmacText as String
      Dim objCrypto as Object
      Dim strHmacGenerated
      Dim loginError as String

      'Initialize the Crypto object
      objCrypto = Server.CreateObject("SonicSSL.Crypto")

      'The text to be encoded
      strHmacText = sessionId & ip & mac & ufi & mgmtBaseUrl & clientRedirectUrl & req

      'Calculate the hash with a key strHmac, the return value is a string converted form the
output sha1 binary.
      'The hash algorithm (MD5 or SHA1) is configured in myvars and in the Extern Guest Auth
config on the SonicWall
      If hmacType = "MD5" Then
         strHmacGenerated = objCrypto.hmac_md5(strHmacText, strHmac)
      Else
         strHmacGenerated = objCrypto.hmac_sha1(strHmacText, strHmac)
      End If

      If strHmacGenerated <> hmac Then
         Dim hmacFail as String
         hmacFail = "<font color=""red"">The HMAC failed validation. Please notify an
attendant.</font><br><br>"
         hmacFail+="<font color=""9CBACE"">Received HMAC: " & hmac & "<br>Calculated HMAC: " &
strHmacGenerated & "<br>"
         hmacFail+="Make sure the digest functions on the SonicWall and LHM server match.<br>"
         hmacFail+="Also make sure the shared secret on the SonicWall and myvars match</font>"
         catchError.Text=hmacFail
      End If

   End If

End Sub

sub OnBtnClearClicked (Sender As Object, e As EventArgs)
   enteredCode.Text = ""
   LHMResult.Text=""
   catchError.Text=""
end sub

Sub btnSubmit_Click(Sender As Object, E As EventArgs)
```

```
   'The following subroutine validates client provided passcodes.
   'The first 6 characters (definable in myars) are grabbed.
   'These characters are then run though a SHA1 hash with a salt that is defined in myvars.

   '3 pairs of substrings are then retrieved from the hash.
   'The code is validated if the 3 pairs concatenated to the randChars (defined in myvars)
characters consist of the following:

   'Validating the 4 0 8 pairs (4,5+0,1+8,9 characters) will provide 1 hour of guest access.
   'Validating the 7 5 2 pairs (7,8+5,6+2,3 characters) will provide 24 hours of guest access.

   grabCode = enteredCode.Text.SubString(0,randChars)

   'Manually compute SHA1 on salt+randomCode, and convert result to base64 - gives stranger
output
   Dim sha1 As sha1 = sha1.Create()
   Dim manualHash As Byte() = sha1.ComputeHash(Encoding.UTF8.GetBytes(salt & grabCode))
   Dim hashResult as String = Convert.ToBase64String(manualHash)

   'Alternatively, use forms hash routine - only provides upper case A-Z + 0-9 output.
   'Dim hashResult as String = FormsAuthentication.HashPasswordForStoringInConfigFile(salt &
randomCode,"SHA1")

   'First try to match on 1 hour code
   passCode  = ""
   passCode = grabCode & hashResult.SubString(4, 2)
   passCode = passCode & hashResult.SubString(0, 2)
   passCode = passCode & hashResult.SubString(8, 2)
   If enteredCode.Text = passCode Then
      sessTimer = "3600"
      authResult.Text="<font color=""green""><b>1 hour code validated.</b></font>"

      'Check the used passcode DB if useDB is enabled in myvars.
      If useDB = "1" Then
         wasItUsed()
      End If
   Else
      'Now try to match on 24 hour code
      passCode = ""
      passCode = grabCode & hashResult.SubString(7, 2)
      passCode = passCode & hashResult.SubString(5, 2)
      passCode = passCode & hashResult.SubString(2, 2)
      If enteredCode.Text = passCode Then
         sessTimer = "86400"
         authResult.Text="<font color=""green""><b>24 hour code validated.</b></font>"

         'Check the used passcode DB if useDB is enabled in myvars.
         If useDB = "1" Then
            wasItUsed()
         End If

      Else
         authResult.Text="<font color=""Red""><b>Passcode cannot be validated.</b><br>The
passcode is case-sensitive.<br>Please try again.</font>"
      End if
   End If

End Sub

Sub wasItUsed ()

   'Check to see if the root (randChars) of the passcode is already in the used database.
   Dim strConn as string = "Provider=Microsoft.Jet.OLEDB.4.0;Data Source=" &
server.mappath("usedcodes.mdb") & ";"
   Dim MySQL as string = "SELECT * From passCodes Where passCode = '" & grabCode & "'"
   Dim MyConn as New OleDBConnection (strConn)
```

```
   Dim cmd as New OleDBCommand (MySQL, MyConn)
   Dim objDR As OleDbDataReader
   Dim isUsed As Boolean

   MyConn.Open()
   objDR = cmd.ExecuteReader()
   isUsed = objDR.Read()
   objDR.Close()
   MyConn.Close()

   'If the passcode is not found in the database
   if isUsed = False
      LHM()
   Else
      authResult.Text="<font color=""Red""><b>Passcode has already been used.</b><br>Please
see an attendant for assistance.</font>"
   End If

End Sub

Sub writeToDB ()

   'Try to write the submitted (only randChars characters instead of the whole passcode) info
to the database file
   Try
      Dim strConn as string = "Provider=Microsoft.Jet.OLEDB.4.0;Data Source=" &
server.mappath("usedcodes.mdb") & ";"

      Dim MySQL as string = "INSERT INTO passCodes (passCode) VALUES ('" & grabCode & "')"
      Dim MyConn as New OleDBConnection (strConn)
      Dim cmd as New OleDBCommand (MySQL, MyConn)
      MyConn.Open ()
      cmd.ExecuteNonQuery ()
      MyConn.Close ()

      Catch ex as Exception
         catchError.Text = "<font color=""9CBACE"">" & ex.ToString & "</font>"
   End Try

End Sub

Sub LHM()

   'The writeToDB sub is in the Response code 50 - Login Succeeded routine, after the LHM
exchange succeeds. You may move it to the top to write the passcode to the DB before the LHM
transaction for testing purposes.
   'writeToDB ()

   enteredCode.Text = "Code Accepted."

   'Let the user know that we are setting up the session, just in case it takes more than a
second
   LHMResult.Text = "Authorizing session. Please wait."

   'The LHM cgi on the SonicWall - this does not change
   Dim loginCgi as String = "externalGuestLogin.cgi"

   'Assemble the data to post back to the SonicWall to authorize the LHM session
   Dim loginParams as String = "sessId=" & sessionId & "&userName=" & mac & "&sessionLifetime="
& sessTimer & "&idleTimeout=" & idleTimer

   'Combine mgmtBaseUrl from the original redirect with the login cgi
   Dim postToSNWL as String = mgmtBaseUrl & loginCgi

   'Convert the loginParams to a well behaved byte array
   Dim byteArray As Byte() = Encoding.UTF8.GetBytes(loginParams)
```

```
   Try
      'Create the webrequest to the SonicWall
      Dim toSNWL as WebRequest = WebRequest.Create(postToSNWL)

      'Calculate the length of the byte array
      toSNWL.ContentLength = byteArray.Length

      'Set the method for the webrequest to POST
      toSNWL.Method = "POST"

      'Set the content type
      toSNWL.ContentType = "application/x-www-form-urlencoded"

      'Open the request stream
      Dim dataStream As Stream = toSNWL.GetRequestStream()

      'Write the byte array to the request stream
      dataStream.Write(byteArray, 0, byteArray.Length)

      'Close the Stream object
      dataStream.Close()

      'Get the response
      Dim snwlReply As WebResponse = toSNWL.GetResponse()

      'Display the status - looking for 200 = OK.
      'Response.Write(CType(snwlReply, HttpWebResponse).StatusCode)

      'Grab the response and stuff it into an xml doc for possible review
      Dim snwlResponse as XmlDocument = New XmlDocument()
      snwlResponse.Load(snwlReply.GetResponseStream())

      'Set the xPath to the SNWL reply, and get the response
      Dim codePath as String = "SonicWallAccessGatewayParam/AuthenticationReply/ResponseCode"

      'Response.Write(snwlResponse.SelectSingleNode(codePath).InnerXml)

      'Response code 50 - Login Succeeded

      If snwlResponse.SelectSingleNode(codePath).InnerXml = "50"

         'Do we want to provide a logout popup window?
         If logoutPopup = "1" Then
            'Popup hack using Javascript for logout window
            Dim sb As New System.Text.StringBuilder()
            sb.Append("<script language='javascript'>")
            sb.Append("window.open('logout.aspx?sessId=")
            sb.Append(Server.URLEncode(CStr(sessionId)))
            sb.Append("&mgmtBaseUrl=")
            sb.Append(Server.URLEncode(CStr(mgmtBaseUrl)))
            sb.Append("&sessTimer=")
            sb.Append(Server.URLEncode(CStr(sessTimer)))
            sb.Append("','logOut','toolbar=no,")
            sb.Append("addressbar=no,menubar=no,")
            sb.Append("width=400,height=250');")
            sb.Append("<")
            sb.Append("/")
            sb.Append("script>")
            RegisterStartupScript("stp", sb.ToString)
         End If

         LHMResult.Text = "<br><b><font color=""green"">Session authorized:</font></b> You may
now go to the URL you originally requested: <a target=""_blank"" href=""" & req & """>" & req &
"</a>"

         'Write the passcode the DB if the LHM session succeeds and if useDB = 1.
         If useDB = "1" Then
```

```
            writeToDB ()
         End If

      'Response code 51 - Session Limit Exceeded
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "51"
         LHMResult.Text = "<br><b><font color=""red"">Session Limit Reached:</font></b> The
maximum number of guest session has been reached. Sorry for the inconvenience. Please close and
relaunch your browser to try again."

      'Response code 100 - Login Failed.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "100"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> Your
session cannot be created at this time. Sorry for the inconvenience. Please close and relaunch
your browser to try again."

      'Response code 251 - Bad HMAC.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "251"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization failed message authentication. Sorry for the inconvenience. Please
close and relaunch your browser to try again."

      'Response code 253 - Invalid SessionID.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "253"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization failed to match a known session identity. Sorry for the
inconvenience. Please close and relaunch your browser to try again."

      'Response code 254 - Invalid CGI.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "254"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization was missing an essential parameter. Sorry for the inconvenience.
Please close and relaunch your browser to try again."

      'Response code 255 - Internal Error.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "255"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization failed due to an unspecified error. Sorry for the inconvenience.
Please close and relaunch your browser to try again."

      End If

      'Close the streams
      dataStream.Close()
      snwlReply.Close()

      'If there is some asp.net error trying to talk to the SonicWall, print it in the same
color as the background.
      Catch ex as Exception
         catchError.Text = "<font color=""9CBACE"">" & ex.ToString & "</font>"
         LHMResult.Text = "<br><b><font color=""red"">Session creation failed:</font></b> The
request for authorization failed due to an unspecified error. Sorry for the inconvenience.
Please close and relaunch your browser to try again. If the problem persists, please notify an
attendant."
      End Try
End Sub


</script>

<STYLE>
body {
  font-size: 10pt;
  font-family: verdana,helvetica,arial,sans-serif;
  color:#000000;
  background-color:#9CBACE;
}
```

```
tr.heading {
  background-color:#006699;
}

.button {
   border: 1px solid #000000;
   background-color: #ffffff;
}
</STYLE>

<HTML>
<HEAD>
<TITLE>LHM Random Script</TITLE>
</HEAD>

<BODY>
<form id="frmValidator" onKeyPress="if(event.keyCode==13)
{document.getElementById('btnSubmit').click(); return false}" runat="server">

<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr class="heading">
      <td colspan=3 align="center"><font color="white"> </td>
   </tr>
   <tr class="heading">
      <td width="50%" valign="center"><font color="white"><b>Algorithmic
Authentication</b></font></td>
      <td><center><img width="216" height="51" src="<%= logo %>"></center></td>
      <td width="50%" align="right" valign="center"><font color="white"><b>Powered by
SonicWall LHM</b> </font></td>
   </tr>
   <tr class="heading">
      <td colspan=3 align="center"><font color="white"> </td>
   </tr>
</table>

<table width="90%" border="0" cellpadding="2" cellspacing="0">
   <tr>
      <td><b>Welcome <%= ip%> to SonicWall's LHM Algorithmic Authenticator.</b><br><br>Enter
your unique randomly generated passcode to obtain secure guest Internet access.<br><br>Valid
passcodes are not stored anywhere, so validation is not performed against any kind of database.
Instead, when a passcode is entered, it is algorithmically validated. After a passcode is
successfully used, it is written to a "used passcode" database so that it cannot be
reused.<br><br>The validator will recognize 1 hour and 24 hour passcodes - these
characteristics were encoded within the passcodes themselves during generation.<br><br>
      </td>
   </tr>
</table>
<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr class="heading">
      <td colspan=3 align="center"><font color="white"> </td>
   </tr>
</table>
<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr><br>
      <td width="30%">Enter your passcode:</td>
      <td width="30%"><asp:TextBox id="enteredCode" runat="server" /></td>
      <td width="40%"><asp:RequiredFieldValidator id="valEnteredCode"
ControlToValidate="enteredCode" ErrorMessage="Please enter your passcode." runat="server"
/></td>
   </tr>
   <tr>
      <td></td><td colspan=2><asp:Label id=authResult runat="server" /> </td>
   </tr>
   <tr>
      <td></td>
      <td><asp:button id="btnSubmit" class="button" text="  Submit  "
onClick="btnSubmit_Click" runat="server" />
```

```
        
      <asp:button id="btnClear" class="button" text=" Clear " CausesValidation="False"
onClick="OnBtnClearClicked" runat="server" />
      </td>
   </tr>
   <tr>
      <td colspan=2><asp:Label id=LHMResult runat="server" /></td>
   </tr>
   <tr>
      <td colspan=2><asp:Label id=catchError runat="server" /></td>
   </tr>
</table>
</form>
</BODY>
</HTML>
```

generator.aspx

```
<%@ Page Language="VB" Debug="true" %>
<%@ Import Namespace="System" %>
<%@ Import Namespace="System.Net" %>
<%@ Import Namespace="System.IO" %>
<%@ Import Namespace="System.Xml" %>
<%@ Import Namespace="System.Text" %>
<%@ Import Namespace="System.Data" %>
<%@ Import Namespace="System.Data.OleDB" %>
<%@ Import Namespace="System.Security" %>
<%@ Import Namespace="System.Security.Cryptography" %>
<%@ Import namespace="System.Security.Cryptography.X509Certificates"%>

<!-- #INCLUDE file="myvars.aspx" -->

<script language="VB" runat="server">

Dim genCodes As New ArrayList()
Dim codeType As String

Sub Page_Load(Source as Object, E as EventArgs)
   If Not isPostBack Then
      Heading.Text=" "
      btnExport.Visible = False
   End If
End Sub

Sub btnSubmit_Click(Sender As Object, E As EventArgs)
   'The following generates passcodes beginning with a random character generator.
   'The number of characters in randomCode is configurable in myvars.
   'The randomCode output is then run though a SHA1 hash with a salt that is defined in myvars.
   'Note: If you are using this in a live environment, it is important to change the salt to
prevent algorithm compromise.

   '3 pairs of substrings are then retrieved from the hash, and concatenated to the randomCode
to form the passcode.

   'In the current sample implementation:
   'The 4 0 8 pairs (4,5+0,1+8,9 characters) from the hash will provide 1 hour of guest access.
   'The 7 5 2 pairs (7,8+5,6+2,3 characters) from the hash will provide 24 hours of guest
access.

   Dim myLooper As Integer
   Dim passCode as String

      For myLooper = 1 to Convert.ToInt32(codeCount.Text)

         Dim x As Integer = 0
         Dim isitRand as boolean = False
```

```
         Dim intRand as Integer = 0
         Dim randomCode as String = ""

         For x = 1 to randChars
            Do Until isItRand = True
               '48 to 57 for numbers, 65 to 90 for uppercase, 97 to 122 for lowercase
               intRand = Int((122 - 48 + 1) * Rnd + 48)
               'Select the legal characterset for randomCode by including legal characters
below.
               If InStr(1, "abcdefgh jk mn pqrstuvwxyzABCDEFGH JKLMN PQRSTUVWXYZ 23456789
",Chr(intRand), 1) Then
                  isItRand = True
               End If
            Loop
            randomCode = randomCode & Chr(intRand)
            isItRand = False
         Next

         'Manually compute SHA1 on salt+randomCode, and convert result to base64 - gives
stranger output
         Dim sha1 As sha1 = sha1.Create()
         Dim manualHash As Byte() = sha1.ComputeHash(Encoding.UTF8.GetBytes(salt &
randomCode))
         Dim hashResult as String = Convert.ToBase64String(manualHash)

         'Alternatively, use forms hash routine - only provides upper case A-Z + 0-9 output.
         'Dim hashResult as String =
FormsAuthentication.HashPasswordForStoringInConfigFile(salt & randomCode,"SHA1")

         If dropDownList1.SelectedItem.Value = "1 Hour" Then
            passCode = randomCode & hashResult.SubString(4, 2)
            passCode = passCode & hashResult.SubString(0, 2)
            passCode = passCode & hashResult.SubString(8, 2)
            genCodes.Add(passCode)
         Else
            passCode = randomCode & hashResult.SubString(7, 2)
            passCode = passCode & hashResult.SubString(5, 2)
            passCode = passCode & hashResult.SubString(2, 2)
            genCodes.Add(passCode)
         End If

      Next

      btnExport.Visible = True
      heading.Text = "Your " & codeCount.Text & " <b>" & dropDownList1.SelectedItem.Value &
"</b> Passcodes:"
      genOutput.DataSource = genCodes
      genOutput.DataBind()
      codeCount.Text=""

      'Store the genCodes array in session state for retrieval for printing and exporting
      Session("myGenCodes") = genCodes
      Session("codeType") = dropDownList1.SelectedItem.Value

   End Sub

   Sub printIt(Src As Object, e As DataListCommandEventArgs)
      If not Session.Item("myGenCodes") is Nothing Then
         genCodes=Session.Item("myGenCodes")
         codeType=Session.Item("codeType")
         'response.write(CStr(genCodes.Item(e.Item.ItemIndex)))

         'Popup hack using Javascript so that individual entries can be printed from the
DataList
         Dim sb As New System.Text.StringBuilder()
         sb.Append("<script language='javascript'>")
         sb.Append("window.open('print.aspx?genCode=")
```

```
            sb.Append(Server.URLEncode(CStr(genCodes.Item(e.Item.ItemIndex))))
            sb.Append("&sessLife=")
            sb.Append(Server.URLEncode(codeType))
            sb.Append("','printCode','toolbar=no,")
            sb.Append("addressbar=no,menubar=no,")
            sb.Append("width=400,height=250');")
            sb.Append("<")
            sb.Append("/")
            sb.Append("script>")
            RegisterStartupScript("stp", sb.ToString)
         End If

      End Sub

      Sub exporter(Sender As Object, E As EventArgs)

      If not Session.Item("myGenCodes") is Nothing Then
         genCodes=Session.Item("myGenCodes")

         'Convert the genCodes array to a string with CRs for later conversion to a byte array
         Dim i as Integer
         Dim genCodeString as String
         for i = 0 To genCodes.Count - 1
            genCodeString += CStr(genCodes.Item(i)) & Chr(13)
         Next

         'response.write(genCodeString)

         'Create the byte array and send it to the browser as genCodes.csv
         Dim data() As Byte = System.Text.ASCIIEncoding.ASCII.GetBytes(genCodeString)
         Response.Clear()
         Response.AddHeader("Content-Type", "application/Excel")
         Response.AddHeader("Content-Disposition", "inline;filename=genCodes.csv")
         Response.BinaryWrite(data)
         Response.End()
      End If

End Sub

</script>

<STYLE>
body {
  font-size: 10pt;
  font-family: verdana,helvetica,arial,sans-serif;
  color:#000000;
  background-color:#9CBACE;
}

tr.heading {
  background-color:#006699;
}

.button {
   border: 1px solid #000000;
   background-color: #ffffff;
}
</STYLE>

<HTML>
<HEAD>
<TITLE>LHM Random Script</TITLE>
</HEAD>

<BODY>
<form id="frmValidator" runat="server">
```

```
<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr class="heading">
      <td colspan=3 align="center"><font color="white"> </td>
   </tr>
   <tr class="heading">
      <td width="50%" valign="center"><font color="white"><b>Algorithmic
Authentication</b></font></td>
      <td><center><img width="216" height="51" src="<%= logo %>"></center></td>
      <td width="50%" align="right" valign="center"><font color="white"><b>Passcode
Generator</b> </font></td>
   </tr>
   <tr class="heading">
      <td colspan=3 align="center"><font color="white"> </td>
   </tr>
</table>

<table width="90%" border="0" cellpadding="2" cellspacing="0">
   <tr>
      <td><b>Welcome to SonicWall's LHM Algorithmic Generator.</b><br><br>This will allow you
to create randomly generated passcodes for secure guest Internet access.<br><br>Valid passcodes
are not stored anywhere, so validation is not performed against any kind of database. Instead,
when a passcode is entered, it is algorithmically validated. Once a passcode is successfully
used, it is written to a "used passcode" database so that it cannot be reused.<br><br>The
validator will recognize 1 hour and 24 hour passcodes - these characteristics were encoded
within the passcodes themselves during generation.<br><br>
      </td>
   </tr>
</table>
<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr class="heading">
      <td colspan=3 align="center"><font color="white"> </td>
   </tr>
</table>
<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr><br>
      <td width="15%">Passcode type:</td>
      <td width="10%"><asp:DropDownList id="DropDownList1" runat="server">
         <asp:ListItem>1 Hour</asp:ListItem>
         <asp:ListItem>24 Hours</asp:ListItem>
      </asp:DropDownList></td>
      <td width="20%">Number to generate:</td>
      <td width="20%"><asp:TextBox id="codeCount" runat="server" /></td>
      <td width="50%"><asp:RequiredFieldValidator id="valcodeCount"
ControlToValidate="codeCount" ErrorMessage="Enter a value." Font-Size="10" Display="Dynamic"
runat="server" />
      <asp:RangeValidator id="Range1" ControlToValidate="codeCount" MinimumValue="1"
MaximumValue="999" Type="Integer" Font-Size="10" ErrorMessage="Values from 1 to 999."
runat="server" /></td>
   </tr>
   <tr>
      <td colspan=3></td>
      <td><asp:button id="btnSubmit" class="button" text=" Submit " onClick="btnSubmit_Click"
runat="server" />  <asp:button id="btnExport" class="button" text=" Export "
CausesValidation="False" onClick="exporter" runat="server" /><br></td>
      <td><br></td>

   </tr>
   <tr><tr class="heading">
      <td colspan=5><font color="white"><asp:Label id=heading runat="server" /></td>
   </tr>
   <tr><td><br></td></tr>
</table>

<asp:DataList id="genOutput" Runat="Server" RepeatColumns="4" RepeatDirection="Horizontal"
CellPadding="0" Cellspacing="0" GridLines="Both" align="center" OnItemCommand="printIt">
  <ItemTemplate>
   <td>
```

```
   <asp:Label Text='<%# Container.DataItem %>' Runat="Server"/>
   </td>
   <td>
   <asp:ImageButton id="print" runat="server" ImageUrl="print.gif" EnableViewState="False"
CausesValidation="False" CommandName='<%# Container.DataItem %>' />
   </td>
  </ItemTemplate>
</asp:DataList>

</form>
</BODY>
</HTML>
```

logout.aspx

```
<%@ Page Language="VB" Debug="true" %>
<%@ Import Namespace="System" %>
<%@ Import Namespace="System.Net" %>
<%@ Import Namespace="System.IO" %>
<%@ Import Namespace="System.Xml" %>
<%@ Import Namespace="System.Text" %>
<%@ Import Namespace="System.Security" %>
<%@ Import namespace="System.Security.Cryptography.X509Certificates"%>

<!-- #INCLUDE file="myvars.aspx" -->

<script language="VB" runat="server">

'This class allows SSL certs signed by unknown CAs to be accepted.
'This is necessary for the POST to the SonicWall authorizing the LHM session.
Public Class acceptAllCerts
   Implements System.Net.ICertificatePolicy
   Public Function CheckValidationResult(ByVal srvPoint As ServicePoint, _
   ByVal cert As X509Certificate, ByVal request As WebRequest, ByVal problem As Integer) _
      As Boolean Implements ICertificatePolicy.CheckValidationResult
      Return True
   End Function
End Class

Dim sessionId as String
Dim mgmtBaseUrl as String
Dim eventId as String = "&eventId=1"

'Grab the code and the session lifetime from the generator page
Sub Page_Load(src as Object, e as EventArgs)
   sessionId=Request.QueryString("sessId")
   mgmtBaseUrl=Request.QueryString("mgmtBaseUrl")
   sessTimer=Request.QueryString("sessTimer")


   'Use the override class in myvars.aspx to accept untrusted certificates from the SonicWall
   'This is necessary for the POST to the SonicWall authorizing the LHM session.
   System.Net.ServicePointManager.CertificatePolicy = New acceptAllCerts

   'When the page loads, make the loggedIn span visible
   loggedIn.Visible=True
   loggedOut.Visible=False

   Me.Button1.Attributes.Add("OnClick", "self.close()")

End Sub

'The Logout button
Sub btnSubmit_Click(Sender As Object, E As EventArgs)
```

```
   'Let the user know that we are setting up the session, just in case it takes more than a
second
   LHMResult.Text = "Authorizing session. Please wait."

   'The LHM cgi on the SonicWall - this does not change
   Dim loginCgi as String = "externalGuestLogoff.cgi"

   'Assemble the data to post back to the SonicWall to authorize the LHM session
   Dim loginParams as String = "sessId=" & sessionId & eventId

   'Combine mgmtBaseUrl from the original redirect with the login cgi
   Dim postToSNWL as String = mgmtBaseUrl & loginCgi

   'Convert the loginParams to a well behaved byte array
   Dim byteArray As Byte() = Encoding.UTF8.GetBytes(loginParams)

   Try
      'Make the loggedOut span visible
      loggedIn.Visible=False
      loggedOut.Visible=True

      'Create the webrequest to the SonicWall
      Dim toSNWL as WebRequest = WebRequest.Create(postToSNWL)

      'Calculate the length of the byte array
      toSNWL.ContentLength = byteArray.Length

      'Set the method for the webrequest to POST
      toSNWL.Method = "POST"

      'Set the content type
      toSNWL.ContentType = "application/x-www-form-urlencoded"

      'Open the request stream
      Dim dataStream As Stream = toSNWL.GetRequestStream()

      'Write the byte array to the request stream
      dataStream.Write(byteArray, 0, byteArray.Length)

      'Close the Stream object
      dataStream.Close()

      'Get the response
      Dim snwlReply As WebResponse = toSNWL.GetResponse()

      'Display the status - looking for 200 = OK.
      'Response.Write(CType(snwlReply, HttpWebResponse).StatusCode)

      'Grab the response and stuff it into an xml doc for possible review
      Dim snwlResponse as XmlDocument = New XmlDocument()
      snwlResponse.Load(snwlReply.GetResponseStream())

      'Set the xPath to the SNWL reply, and get the response
      Dim codePath as String = "SonicWallAccessGatewayParam/LogoffReply/ResponseCode"

      'Response.Write(snwlResponse.SelectSingleNode(codePath).InnerXml)

      'Response code 150 - Logout Succeeded
      If snwlResponse.SelectSingleNode(codePath).InnerXml = "150"
         LHMResult.Text = "<br><b><font color=""green"">Your session has been logged
out.<br><br>Thank you for using LHM Guest Services.</font></b>"

      'Response code 251 - Bad HMAC.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "251"
         LHMResult.Text = "<br><b><font color=""red"">Session logout failed:</font></b> The
request failed message authentication. Sorry for the inconvenience. Please close and relaunch
your browser to try again."
```

```
      'Response code 253 - Invalid SessionID.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "253"
         LHMResult.Text = "<br><b><font color=""red"">Session logout failed:</font></b> The
request failed to match a known session identity. Sorry for the inconvenience. Please close and
relaunch your browser to try again."

      'Response code 254 - Invalid CGI.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "254"
         LHMResult.Text = "<br><b><font color=""red"">Session logout failed:</font></b> The
request was missing an essential parameter. Sorry for the inconvenience. Please close and
relaunch your browser to try again."

      'Response code 255 - Internal Error.
      ElseIf snwlResponse.SelectSingleNode(codePath).InnerXml = "255"
         LHMResult.Text = "<br><b><font color=""red"">Session logout failed:</font></b> The
request failed due to an unspecified error. Sorry for the inconvenience. Please close and
relaunch your browser to try again."

      End If

      'Close the streams
      dataStream.Close()
      snwlReply.Close()

      'If there is some asp.net error trying to talk to the SonicWall, print it in the same
color as the background.
      Catch ex as Exception
         catchError.Text = "<font color=""9CBACE"">" & ex.ToString & "</font>"
         LHMResult.Text = "<br><b><font color=""red"">Session logout failed:</font></b> The
request failed due to an unspecified error. Sorry for the inconvenience. Please close and
relaunch your browser to try again. If the problem persists, please notify an attendant."
      End Try
End Sub

</script>
<STYLE>
body {
   font-size: 10pt;
   font-family: verdana,helvetica,arial,sans-serif;
   color:#000000;
   background-color:#9CBACE;
}

tr.heading {
   font-size: 10pt;
   background-color:#006699;
}

tr.smalltext {
   font-size: 8pt;
}

.button {
   border: 1px solid #000000;
   background-color: #ffffff;
   font-size: 8pt;
}
</STYLE>

<HTML>
<HEAD>
<TITLE>LHM Logout Page</TITLE>

<SCRIPT LANGUAGE="Javascript">

//'Javascript Seconds Countdown Timer
```

```
var SecondsToCountDown = <%= sessTimer%>;
var originalTime=" ";

function CountDown()
{
   clockStr="";

   dayStr=Math.floor(SecondsToCountDown/86400)%100000
   if(dayStr>0){
      if(dayStr>1){
         dayStr+=" days ";
      } else dayStr+=" day ";
      clockStr=dayStr;
   }
   hourStr=Math.floor(SecondsToCountDown/3600)%24
   if(hourStr>0){
      if(hourStr>1){
         hourStr+=" hours ";
      } else hourStr+=" hour ";
      clockStr+=hourStr;
   }
   minuteStr=Math.floor(SecondsToCountDown/60)%60
   if(minuteStr>0){
      if(minuteStr>1){
         minuteStr+=" minutes ";
      } else minuteStr+=" minute ";
      clockStr+=minuteStr;
   }
   secondStr=Math.floor(SecondsToCountDown/1)%60
   if(secondStr>0){
      if(secondStr>1){
         secondStr+=" seconds ";
      } else secondStr+=" second ";
      clockStr+=secondStr;
   }

   if(SecondsToCountDown > 0)
   {
      --SecondsToCountDown;
   }

   if(originalTime.length < 2)
   {
      originalTime = clockStr;
   }

   // Make sure the form is still there before trying to set a value
   if(document.frmValidator){
      document.frmValidator.originalTime.value = originalTime;
      document.frmValidator.countdown.value = clockStr;
   }

   setTimeout("CountDown()", 1000);
   if(SecondsToCountDown == 0)
   {
      document.frmValidator.countdown.value = "Session Expired";
   }
}

//'Disable right-click so that the window doesn't get refreshed since the countdown is
clientside.
document.oncontextmenu = disableRightClick;
function disableRightClick()
{
  return false;
}
```

```
//'Disable F5 key, too, on IE at least.
function noF5()
{
   var key_f5 = 116;
   if (key_f5==event.keyCode)
   {
      event.keyCode=0;
      return false;
   }
   return false;
}

document.onkeydown=noF5
document.onmousedown=disableRightClick

</SCRIPT>

</HEAD>

<BODY onload='CountDown()'>
<span id="loggedIn" runat="server">
<form id="frmValidator" runat="server">
<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr class="heading">
      <td colspan=2 align="center"> </td>
   </tr>
   <tr class="heading">
      <td colspan=2 align="center"><font color="white"><b>SonicWall LHM Logout
Window</b></font></td>
   </tr>
   <tr class="heading">
      <td colspan=2 align="center"> </td>
   </tr>
   <tr class="smalltext"><td><br></td></tr>
   <tr class="smalltext">
      <td>Original Session Time:</td>
      <td><asp:textbox width=250 id="originalTime" runat="server" /></td>
   </tr>
   <tr class="smalltext">
      <td>Remaining Session Time:</td>
      <td><asp:textbox width=250 id="countdown" runat="server" /></td>
   </tr>
   <tr class="smalltext">
      <td colspan=2><br>You may use this window to manually logout your session at any time, or
you may safely close this window if you prefer to let your session timeout
automatically.</font></td>
   </td>
   <tr>
      <td colspan=2><center><asp:button id="btnSubmit" class="button" text="  Logout  "
onClick="btnSubmit_Click" runat="server" /></center></td>
   </tr>
</table>
</form>
</span>

<span id="loggedOut" runat="server">
<form id="logout" runat="server">
<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr class="heading">
      <td colspan=2 align="center"> </td>
   </tr>
   <tr class="heading">
      <td colspan=2 align="center"><font color="white"><b>SonicWall LHM Logout
Window</b></font></td>
   </tr>
   <tr class="heading">
      <td colspan=2 align="center"> </td>
```

```
   </tr>
   <tr>
      <td><asp:Label id=LHMResult runat="server" /></td>
   </tr>
   <tr>
      <td><asp:Label id=catchError runat="server" /></td>
   </tr>
   <tr><td><br></td></tr>
   <tr>
      <td><center><asp:button id="Button1" class="button" text="  Close  " runat="server"
/></center></td>
   </tr>
</table>
</form>
</span>

</BODY>
</HTML>
```

myvars.aspx

```
<script language="VB" runat="server">

'Set the logoutPopup window flag - 0 = no popup, 1 = popup
'The use of the logoutPopup in this script is discouraged because the login event is exclusive.
'The login event can be made non exclusive in this script by setting useDB to 0.
Dim logoutPopup as String = "0"

'Set the use of the database for storing and checking used passcodes. 0 = do not use DB, 1 = use
DB.
Dim useDB as String = "1"

'The number of characters in the randomCode
Dim randChars as Integer = 6

'Set the salt the generation of the SHA1 hash
Dim salt as String = "moosifer"

'The LHM Session Timeout is set by the passcode in this script
Dim sessTimer as String

'Set the LHM Idle Timeout
Dim idleTimer as String = "300"

'Set the secret for use with optional HMAC auth, as configured in the Extern Guest Auth config
on the SonicWall
Dim strHmac as String = "password"

'Set the digest method for the HMAC, either MD5 or SHA1
Dim hmacType as String = "MD5"
'Dim hmacType as String = "SHA1"

'Set the logo image to use
Dim logo as String = "SonicWall.gif"

'-----------------------End of Configurable Settings-------------------------

</script>
```

print.aspx

```
<!-- #INCLUDE file="myvars.aspx" -->

<script language="VB" runat="server">
```

```
Dim genCode as String
Dim sessLife as String

'Grab the code and the session lifetime from the generator page
Sub Page_Load(src as Object, e as EventArgs)
   genCode=Request.QueryString("genCode")
   sessLife=Request.QueryString("sessLife")
End Sub

</script>
<STYLE>
body {
  font-size: 10pt;
  font-family: verdana,helvetica,arial,sans-serif;
  color:#000000;
  background-color:#9CBACE;
}
tr.heading {
  background-color:#006699;
}
</STYLE>
<BODY>
<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr class="heading">
      <td colspan=2 align="center"><font color="white"> </td>
   </tr>
   <tr class="heading">
      <td colspan=2 align="center"><img width="216" height="51" src="<%= logo %>"></td>
   </tr>
   <tr class="heading">
      <td colspan=2 align="center"><font color="white"> </td>
   </tr>
   <tr><td><br><br></td></tr>
   <tr>
      <td>Your Pass Code is:</td>
      <td><b><%= genCode%></b></td>
   </tr>
   <tr><td><br></td></tr>
   <tr>
      <td>Session Lifetime is:</td>
      <td><b><%= sessLife%></b></td>
   </tr>
</table>

<script language='javascript'>window.print();</script>

</BODY>
</HTML>
```

Chooser.aspx Script

```
<%@ Page Language="VB" Debug="true" %>
<%@ Import Namespace="System" %>
<%@ Import Namespace="System.Net" %>
<%@ Import Namespace="System.IO" %>
<%@ Import Namespace="System.Xml" %>
<%@ Import Namespace="System.Text" %>
<%@ Import Namespace="System.Security" %>
<%@ Import namespace="System.Security.Cryptography.X509Certificates"%>

<script language="VB" runat="server">

Dim ip as String
Dim sessionId as String
Dim mac as String
Dim ufi as String
```

```
Dim mgmtBaseUrl as String
Dim clientRedirectUrl as String
Dim req as String
Dim hmac as String
Dim customCode as String
Dim qString as String

Sub Page_Load(src as Object, e as EventArgs)

   'Grab the querystring one element at a time since we need to do a custom URL encode on the
req variable
   sessionId=Request.QueryString("sessionId")
   ip=Request.QueryString("ip")
   mac=Request.QueryString("mac")
   ufi=Request.QueryString("ufi")
   mgmtBaseUrl=Request.QueryString("mgmtBaseUrl")
   clientRedirectUrl=Request.QueryString("clientRedirectUrl")
   req=Request.QueryString("req")
   hmac=Request.QueryString("hmac")
   customCode=Request.QueryString("cc")

   'SonicWall URL Encode routine is different from Microsoft - this is the SonicWall method
   req=Replace(req,"%","%25")
   req=Replace(req,":","%3A")
   req=Replace(req," ","%20")
   req=Replace(req,"?","%3F")
   req=Replace(req,"+","%2B")
   req=Replace(req,"&","%26")
   req=Replace(req,"=","%3D")

   'Rebuild the querystring variable
   qString = "sessionId=" & sessionId & "&ip=" & ip & "&mac=" & mac & "&ufi=" & ufi &
"&mgmtBaseUrl=" & mgmtBaseUrl & "&clientRedirectUrl=" & clientRedirectUrl & "&req=" & req

   'Add the optional hmac and cc vars if they are there.
   If hmac <> "" Then
      qString+="&hmac=" & hmac
   End If

   If customCode <> "" Then
      qString+="&cc=" & customCode
   End If

   'Bind the directory data
   Dim lhmDir As New DirectoryInfo(Server.MapPath("."))
   lhmList.DataSource = lhmDir.GetDirectories
   lhmList.DataBind()


End Sub

</script>

<STYLE>
body {
  font-size: 10pt;
  font-family: verdana,helvetica,arial,sans-serif;
  color:#000000;
  background-color:#9CBACE;
}

tr.heading {
  background-color:#006699;
}

.button {
   border: 1px solid #000000;
```

```
   background-color: #ffffff;
}

tr.hidden {
   font-size: 5pt;
   color:#9CBACE;
}

</STYLE>

<HTML>
<HEAD>
<TITLE>LHM Script Chooser</TITLE>
</HEAD>

<BODY>

<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr class="heading">
      <td colspan=3 align="center"><font color="white"> </td>
   </tr>
   <tr class="heading">
      <td width="50%" valign="center"><font color="white"><b>LHM Script
Chooser</b></font></td>
      <td><center><img width="216" height="51" src="SonicWall.gif"></center></td>
      <td width="50%" align="right" valign="center"><font
color="white"><b></b> </font></td>
   </tr>
   <tr class="heading">
      <td colspan=3 align="center"><font color="white"> </td>
   </tr>
</table>

<table width="100%" border="0" cellpadding="2" cellspacing="0">
   <tr><td><br></td></tr>
   <tr><td><H3>Please select one of the LHM Scripts below</H3></td></tr>
   <tr><td>Your original querystring information will be passed to the target script, and it
will open in a new window.</td></tr>
   <tr><td><br></td></tr>
</table>

<asp:Repeater id="lhmList" runat="server">
   <ItemTemplate >
       <li><a href = <%# DataBinder.Eval(Container.DataItem, "Name").ToString() &
"/default.aspx?" & qString & " target=""_blank""" %> >
       <%# DataBinder.Eval(Container.DataItem, "Name").ToString() %>
       </a>
       </li>
    </ItemTemplate>
 </asp:Repeater>

<table>
<tr class="hidden">
<td>default.aspx?sessionId=0b712fd83b9f5313db5af1cea6b1004f&ip=10.50.165.231&mac=00:0e:35:bd:c
9:37&ufi=0006b11184300&mgmtBaseUrl=https://10.50.165.193:4043/&clientRedirectUrl=https://10.50
.165.193:444/&req=http%3A//www.google.com/ig</td></tr>
</table>

</BODY>
</HTML>
```

B

IPv6

Topics:

- IPv6 on page 932
  - About IPv6 on page 932
  - Configuring IPv6 on page 938
  - IPv6 Visualization on page 961
  - IPv6 High Availability Monitoring on page 962
  - IPv6 Diagnostics and Monitoring on page 963

IPv6

This appendix provides an overview of the SonicOS implementation of IPv6, how IPv6 operates, and how to configure IPv6 for your network.

Topics:

- About IPv6 on page 932
- Configuring IPv6 on page 938
- IPv6 Visualization on page 961
- IPv6 High Availability Monitoring on page 962
- IPv6 Diagnostics and Monitoring on page 963

About IPv6

Topics:

- IPv6 Ready Certification on page 933
- IPv6 Technology Overview on page 933
- IPv6 Benefits on page 935
- SonicWall IPv6 Services and Features Currently Supported on page 936
- SonicWall IPv6 Features Not Currently Supported on page 936
- Supported IPv6 RFCs on page 936
- Non-Supported IPv6 RFCs on page 937

IPv6 Ready Certification

SonicWall has met the requirements for “IPv6 Ready” Phase-1 and Phase-2, as specified by the IPv6 Forum, a world-wide consortium providing technical guidance for the deployment of IPv6. The IPv6 Ready Logo Program is a conformance and interoperability testing program intended to increase user confidence by demonstrating that IPv6 is available now and ready to be used.

The IPv6 Ready series of tests extends from a basic level of minimum coverage in Phase-1 to a more complete coverage with Phase-2:

- Phase-1 (Silver) Logo: In a first stage, the Logo indicates that the product includes IPv6 mandatory core protocols and can interoperate with other IPv6 implementations.
- Phase-2 (Gold) Logo: The “IPv6 ready” step implies a proper care, technical consensus and clear technical references. The IPv6 Ready Logo indicates a product has successfully satisfied strong requirements stated by the IPv6 Logo Committee (v6LC).

SonicWall has been certified for Phase 2 (Gold) IPv6 Ready status. A future Phase-3 level of IPv6 Ready coverage is currently being developed.

For more information, see: http://www.ipv6ready.org/

NOTE: Wizards for IPv6 are not supported in SonicOS.

IPv6 Technology Overview

Every device connected to the Internet (computer, printer, smart phone, smart meter, etc.) requires an IP address. The Internet Protocol version 4 (IPv4) provides for approximately 4.3 billion unique IP addresses. The rapid global expansion in usage of the Internet, mobile phones, and VoIP telephony soon leads to the exhaustion of these 4.3 billion IP addresses.

On February 3rd, 2011, the Internet Assigned Numbers Authority (IANA) distributed the last-remaining blocks of IPv4 addresses to the Regional Internet Registries (RIRs). After the RIRs distribute these addresses to ISPs later this year, the world’s supply of new IPv4 addresses will be exhausted.

Luckily, the Internet Engineering Task Force (IETF) began planning for this day back around 1992, and in 1998, RFC 2460 was published to define Internet Protocol, Version 6 (IPv6). By increasing the address length from 32 bits to 128 bits, IPv6 dramatically increases the number of available addresses compared to IPv4:

- IPv4: 4,294,967,296 addresses
- IPv6: 340,282,366,920,938,463,463,374,607,431,768,211,456 addresses

Understanding IPv6 Addresses

IPv6 addresses are written in eight groups of four hexadecimal digits separated by colons:

```
XXXX:XXXX:XXXX:XXXX:XXXX:XXXX:XXXX:XXXX
```

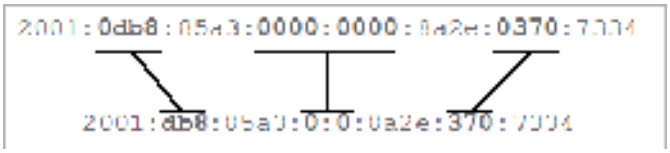
IPv6 addresses are logically divided into two parts: a 64-bit (sub-)network prefix, and a 64-bit interface identifier. Here is an example of an IPv6 address:

```
2001:0db8:85a3:0000:0000:8a2e:0370:7334
```

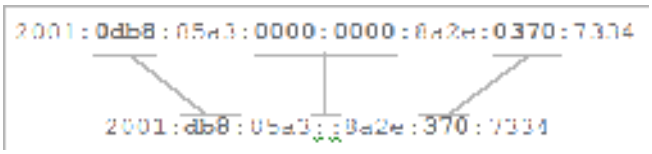
NOTE: The hexadecimal digits in IPv6 addresses are case-insensitive.

IPv6 address can be abbreviated using the following two rules:

1 Leading zeroes within a 16-bit value may be omitted. Thus, our example address can be abbreviated from the full form as follows:



2 Any number of consecutive groups of four zeros (technically 16-bits of zeros) can be expressed by a double colon (::). Combining these two rules, our example address can be abbreviated from the full form as follows:



TIP: The abbreviation for an empty address, or 0:0:0:0:0:0:0:0, is ::.

Types of IPv6 Addresses

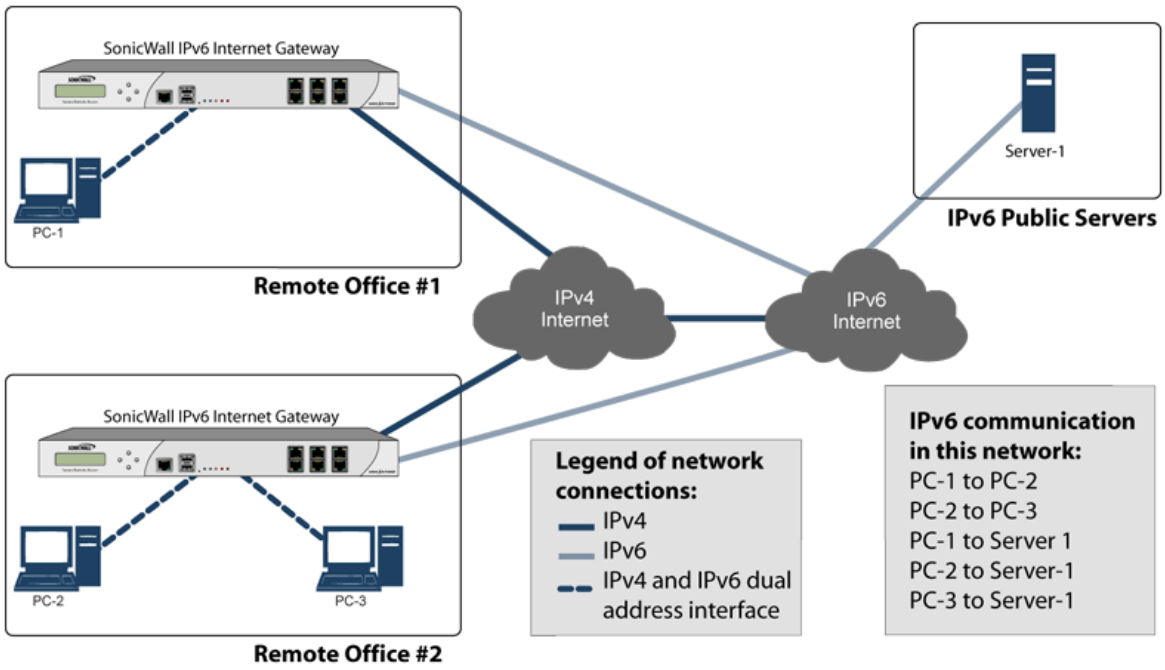
Type of Address	Full Address	Abbreviated Address
unicast address	1080:0:0:0:8:800:200C:417A	1080::8:800:200C:417A
multicast address	FF01:0:0:0:0:0:0:101	FF01::101
loopback address	0:0:0:0:0:0:0:1	::1
unspecified address	0:0:0:0:0:0:0:0	::

NOTE: Networks must have IPv4 Internet connectivity to get connected to IPv6 Internet.

NOTE: IPv6 stack must be enabled for computers at the local network sites.

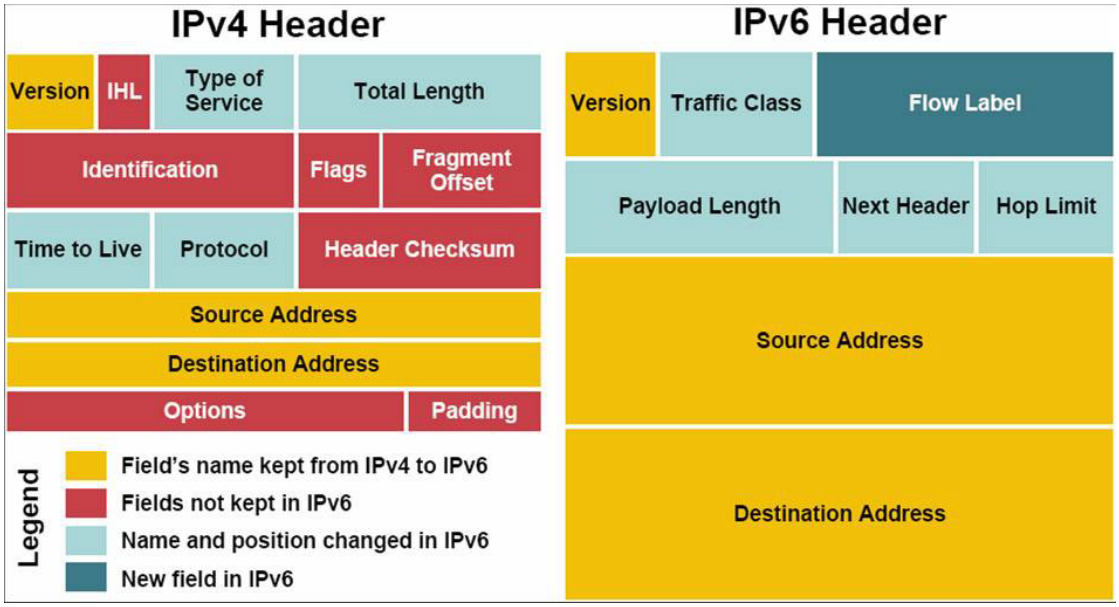
Typical IPv6 deployment is a simplified picture showing connectivity model for a typical IPv6 deployment.

Typical IPv6 deployment



Comparison of IPv4 and IPv6 header elements compares of the header elements between IPv4 and IPv6.

Comparison of IPv4 and IPv6 header elements



IPv6 Benefits

IPv6 brings some key features to improve the limitations exposed by IPv4. The new IP standard extends IPv4 in a number of important aspects:

- 6to4 tunnel (allows IPv6 nodes to connect to outside IPv6 services over an IPv4 network)
  - 6to4 Auto Tunnel
  - GRE Tunnel
- IPv6 Manual Tunnel
- New, simplified IPv6 header format
- Massively large number of available IPv6 addresses
- Efficient and hierarchical addressing and routing infrastructure
- Auto address assignment to hosts and routers using Neighbor Discovery Protocol (NDP) and DHCPv6
- Stateless and stateful address configuration
- Built-in security - AH and ESP strongly recommended
- Better support for QoS - Flow label in the header
- New protocol for neighboring node interaction
- Extensibility for new features using extension headers
  - Extension header detection report an log support
  - Extension header order check enforcement
  - Hop-by-hop extension header support
  - Inbound type 0 routing header packet check

SonicWall IPv6 Services and Features Currently Supported

For a complete list of currently supported IPv6 services and features, see the Knowledge Base article, *Supported/Unsupported IPv6 Features in SonicOS 6.x.x firmware*.

SonicWall IPv6 Features Not Currently Supported

NOTE: SonicOS is a dual IP stack firmware. Features that are not supported for IPv6 are still supported for IPv4.

For a complete list of IPv6 services and features currently not supported, see the Knowledge Base article, *Supported/Unsupported IPv6 Features in SonicOS 6.x.x firmware*.

Supported IPv6 RFCs

This section lists the IPv6 RFCs supported in SonicOS:

- TCP/IP stack and Network Protocols on page 936
- IPsec Conformance on page 937
- NAT Conformance on page 937
- DNS Conformance on page 937

TCP/IP stack and Network Protocols

- RFC 1886 DNS Extensions to support IP version 6 [IPAPPL DNS client]

- RFC 1981 Path MTU Discovery for IPv6
- RFC 2113 IP Router Alert Option
- RFC 2373 IPv6 Addressing Architecture
- RFC 2374 An IPv6 Aggregatable Global Unicast Address Format (obsoleted by 3587)
- RFC 2375 IPv6 Multicast Address Assignments
- RFC 2460 IPv6 specification
- RFC 2461 Neighbor discovery for IPv6
- RFC 2462 IPv6 Stateless Address Autoconfiguration
- RFC 2463 ICMPv6 for IPv6 specification
- RFC 2464 Transmission of IPv6 Packets over Ethernet Networks
- RFC 2473 Generic Packet Tunneling in IPv6 Specification
- RFC 2474 Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers
- RFC 2545 Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing
- RFC 2553 Basic Socket Interface Extensions for IPv6
- RFC 2710 Multicast Listener Discovery (MLD) for IPv6
- RFC 2711 IPv6 Router Alert Option
- RFC 2784 Generic Routing Encapsulation
- RFC 2893 Transition Mechanisms for IPv6 Hosts and Routers
- RFC 2991 Multipath Issues in Unicast and Multicast Next-Hop Selection
- RFC 3056 Connection of IPv6 Domains through IPv4 Clouds
- RFC 3484 Default Address Selection for Internet Protocol version 6 (IPv6) (no policy hooks)
- RFC 3493 Basic Socket Interface Extensions for IPv6
- RFC 3513 Internet Protocol Version 6 (IPv6) Addressing Architecture
- RFC 3542 Advanced Sockets Application Program Interface (API) for IPv6
- RFC 3587 IPv6 Global Unicast Address Format (obsoletes 2374)

IPsec Conformance

- RFC 1826 IP Authentication Header [old AH]
- RFC 1827 IP Encapsulating Security Payload (ESP) [old ESP]

NAT Conformance

- RFC 2663 IP Network Address Translator (NAT) Terminology and Considerations.
- RFC 3022 Traditional IP Network Address Translator (Traditional NAT).

DNS Conformance

- RFC 1886 DNS Extensions to support IP version 6

Non-Supported IPv6 RFCs

This section lists the IPv6 RFCs currently not supported in SonicOS:

- RFC 2002 IP Mobility Support
- RFC 2766 Network Address Translation - Protocol Translation (NAT-PT)
- RFC 2472 IP Version 6 over PPP
- RFC 2452 IP Version 6 Management Information Base for the Transmission Control Protocol.
- RFC 2454 IP Version 6 Management Information Base for the User Datagram Protocol.
- RFC 2465 Management Information Base for IP Version 6: Textual Conventions and General Group.

Configuring IPv6

Topics:

- IPv6 Interface Configuration on page 938
- Configuring IPv6 Tunnel Interfaces on page 948
- Accessing the SonicWall Management Interface Using IPv6 on page 958
- IPv6 Network Configuration on page 958
- IPv6 Access Rules Configuration on page 960
- IPv6 Advanced Firewall Settings on page 960
- IPv6 IPSec VPN Configuration on page 960
- SSL VPN Configuration for IPv6 on page 961

IPv6 Interface Configuration

IPv6 interfaces are configured on the **MANAGE | System Setup | Network > Interfaces** page by clicking **IPv6** in the **View IP Version** at the top right corner of the **Interface Settings** table.

By default, all IPv6 interfaces appear as routed with no IP address. Multiple IPv6 addresses can be added on the same interface. Auto IP assignment can only be configured on WAN interfaces.

NOTE: PortShield interfaces are not supported in IPv6.

Each interface can be configured to receive router advertisement or not. IPv6 can be enabled or disabled on each interface.

NOTE: The zone assignment for an interface must be configured through the IPv4 interface page before switching to IPv6 mode.

Topics:

- IPv6 Interface Configuration Constraints on page 939
- Configuring an Interface for IPv6 Static Mode on page 939
- Configuring Advanced IPv6 Interface Options and Multiple IPv6 Addresses on page 940
- Configuring Router Advertisement Settings on page 941
- Configuring Router Advertisement Prefix Settings on page 942
- Configuring an Interface for DHCPv6 Mode on page 943
- Configuring Advanced Settings for an IPv6 Interface on page 945
- Viewing DHCPv6 Protocol Information on page 946

- Configuring an Interface for Auto Mode on page 947
- PPPoE on page 948
- Configuring a VLAN Sub-Interface on page 948
- Configuring an Interface for Wire Mode on page 948

IPv6 Interface Configuration Constraints

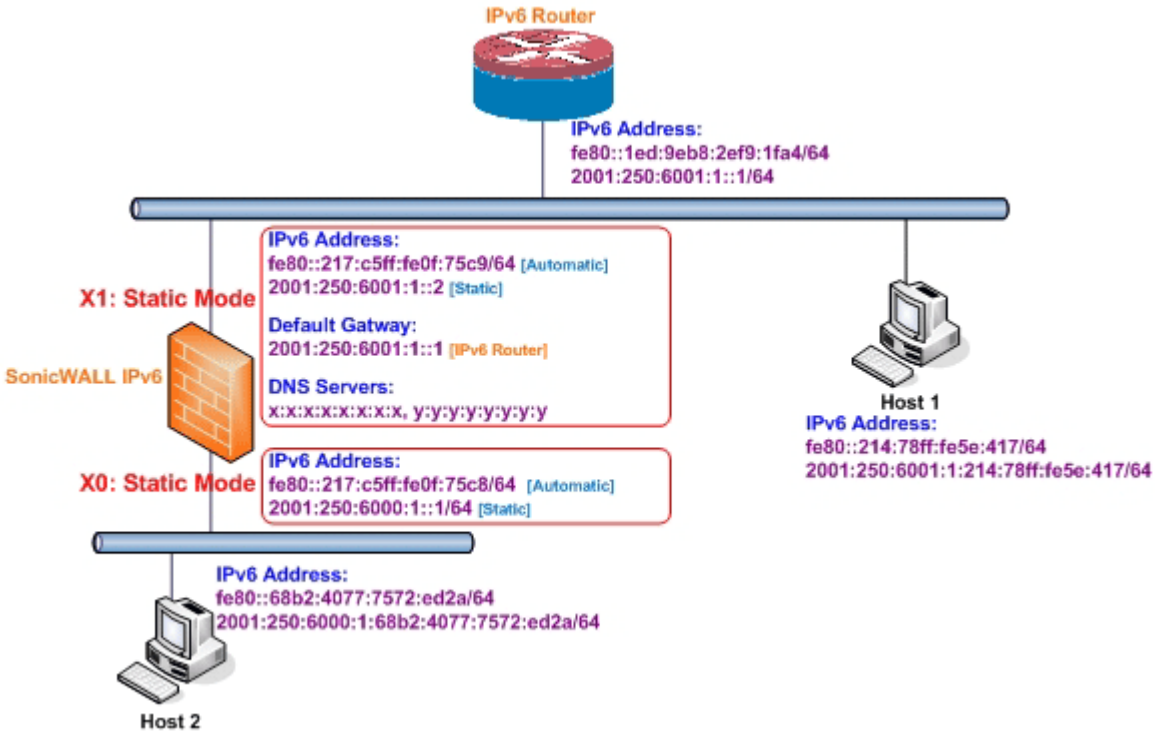
- The HA interface cannot be configured for IPv6.
- Only the parent interface of a SwitchPort group can be configured as an IPv6 interface; hence, all children of a switch port group must be excluded from this list.
- Zone and Layer 2 Bridge groups are shared configurations between by IPv4 and IPv6 on an interface. When they are configured on the IPv4 side, the IPv6 side of the interface uses the same configuration.
- Default Gateway and DNS Servers can only be configured for WAN zone interfaces.
- Wire mode is supported for IPv6, but you can not edit any settings. Instead, SonicOS uses the same configuration options set for IPv4.

Configuring an Interface for IPv6 Static Mode

Static mode provides user a way to assign static IPv6 address as opposed to an auto-assigned address. Using static mode, the IPv6 interface can still listen for Router Advertisements and learn an autonomous address from the appropriate prefix option. Static Mode does not disturb the running of Stateless Address Autoconfiguration on IPv6 interface unless the user manually disables it.

IPv6 static mode configuration shows a sample topology with IPv6 configured in static mode.

IPv6 static mode configuration



Three types of IPv6 address are possible to assign under this mode:

- Automatic Address
- Autonomous Address
- Static Address

To configure an interface for a static IPv6 address:

1. Navigate to the **MANAGE | System Setup | Network > Interfaces** page.
2. Click **IPv6** at the top right corner of the page. IPv6 addresses for the appliance are displayed.
3. Click **Configure** for the interface you want to configure an IPv6 address for. The **Edit Interface** dialog displays.

   NOTE: The zone assignment for interfaces must be configured on the IPv4 addressing page. To modify the zone assignment for an IPv6 interface, click **IPv4** at the top right of the page, modify the zone for the interface, and then return to the IPv6 interface page.

4. In the **IP Assignment** drop-down menu, select **Static**.
5. Enter the **IPv6 Address** for the interface.
6. Enter the **Prefix Length** for the address.
7. If this is the primary WAN interface, enter the IPv6 address of the **Default Gateway**. If this is not the primary WAN interface, any Default Gateway entry is ignored, so you can leave this as `::`. (The double colon is the abbreviation for an empty address, or `0:0:0:0:0:0:0:0`.)
8. If this is the primary WAN interface, enter up to three **DNS Server** IPv6 addresses. Again, if this is not the primary WAN interface, any DNS Server entries are ignored.
9. Select **Enable Router Advertisement** to make this an advertising interface that distributes network and prefix information.
10. Select **Advertise Subnet Prefix of IPv6 Primary Static Address** to add a default prefix into the interface advertising prefix list. This prefix is the subnet prefix of interface IPv6 primary static address. This option helps all hosts on the link stay in the same subnet.

Configuring Advanced IPv6 Interface Options and Multiple IPv6 Addresses

To modify Advanced IPv6 interface options or to configure multiple static IPv6 addresses:

1. In the **Edit Interface** dialog, click on the **Advanced** tab.
2. Click **Add Address** to configure multiple static IPv6 addresses for the interface. The **Add Interface IPv6 Address** dialog displays.

   NOTE: Multiple IPv6 addresses can only be added for an interface that is configured for Static IPv6 address mode. Multiple IPv6 addresses cannot be configured for **Auto** or **DHCPv6** modes.

3. Enter the **IPv6 Address** for the additional address for the interface.
4. Enter the **Prefix Length** for the address.
5. Select **Advertise Subnet Prefix of IPv6 Address** to add a default prefix into the interface advertising prefix list. This prefix is the subnet prefix of interface IPv6 primary static address. This option helps all hosts on the link stay in the same subnet.
6. Click **OK**.
7. The following additional options can be configured on the **Advanced** tab under the **Advanced Settings** heading:

- Select **Disable all IPv6 Traffic on the Interface** to stop the interface from handling all IPv6 traffic. Disabling IPv6 traffic can improve firewall performance for non-IPv6 traffic. This option is not selected by default.

  TIP: If the firewall is deployed in a pure IPv4 environment, SonicWall recommends enabling this option.

- Select **Enable Listening to Router Advertisement** to have the firewall receive router advertisement. If disabled, the interface filters all incoming Router Advertisement messages, which can enhance security by eliminating the possibility of receiving malicious network parameters (for example, prefix information or default gateway). This option is selected by default.

  NOTE: When this option is disabled, all assigned autonomous IPv6 address are removed from this interface.

  This option is not visible for **Auto** mode. In **Auto** mode, it is always enabled.

  - Select **Enable Stateless Address Autoconfiguration** to allow autonomous IPv6 addresses to be assigned to this interface. If unchecked, all assigned autonomous IPv6 address are removed from this interface.

    This option is not visible for **Auto** mode. In **Auto** mode, it is always enabled.

- Enter a numeric value for **Duplicate Address Detection Transmits** to specify the number of consecutive Neighbor Solicitation messages sent while performing Duplicate Address Detection (DAD) before assigning a tentative address to the interface. The minimum number is 0, the maximum is 9, and the default is **1**. A value of 0 indicates that DAD is not performed on the interface.
- In **Neighbor Discovery Base Reachable Time (seconds)**, enter a base value, in seconds, to use for computing the random Reachable Time value for the interface. The minimum value is 0, the maximum is 9999, and the default is **30**.

  A value of 0 indicates the parameter is unspecified, and the global setting in **Network > Neighbor Discovery** is used. If RA is enabled on this interface, however, the value in the **Reachable Time** option in the **Router Advertisement** tab is used.
- Select **Enable Max NDP Size Per Interface** to enable a maximum NDP size per interface. Every interface should have a maximum NP size for preventing system resources from being exhausted.
  - Enter the maximum NDP size in the Max NDP Size Per Interface field. The minimum value is 64, the maximum value is 9999, and the default values are **128** for WAN interfaces and **1200** for others.
- Similar with IPv4 gratuitous ARP, IPv6 node uses Neighbor Solicitation message to detect duplicate IPv6 address on the same link. DAD must be performed on any Unicast address (except Anycast address) before assigning a tentative to an IPv6 interface.

Configuring Router Advertisement Settings

Router Advertisement allows IPv6 routers to advertise DNS recursive server addresses to IPv6 hosts. Router Advertisement-based DNS configuration is a useful, optional alternative in networks where an IPv6 host's address is autoconfigured through IPv6 stateless address autoconfiguration, and where the delays in acquiring server addresses and communicating with the servers are critical. Router Advertisement allows the host to acquire the nearest server addresses on every link. Furthermore, it learns these addresses from the same RA message that provides configuration information for the link, thereby avoiding an additional protocol run. This can be beneficial in some mobile environments, such as with Mobile IPv6. SonicWall's implementation of IPv6 is full conformable with RFC 4861 in Router and Prefix Discovery.

NOTE: Router Advertisement can only be enabled when interface is under Static mode.

To configure Router Advertisement for an IPv6 interface:

1 In the **Edit Interface** dialog, click on the **Router Advertisement** tab.

2 Select the **Enable Router Advertisement** checkbox to make this an advertising interface that distributes network and prefix information.

3 Optionally, you can modify the following Router Advertisement settings:

- **Router Adv Interval Range (seconds)** – Enter the time interval allowed between unsolicited multicast Router Advertisements sent from the interface, in seconds. Advertisements are sent at a random value between the minimum and maximum interval:
  - Minimum interval – Enter the shortest interval allowed between Router Advertisements. The minimum time is 3 seconds, the maximum is 1350 seconds, and the default minimum time is **200** seconds.
  - Maximum interval – Enter the longest interval allowed between Router Advertisements. The minimum time is 4 seconds, the maximum is 1800 seconds, and the default maximum time is **600** seconds.
- **Link MTU** – Enter the recommended MTU for the interface link. The minimum value is 0, the maximum value is 99999, and the default value is **0**, which means the firewall does not advertise link MTU for the link.
- **Reachable Time (seconds)** – Enter the time that a node assumes a neighbor is reachable after having received a reachability confirmation. The minimum value is 0, the maximum value is 9999999999, and the default value is **0**, which means this parameter is unspecified by this firewall.
- **Retrans Time** – Enter the time between retransmitted Neighbor Solicitation messages. The minimum value is 0, the maximum value is 9999999999, and the default value is **0**, which means this parameter is unspecified by this firewall.
- **Current Hop Limit** – Enter the default value that should be placed in the Hop Count field of the IP header for outgoing IP packets. The minimum value is 0, which means this parameter is unspecified by this firewall; the maximum value is 255; and the default value is **64**.
- **Router Lifetime (seconds)** – Enter the lifetime when the firewall is accepted as a default router. The minimum value is 0 seconds, which means that the router is not a default router; the maximum time is 9000 seconds, and the default value is **1800** seconds.
- **Router Preference** – Indicates whether the advertising default router should be preferred over other default routers. Select **High**, **Medium** (default), or **Low** from the drop-down menu.

4 Select the **Managed** checkbox to set the managed address configuration flag in the Router Advertisement message. If set, the flag indicates that IPv6 addresses are available through Dynamic Host Configuration Protocol.

5 Select the **Other Configuration** checkbox to set the Other configuration flag in Router Advertisement message. If set, the flag indicates that other configuration information is available through Dynamic Host Configuration Protocol.

Configuring Router Advertisement Prefix Settings

Advertising prefixes provide hosts with prefixes for on-link determination and Address Autoconfiguration.

To configure a router advertisement prefix:

1 Go to the **Prefix List Settings** table on the **Router Advertisement** tab of the **Edit Interface** dialog.

2 Click **Add Prefix**. The **Add Advertising Prefix** dialog displays.

3 Enter the **Prefix** that is to be advertised with the Router Advertisement message.

4 Enter the **Valid Lifetime (minutes)** to set the length of time that the prefix is valid for the purpose of on-link determination. The minimum value is 1; the maximum value is 9999, which means the lifetime is infinite, and the default value is **43200** minutes.

5 Enter the **Preferred Lifetime (minutes)** to set the length of time that addresses generated from the prefix through stateless address autoconfiguration remain preferred. The minimum value is 1; the maximum value is 9999, which means the lifetime is infinite; and the default value is **10080** minutes.

6 Optionally select the **On-link** checkbox to enable the on-link flag in the Prefix Information option to indicate that this prefix can be used for on-link determination.

7 Optionally select the **Autonomous** checkbox to enable the autonomous address-configuration flag in Prefix Information option to indicate that this prefix can be used for stateless address configuration.

8 Click **OK**.

Configuring an Interface for DHCPv6 Mode

DHCPv6 (DHCP for IPv6) is a client/server protocol that provides stateful address configuration or stateless configuration setting for IPv6 hosts. DHCPv6 client is enabled to learn IPv6 address and network parameters when the interface is configured to DHCPv6 mode.

DHCPv6 defines two different configuration modes:

- **DHCPv6 stateful mode**: DHCPv6 clients require IPv6 address together with other network parameters (for example, DNS Server, Domain Name).
- **DHCPv6 stateless mode**: DHCPv6 client only obtains network parameters other than IPv6 address.

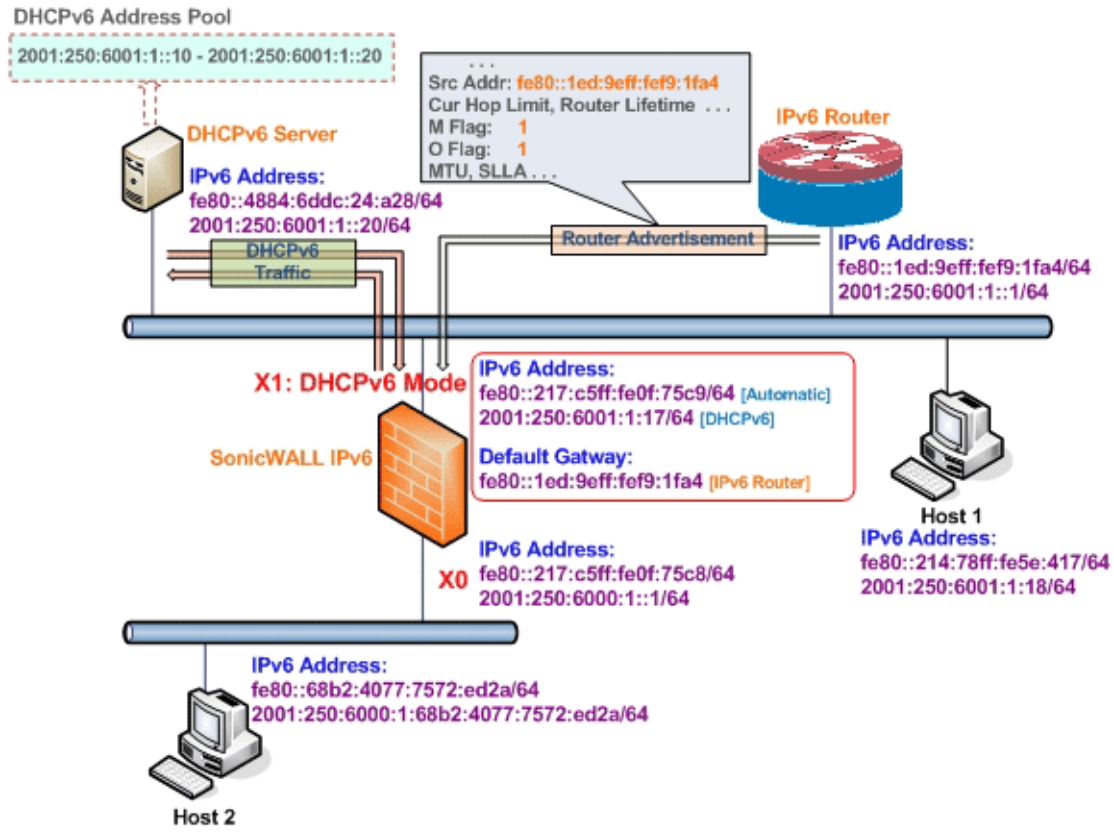
Choosing which mode depends on the Managed (M) Address Configuration and Other (O) Configuration flags in the advertised Router Advertisement message:

DHCPv6 infrastructure

Flag		Configuration
M	O	
0	0	No DHCPv6 infrastructure.
1	1	IPv6 host uses DHCPv6 for both IPv6 address and other network parameter settings.
0	1	IPv6 host uses DHCPv6 only for IPv6 address assignment.
1	0	IPv6 host uses DHCPv6 only for other network parameter settings, known as DHCPv6 stateless.

DHCPv6 topology shows a sample DHCPv6 topology.

DHCPv6 topology



There are three types of IPv6 addresses that can be assigned under DHCPv6:

- Automatic Address
- Autonomous Address
- IPv6 Address assigned through DHCPv6 client

To configure an interface for a DHCPv6 address:

1 Navigate to **MANAGE | System Setup | Network > Interfaces**.

2 If you are configuring an unassigned interface, click **IPv4** at the top right corner of the page.

3 Click **Edit** for the interface to be configured. The **Edit Interface** dialog displays.

4 Select **WAN** from the **Zone** drop-down menu. More options appear.

5 Select **DHCP** from the **IP Assignment** drop-down menu.

6 Click **OK**.

7 Click **IPv6** at the top right corner of the page. IPv6 addresses for the appliance are displayed.

8 Click **Configure** for the interface you want to configure an IPv6 address for. The **Edit Interface** dialog displays.

9 In the **IP Assignment** drop-down menu, select **DHCPv6**. The options change.

10 The following options can be configured for IPv6 interfaces configured for DHCPv6 mode:

- **Enable DHCPv6 prefix delegation** - If enabled, these options become available:

- **Send preferred delegated prefix** - Select this option to require a DHCAPv6 client to try to send the preferred delegated prefix specified in the two fields.
- **Send hints for renewing previous delegated prefix on startup** - Select this option to require a DHCPv6 client to try to renew the delegated prefix assigned before when the firewall started up.

- **Use Rapid Commit Option** - If enabled, DHCPv6 client use Rapid Commit Option to use the two message exchange for address assignment.
- **Send hints for renewing previous IP on startup** - If enabled, DHCPv6 client tries to renew the address assigned before when firewall startup.

11 Select the **DHCPv6 Mode** for the interface. As required by RFC, DHCPv6 client depends on the Router Advertisement message to decide which mode (stateful or stateless) it should choose. This definition limits the user's choice to determine the DHCPv6 mode by itself. SonicWall's implementation of DHCPv6 defines two different modes to balance the conformance and flexibility:

- **Automatic** - The IPv6 interface configures IPv6 addresses using stateless/stateful autoconfiguration in accord with the M and O settings in the most recently received router advertisement message. See DHCPv6 infrastructure.
- **Manual** - The DHCPv6 mode is manually configured regardless of any received Router Advertisement.

The **Only Request Stateless Information** option determines which DHCPv6 mode is used. If this option is unchecked, DHCPv6 client is under stateful mode; if it is checked, DHCPv6 client is under stateless mode and only obtains network parameters.

12 Optionally, select the **Only Request Stateless Information** checkbox to have DHCPv6 clients only request network parameter setting from the DHCPv6 server. The IPv6 address is assigned through stateless auto-configuration.

13 Optionally, you can configure **Management** login or **User Login**.

14 Optionally click the **Advanced** tab to configure Advanced options and/or click the **Protocol** tab to view DHCPv6 stateful and stateless configuration information.

15 Click **OK** to complete the configuration.

Configuring Advanced Settings for an IPv6 Interface

To configure advanced IPv6 interface settings:

1 On the **Edit Interface** dialog, click the **Advanced** tab.

2 Select **Disable all IPv6 Traffic on the Interface** to stop the interface from handling all IPv6 traffic. Disabling IPv6 traffic can improve firewall performance for non-IPv6 traffic. This option is not selected by default.

> **TIP:** If the firewall is deployed in a pure IPv4 environment, SonicWall recommends enabling this option.

3 Select **Enable Listening to Router Advertisement** to have the firewall receive router advertisement. If disabled, the interface filters all incoming Router Advertisement messages, which can enhance security by eliminating the possibility of receiving malicious network parameters (for example, prefix information or default gateway). This option is not selected by default.

> **NOTE:** If this option is disabled, all assigned autonomous IPv6 addresses are removed from this interface.

This option is not visible for Auto mode. In Auto mode, it is always enabled.

When this option is selected the Enable Stateless Address Autoconfiguration option becomes available.

- Select **Enable Stateless Address Autoconfiguration** to allow autonomous IPv6 addresses to be assigned to this interface. If unchecked, all assigned autonomous IPv6 addresses are removed from this interface.

  NOTE: If this option is disabled, all assigned autonomous IPv6 addresses are removed from this interface.

  This option is not visible for Auto mode. In Auto mode, it is always enabled.

4 Enter a numeric value for **Duplicate Address Detection Transmits** to specify the number of consecutive Neighbor Solicitation messages sent while performing Duplicate Address Detection (DAD) before assigning a tentative address to this interface. The minimum value is 0, which indicates that DAD is not performed on the interface; the maximum number is 9; and the default number is **1**.

Similar to IPv4 gratuitous ARP, IPv6 node uses a Neighbor Solicitation message to detect a duplicate IPv6 address on the same link. DAD must be performed on any Unicast address (except Anycast address) before assigning a tentative to an IPv6 interface.

5 In **Neighbor Discovery Base Reachable Time (seconds)**, enter a base value, in seconds, to use for computing the random Reachable Time value for the interface. The minimum value is 0, the maximum is 9999, and the default is **30**.

A value of 0 indicates the parameter is unspecified, and the global setting in **Network > Neighbor Discovery** is used. If RA is enabled on this interface, however, the value in the **Reachable Time** option in the **Router Advertisement** tab is used.

6 Select **Enable Max NDP Size Per Interface** to enable a maximum NDP size per interface. Every interface should have a maximum NP size for preventing system resources from being exhausted. This option is selected by default.

Enter the maximum NDP size in the Max NDP Size Per Interface field. The minimum value is 64, the maximum value is 9999, and the default values are **128** for WAN interfaces and **1200** for others.

Viewing DHCPv6 Protocol Information

When configuring an IPv6 interface in DHCpv6 mode, the **Protocol** tab displays additional DHCPv6 information.

- **DHCPv6 General Information**
  - **DHCPv6 State**: If the interface is configured for:
    - Stateless mode, the DHCPv6 State is Stateless.
    - Stateful mode, the DHCPv6 State is either **Enabled** or **Disabled**.

      When the interface is in Stateful DHCPv6 mode, mousing over the **Comment** icon displays current Router Advertisement information for the interface.
  - **DHCPv6 Server**: The IPv6 address of the DHCPv6 server.
  - **DHCPv6 DUID**: The DUID (DHCP Unique Identifier) or host identifier.
- **Stateful Addresses Acquired via DHCPv6**: Displays information on any acquired stateful IPv6 addresses:
  - IAID (Identity Association Identifier)
  - Type
  - IPv6 Address
  - Lease Expires
- **Stateless Configuration Settings Acquired via DHCPv6**
  - **DNS Servers 1/2/3**: The IPv6 addresses of any DNS Servers.

You can renew, release, or refresh the DNS servers by clicking the appropriate button.

- **Delegated Prefixes Acquired via DHCPv6**: Displays information on any acquired delegated prefixes for stateful IPv6 addresses:

    - IAID
    - Type
    - IPv6 Prefix
    - Prefix Length
    - Lease Expires

    You can renew, release, or refresh the prefixes by clicking the appropriate button.

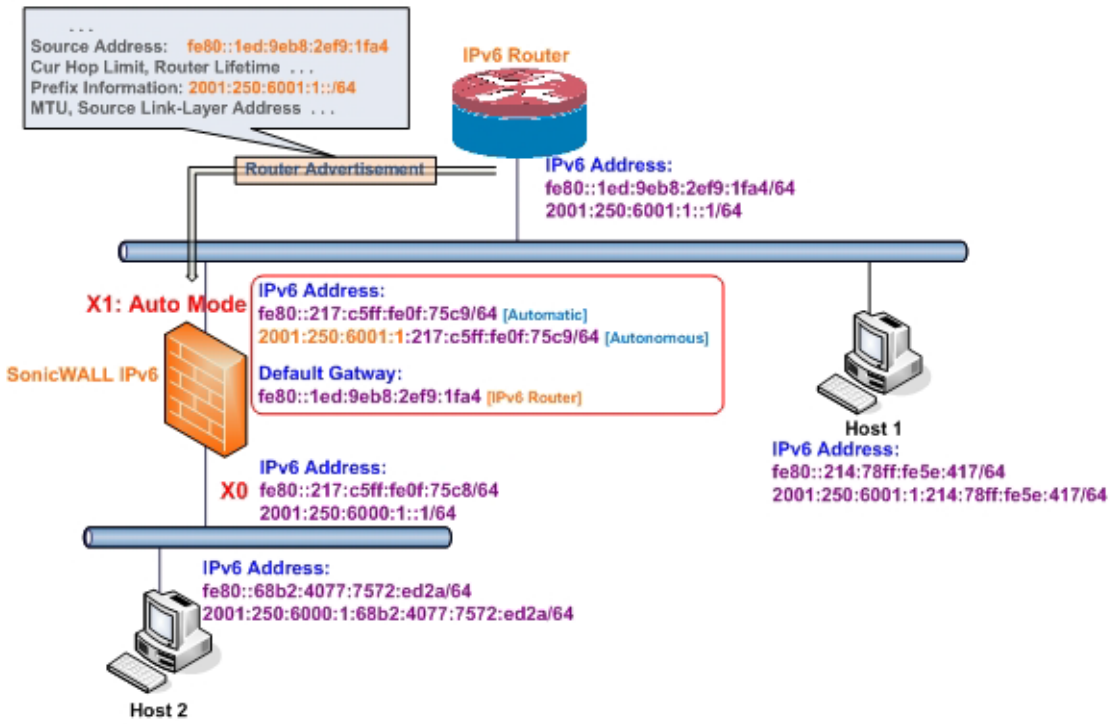
Configuring an Interface for Auto Mode

Auto mode utilities IPv6's Stateless Address Autoconfiguration to assign IPv6 address. This mode does not require any manual address configuration by the network administrator. The Security Appliance listens to the network and receives prefix information from neighboring routers. The IPv6 Stateless Address Autoconfiguration feature performs all configuration details, such as IPv6 address assignment, address deleting for address conflicting or lifetime expiration, and default gateway selection based on the information collected from on-link router.

NOTE: Auto mode can only be configured for the WAN zone. For security consideration, Auto mode is not available on LAN zone interface.

IPv6 auto mode configuration shows a sample topology for IPv6 configured in Auto mode.

IPv6 auto mode configuration



In this mode, 2 types of IPv6 address are possible to assign:

- Automatic Address - The interface default link-local address. It is never timed out and is not able to be edited or deleted.
- Autonomous Address - Assigned from Stateless Address Autoconfiguration. Users can manually delete the address if they do not want to wait for its valid lifetime expires.

To configure an IPv6 interface for Auto mode:

1. Navigate to **MANAGE | System Setup | Network > Interfaces**.
2. Click **IPv6** at the top right corner of the page to display IPv6 addresses.
3. Click **Configure** for the interface you want to configure an IPv6 address for. The **Edit Interface** dialog displays.
4. In the **IP Assignment** drop-down menu, select **Auto**.
5. Optionally, you can select enter a numeric value for **Duplicate Address Detection Transmits** on the **Advanced** tab to specify the number of consecutive Neighbor Solicitation messages sent while performing Duplicate Address Detection (DAD) before assigning a tentative address to interface. A value of 0 indicates that DAD is not performed on the interface.
6. Click **OK**.

PPPoE

Only PPPoE Client Mode is supported in IPv6.

Configuring a VLAN Sub-Interface

The procedure for configuring a VLAN Sub-interface in IPv6 is identical to that in IPv4. Refer to Configuring Virtual Interfaces (VLAN Subinterfaces) on page 305 for details.

All VLAN Sub-interfaces must be configured in IPv4, before configuring them in IPv6.

Configuring an Interface for Wire Mode

NOTE: Wire mode is supported on NSA 2600 and higher appliances.

The procedure for configuring a Wire Mode interface in IPv6 is identical to that in IPv4. Refer to Configuring an Interface for Wire Mode on page 312 for details.

All Wire Mode interfaces must be configured in IPv4; you can not edit Wire Mode settings in IPv6. Any functionality enabled in IPv4 (for example, Link State Propagation) applies to IPv6.

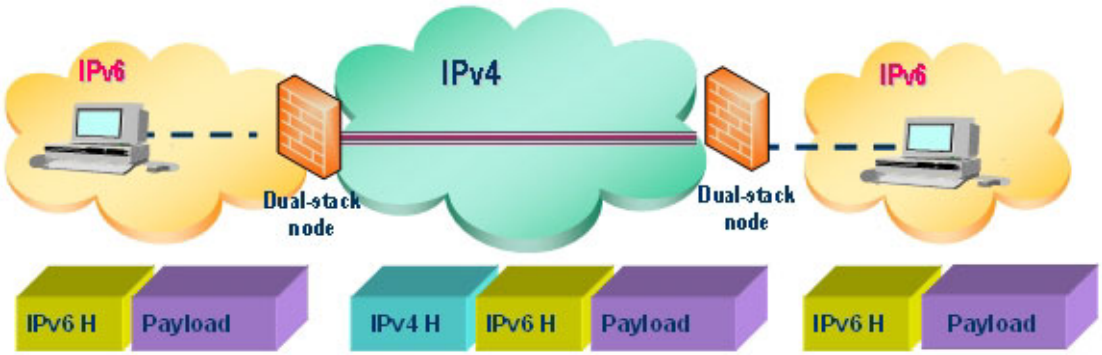
Configuring IPv6 Tunnel Interfaces

This section describes how to tunnel IPv4 packets through IPv6 networks and IPv6 packets through IPv4 networks. For instance, to pass IPv6 packets through the IPv4 network, the IPv6 packet is encapsulated into an IPv4 packet at the ingress side of a tunnel. When the encapsulated packet arrives at the egress of the tunnel, the IPv4 packet is de-capsulated.

Tunnels can be either automatic or manually configured. A configured tunnel determines the endpoint addresses by configuration information on the encapsulating node. An automatic tunnel determines the IPv4 endpoints from the address of the embedded IPv6 datagram. IPv4 multicast tunneling determines the endpoints through Neighbor Discovery.

IPv6-to-IPv4 tunnel interface depicts an IPv6-to-IPv4 tunnel.

IPv6-to-IPv4 tunnel interface



Topics:

- Configuring the 6to4 Auto Tunnel on page 949
- Configuring 6to4 Relay for Non-2002 Prefix Access on page 950
- Configuring a Manual IPv6 Tunnel on page 951
- Configuring a GRE IPv6 Tunnel on page 952
- IPv6 Prefix Delegation on page 952
- 6rd Tunnel Interfaces on page 954
- Configuring an ISATAP Tunnel on page 955

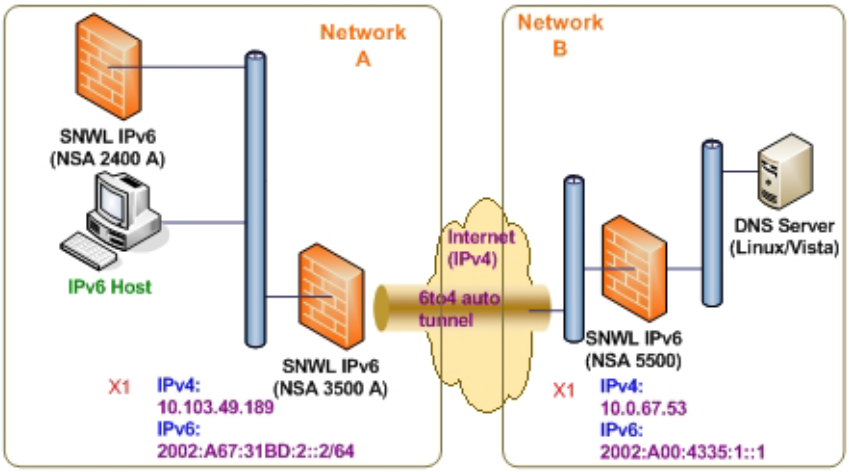
Configuring the 6to4 Auto Tunnel

The 6to4 Auto Tunnel is an automatic tunnel: tunnel endpoints are extracted from the encapsulated IPv6 datagram. No manual configuration is necessary.

6to4 tunnels use a prefix of the form `2002:tunnel-IPv4-address::/48` to tunnel IPv6 traffic over IPv4 (for example, if the tunnel's IPv4 endpoint has the address a01:203, the 6to4 tunnel prefix is 2002:a01:203::1). Routers advertise a prefix of the form `2002:[IPv4]:xxxx/64` to IPv6 clients. For complete information, see RFC 3056.

6to4 auto tunnel topology shows a sample 6to4 auto tunnel topology.

6to4 auto tunnel topology



In IPv6-to-IPv4 tunnel interface, customers do not need to specify the tunnel endpoint, but only need to enable the 6to4 auto tunnel. All packets with a 2002 prefix are routed to the tunnel, and the tunnel's IPv4 destination is extracted from the destination IPv6 address.

6to4 tunnels are easy to configure and use. Users must have a global IPv4 address and IPv6 address, which must also have a 2002 prefix. Therefore, in general, a user can only access network resources with a 2002 prefix.

NOTE: Only one 6to4 auto tunnel can be configured on the Security Appliance.

NOTE: VPN Tunnel Interfaces have automatically created IPv6 link local addresses.

To configure the 6to4 auto tunnel on the firewall:

1 Navigate to **MANAGE | System Setup | Network > Interfaces**.

2 Either:

- Click **Add Interface**.
- Select **Tunnel Interface** from the **Add Interface** drop-down menu.

The **Edit Interface** dialog displays.

3 Select the **Zone** for the 6to4 tunnel interface. This is typically the WAN interface.

4 In the **Tunnel Type** drop-down menu, select **6to4 Auto Tunnel Interface**.

5 Specify a name in the **Name** field. By default, the interface **Name** is set to **6to4AutoTun**.

6 Select the **Enable IPv6 6to4 Tunnel** checkbox. By default, this checkbox is selected.

7 Optionally, you can configure one or more **Management** login protocols: **HTTPS**, **Ping**, or **SNMP**.

NOTE: Selecting **HTTPS** enables the **Add rule to enable redirect from HTTP to HTTPS** option automatically. This option cannot be selected for the other protocols. For more information about this option, see HTTP/HTTPS Redirection on page 267.

8 Optionally, you can configure either or both **User Login** protocols: **HTTP** or **HTTPS**.

NOTE: Selecting only **HTTPS** enables the **Add rule to enable redirect from HTTP to HTTPS** option automatically. For more information about this option, see HTTP/HTTPS Redirection on page 267. If you also select **HTTP**, the **Add rule to enable redirect from HTTP to HTTPS** option is deselected and cannot be selected.

9 Click **OK**.

Configuring 6to4 Relay for Non-2002 Prefix Access

By default, 6to4 auto tunnel can only access the destination with a 2002 prefix. The 6to4 relay feature can be used to access non-2002 prefix destinations.

To enable 6to4 relay:

1 Navigate to **MANAGE | System Setup | Network > Routing**.

2 Click **Add** to create a Route Policy that can route all traffic destined for 2003 prefixes over the 6to4 auto tunnel interface:

This static route can be added on the 6to4 auto tunnel interface to enable the relay feature, which makes it possible to access the IPv6 destination with non-`2002:` prefix through 6to4 tunnel.

NOTE: The gateway must be the IPv6 address with the `2002:` prefix.

Configuring a Manual IPv6 Tunnel

To configure the 6to4 tunnel on the firewall:

1 Navigate to **MANAGE | System Setup | Network > Interfaces**.

2 Click **Add Interface**. The **Edit Interface** dialog displays.

3 Select the **Zone** for the tunnel interface.

4 In the **Tunnel Type** drop-down menu, select **IPv6 Manual Tunnel Interface**. This is the default.

5 Enter a **Name** for the tunnel interface.

6 Enter an address in the **Tunnel Interface IPv6 Address** field. The field starts with **::** already.

7 Select an interface to which the tunnel is bound from the **Bound to** drop-down menu. The default is **X1**.

8 From the **Remote IPv4 Address** drop-down menu, select an IPv4 address object for the tunnel endpoint.

9 From the **Remote IPv6 network** drop-down menu, select an IPv6 Address object, which can be a group, range, network, or host.

10 Optionally, you can configure one or more **Management** login protocols: **HTTPS**, **Ping**, or **SNMP**.

> **NOTE:** Selecting **HTTPS** enables the **Add rule to enable redirect from HTTP to HTTPS** option automatically. For more information about this option, see HTTP/HTTPS Redirection on page 267. This option cannot be selected for the other protocols.

11 Optionally, you can configure either or both **User Login** protocols: **HTTP** or **HTTPS**.

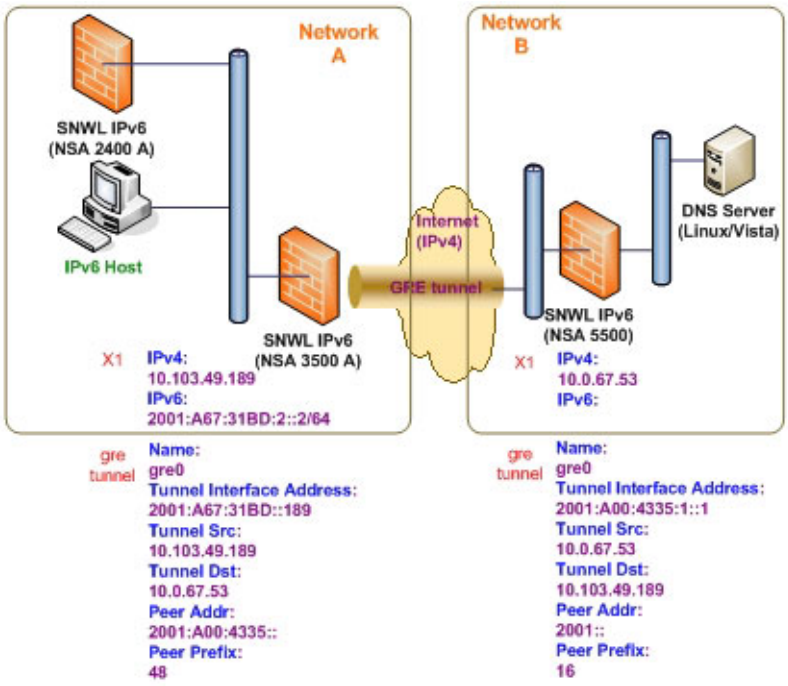
> **NOTE:** Selecting only **HTTPS** enables the **Add rule to enable redirect from HTTP to HTTPS** option automatically. For more information about this option, see HTTP/HTTPS Redirection on page 267. If you also select **HTTP**, the **Add rule to enable redirect from HTTP to HTTPS** option is deselected and cannot be selected.

12 Click **OK**.

Configuring a GRE IPv6 Tunnel

GRE can be used to tunnel IPv4 and IPv6 traffic over IPv4 or IPv6. GRE tunnels are static tunnels where both endpoints are specified manually. GRE IPv6 tunnel configuration shows a sample GRE IPv6 tunnel.

GRE IPv6 tunnel configuration



The configuration of a GRE tunnel is similar to a manual tunnel, except **GRE Tunnel Interface** is selected for the **Tunnel Type**.

IPv6 Prefix Delegation

IPv6 Prefix Delegation, also known as DHCPv6 Prefix Delegation (DHCPv6-PD), is an extension to DHCPv6. In DHCPv6, addresses are assigned by a DHCPv6 server to an IPv6 host. In DHCPv6-PD, complete IPv6 subnet addresses and other parameters are assigned by a DHCPv6-PD server to a DHCPv6-PD client.

When DHCPv6-PD is enabled, it is applied to all DHCPv6 interfaces attached to the WAN zone. DHCPv6-PD is an additional subnet-configuration mode that co-exists with DHCPv6.

The IPv6 address is a combination of the prefix provided by the DHCPv6-PD server and the suffix provided by the DHCPv6-PD client. The prefix length is 64 by default, but can be edited.

When the firewall starts, a default address object group called *Prefixes from DHCPv6 Delegation* is automatically created. Prefixes delegated from the upstream interface are members of this group.

IPv6 Prefix Delegation is configured on:

- An Upstream Interface
- One or More Downstream Interfaces

When the upstream interface learns the prefix delegation from the DHCPv6-PD server, SonicOS calculates and applies the IPv6 address prefixes to all the downstream interfaces, and the downstream interfaces advertise this information to all the hosts in their network segments.

This section contains the following configuration procedures:

- Configuring IPv6 Prefix Delegation on the Upstream Interface on page 953

- Configuring IPv6 Prefix Delegation on the Downstream Interface on page 953

IMPORTANT: Before you disable prefix delegation in your network, we recommend that you release the prefix delegation in the upstream interface first.

Configuring IPv6 Prefix Delegation on the Upstream Interface

To configure IPv6 Prefix Delegation on the upstream interface:

1. Go to **MANAGE | System Setup | Network > Interfaces**.
2. At **View IP Version**, select **IPv6**.
3. Click the **Edit** icon in the **Configure** column for the Interface you want to configure as the upstream interface. The **Edit Interface** dialog appears.

   NOTE: The **Zone** is always **WAN**.

4. From the **IP Assignment** menu, select **DHCPv6**.
5. Select the **Enable DHCPv6 prefix delegation** option.
6. From the **DHCPv6 Mode** menu, select **Manual**.
7. To see the configured DHCPv6 information, click the **Protocol** tab.

   In the **DHCPv6 General Information** panel, the **DHCPv6 DUID** is displayed.

   In the **Stateful Addresses Acquired via DHCPv6** panel, the stateful **IAID** is displayed.

   In the **Delegated Prefixes Acquired via DHCPv6** panel, the delegated **IAID** is displayed.
8. Click **Renew**. The information for the other columns is displayed.

Configuring IPv6 Prefix Delegation on the Downstream Interface

To configure IPv6 Prefix Delegation on the downstream interface:

1. Go to **MANAGE | System Setup | Network > Interfaces**.
2. Select the **IPv6** option.
3. Click the **Edit** icon in the **Configure** column for the Interface you want to configure as the downstream interface. The **Edit Interface** dialog appears.
4. Select the **Enable Router Advertisement** option.
5. Click the **Advanced** tab.

   If the upstream prefix is obtained, it is displayed in the **IPv6 Addresses** panel.
6. If the upstream prefix cannot be obtained, an alternate address is displayed in the **IPv6 Addresses** panel.
7. Click **Add Address** to display the **Add IPv6 Address** dialog.
8. Select the **Add Downstream Delegated IPv6 Address** option.
9. (Optional) Select the **Advertise Subnet Prefix of Static IPv6 Address** option.
10. Click the **Router Advertisement** tab.
11. Select the **Enable Router Advertisement** option.

    If you selected **Advertise Subnet Prefix of Static IPv6 Address** option under the **General** tab, the prefix is listed in the **Prefix List Settings** panel.
12. To see your new IPv6 PD interfaces, go to **MANAGE | System Setup | Network > Routing**.

13 Select the **IPv6** option.

The two new IPv6 interfaces with prefix delegation (upstream and downstream) are displayed.

6rd Tunnel Interfaces

IPv6 Rapid Deployment (6rd) enables IPv6 to be deployed across an IPv4 network quickly and easily. 6rd utilizes a Service Provider's existing IPv6 address prefixes, ensuring that the 6rd operational domain is limited to the Service Provider's network and is under the Service Provider's direct control.

A 6rd tunnel interface is a virtual interface that transports 6rd encapsulated IPv6 packets in an IPv4 network.

NOTE: A 6rd tunnel interface must be bound to a physical or a virtual interface.

When 6rd is deployed, the IPv6 service is equivalent to native IPv6. 6rd mapping of IPv6 addresses to IPv4 addresses provides automatic determination of IPv4 tunnel endpoints from IPv6 prefixes, allowing stateless operation of 6rd.

A 6rd domain consists of several 6rd customer edge (CE) routers and one or more 6rd border relay (BR) routers. IPv6 packets encapsulated by 6rd follow the IPv4 routing topology within the service provider network.

A typical 6rd implementation using customer edge routers and border relay routers requires only one 6rd tunnel interface. A border relay router servicing multiple 6rd domains may have more than one 6rd tunnel interface. However, each 6rd domain can have only one 6rd tunnel interface.

IPv6 packets traverse the border relays when they enter or exit a Service Provider's 6rd domain. Because 6rd is stateless, packets can be sent to the border relays using the Anycast method, where packets from a single source are routed to the nearest node in a group of potential receivers, or to several nodes, all identified by the same destination address.

Service Providers may deploy 6rd in a single domain or in multiple domains. A 6rd domain can have only one 6rd prefix. Different 6rd domains must use different 6rd prefixes.

On **MANAGE | System Setup | Network > Routing**, in the **Route Policies** panel, there are four default route policies for 6rd tunnel interfaces.

There are two configuration modes:

- Manual
- DHCP

The following four 6rd parameters can be set manually, or they can be set automatically by the DHCPv4 server if you select DHCP as the configuration mode.

- IPv4 Mask Length
- 6rd Prefix
- 6rd Prefix Length
- 6rd BR IPv4 Address

In DHCP mode, the 6rd parameters are received from the bound interface. In Manual mode, the 6rd parameters must be configured manually.

Configuring a 6rd Tunnel Interface

A 6rd tunnel interface is configured in the same way as other IPv6 tunnel interfaces. A bound interface is required to configure a 6rd tunnel interface.

To configure a 6rd tunnel interface:

1 Go to **MANAGE | System Setup | Network > Interfaces**.

2 At **View IP Version**, select **IPv6**.

3 At the bottom of the **Interface Settings** panel, click **Add Interface**.

NOTE: The **Protocol** tab is shown only when you select DHCP as the Configure Mode.

4 From the **Zone** drop-down menu, select **WAN**.

5 The **Interface Type** menu is disabled. It already has **Tunnel Interface** selected as it was selected from the **Add Interface** menu in Step 3.

6 From the **Tunnel Type** menu, select **6rd Tunnel Interface**.

7 In the name box, enter a name for your tunnel interface, or example, **6rd Tunnel**.

8 In the **Tunnel Interface IPv6 Address** field, enter the IPv6 address of the tunnel interface. For example, **2001::2**.

9 In the **Prefix Length** field, enter the length for the IPv6 prefix. For example, **64**.

10 From the **Bound to** drop-down menu, select the interface that you want, such as **X1**.

11 From the **Configure Mode** drop-down menu, select the mode you want: **Manual** or **DHCP**.

NOTE: If you select **Manual** as the **Configure Mode**, do Step 12 through Step 15.

If you select **DHCP** as the **Configure Mode**, skip Step 12 through Step 15.

12 In the **6rd Prefix** field, enter the 6rd prefix, such as **2222:2222::** (**Manual** mode only).

13 In the **6rd Prefix Length** field, enter the length for the 6rd prefix, such as **32** (**Manual** mode only).

14 In the **IPv4 Mask Length** field, enter the length of the IPv4 subnet mask (**Manual** mode only).

15 In the **BR IPv4 Address** field, enter the IPv4 address of the 6rd border relay (**Manual** mode only).

16 (Optional) In the **Comment** field, enter a comment to describe the tunnel interface.

17 Select the **Add Default Route Automatically** option.

18 Select the **Management** options that you want, or select the **User Login** options that you want.

If you selected **Manual** as the **Configure Mode**, your 6rd Tunnel Interface settings are shown under the **General** tab.

If you selected **DHCP** as the **Configure Mode**, your 6rd Tunnel Interface settings are shown under the **Protocol** tab.

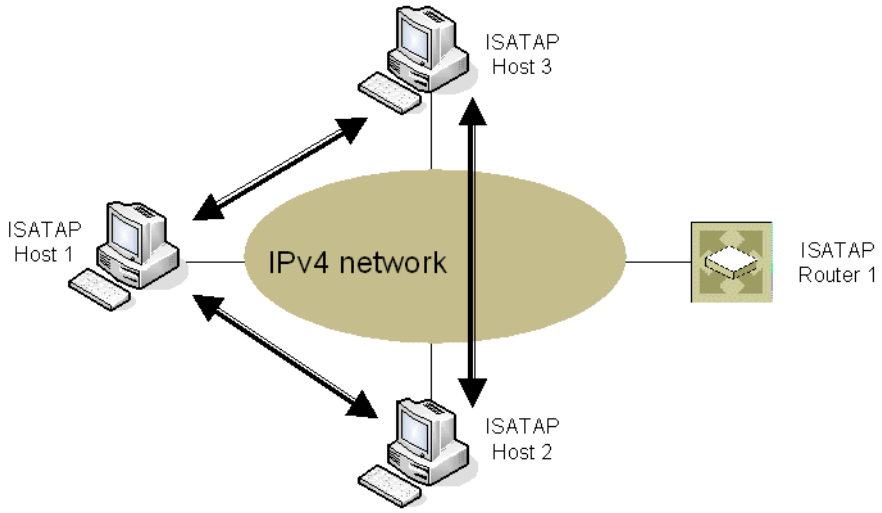
Configuring an ISATAP Tunnel

ISATAP (Intra-Site Automatic Tunnel Addressing Protocol) can be used to provide IPv6 connectivity through an IPv4-only infrastructure. ISATAP is a simple tunneling mechanism that connects dual-stack (IPv6/IPv4) node to other dual-stack nodes or IPv6 nodes over IPv4 networks. The IPv4 network is viewed by ISATAP as a link layer for IPv6.

ISATAP can be used in several scenarios to provide unicast connectivity between ISATAP hosts, and ISATAP host and hosts on IPv6 networks.

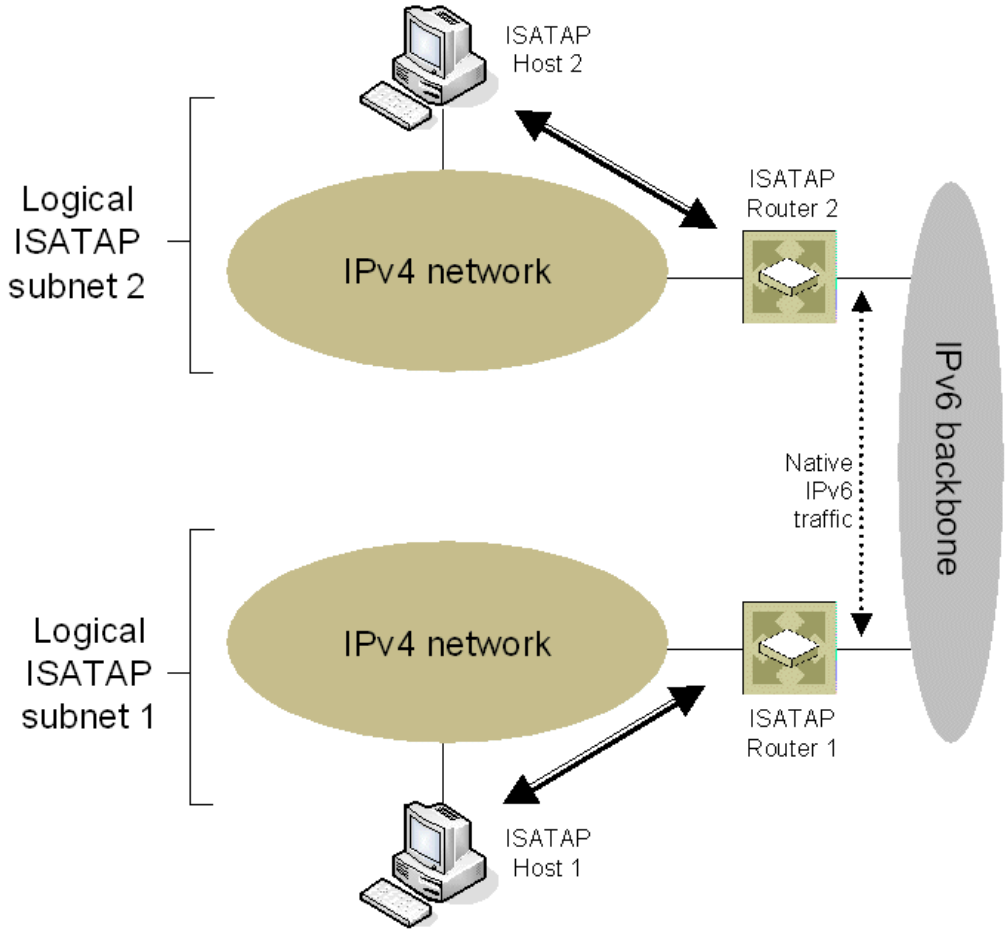
Delivery of traffic between ISATAP hosts and same logical ISATAP subnet shows the delivery of ISATAP traffic between ISATAP hosts on the same logical ISATAP subnet:

Delivery of traffic between ISATAP hosts and same logical ISATAP subnet



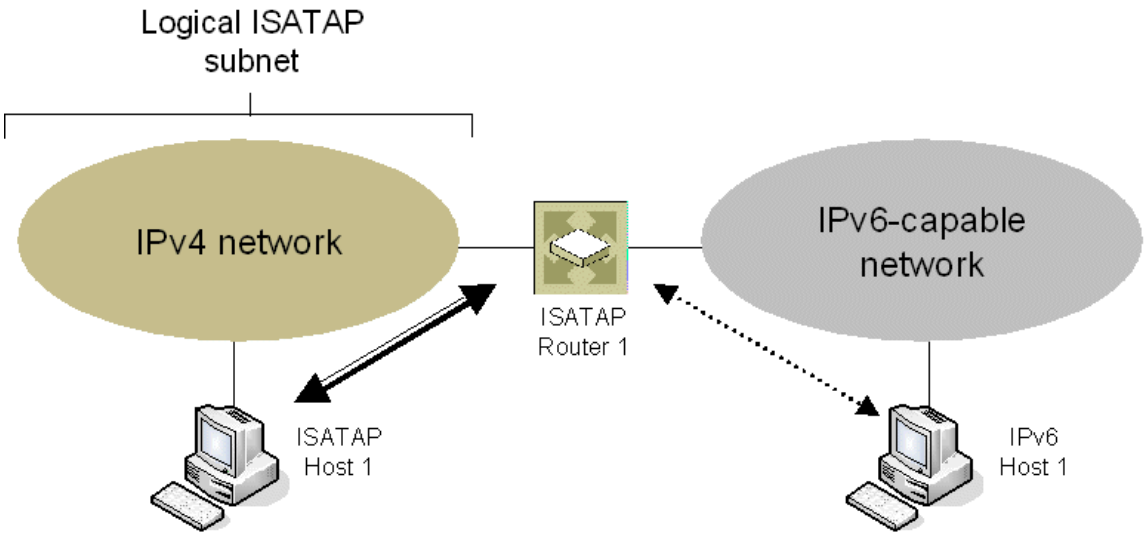
Delivery of traffic between ISATAP hosts and different ISATAP subnets shows the delivery of ISATAP traffic between hosts on different ISATAP subnets:

Delivery of traffic between ISATAP hosts and different ISATAP subnets



Delivery of packets between ISATAP hosts and hosts on IPv6-capable network shows the delivery of packets between ISATAP hosts and hosts on an IPv6-capable network.

Delivery of packets between ISATAP hosts and hosts on IPv6-capable network



In the scenario presented in Delivery of packets between ISATAP hosts and hosts on IPv6-capable network, the ISATAP hosts can communicate directly to each other without going through the ISATAP router or IPv6 network. This allows an IPv6-capable application to leverage connectivity of an existing IPv4 infrastructure.

The other two scenarios require the ISATAP router to have an IPv6 interface connected to the IPv6 network which supports forwarding between the ISATAP interface-facing IPv4 network and the IPv6 interface.

ISATAP needs to be implemented and run in both the host and router. Dual-stack node support is enabled by default on the Windows XP and Windows 7 platforms.

ISATAP support in SonicOS allows the Security Appliance to function as an ISATAP router on LAN- facing interfaces and forward IPv6 packets between the ISATAP tunneling interface and IPv6 interface connected to the IPv6 network.

To configure an ISATAP tunnel:

1. In **MANAGE | System Setup | Network > Interfaces**, at **View IP Version**, select **IPv6**.
2. Click **Add Interface**.
3. In the **General** tab, Select the **Zone** for the tunnel interface.
4. In the **Tunnel Type** drop-down list, select **ISATAP Tunnel Interface**.
5. Enter a **Name** for the tunnel interface.
6. **Bound to IPv4 Address of** - Select an interface from the drop-down menu. The ISATAP tunnel uses the IPv4 address of the bound interface as the IPv4 end address of 6over4 tunnel.
7. **IPv6 Subnet Prefix** - Select an address object from the drop-down menu (or select Create a new address object). The IPv6 subnet prefix is a 64 bit prefix, and is used by ISATAP hosts for ISATAP address auto configuration.
8. **Tunnel Interface Link MTU** - The recommended MTU for the interface link. A value of 0 means firewall does not advertise link MTU for the link.
9. Enter any optional comment text in the **Comment** field. This text is displayed in the **Comment** column of the **Interface** table.
10. If you want to enable remote management of the firewall from this interface, select the supported management protocol(s): **HTTPS**, **Ping**, or **SNMP**.

11 If you want to allow selected users with limited management rights to log in to the Security Appliance, select **HTTP** and/or **HTTPS** in **User Login**.

Additionally, you can specify how SonicOS resolves ISATAP host queries on **MANAGE | Security Configuration | Firewall Settings | Advanced Settings**. For information about configuring advanced firewall settings, see *SonicOS Security Configuration*.

12 Locate the **IPv6 Advanced Configurations** section.

- **Enable NetBIOS name query response for ISATAP** – Select this to if you want the Security Appliance to answer a NetBIOS query in order to help ISATAP hosts resolve the name into an IPv4 address.
- **Resolved name ISATAP is valid for (seconds)** – Enter a time period (in seconds).

Accessing the SonicWall Management Interface Using IPv6

After IPv6 addressing has been configured on the Security Appliance, the SonicWall Management Interface can be accessed by entering the IPv6 of the Security Appliance in your browser's URL field.

IPv6 Network Configuration

Topics:

- IPv6 DNS on page 958
- Address Objects on page 958
- Policy-based Routing on page 959
- IPv6 NAT Policies on page 959
- Neighbor Discovery Protocol on page 959
- DHCPv6 Configuration on page 960

IPv6 DNS

DNS for IPv6 is configured using the same method as for IPv4. Click **IPv6** in **View IP Version** at the top left of **MANAGE | System Setup | Network > DNS**.

Address Objects

IPv6 address objects or address groups can be added in the same manner as IPv4 address objects. For information about configuring address objects, see *SonicOS Policies*.

NOTE: Address Objects of type Host, Range and Network are supported. Dynamic address objects for MAC and FQDN are not currently supported for IPv6 hosts.

IPv4 interfaces define a pair of a default Address Object (DAO) and an Address Object Group for each interface. The basic rule for IPv4 DAO is each IPv4 address corresponds to 2 address objects: Interface IP and Interface Subnet. There are also couples of AO groups for Zone Interface IP, Zone Subnets, All Interface IP, All Interface Management IP, and so forth.

IPv6 interface prepares the same DAO set for each interface. Because multiple IPv6 can be assigned to one interface, all of those address can be added, edited, and deleted dynamically. Therefore, IPv6 DAOs need to be created and deleted dynamically.

To address this, DAOs are not generated dynamically for IPv6 interfaces. Only limited interface DAO are created, which results in limitation support for other module which needs to refer interface DAO.

Policy-based Routing

Policy-based Routing is fully supported for IPv6 by selecting IPv6 address objects and gateways for route policies on **MANAGE | System Setup | Network > Routing**.

Routing Information Protocol next generation (RIPng) is an information routing protocol for IPv6, which allows routers to exchange information for computing routes through an IPv6-based network.

A radio button is added to switch between RIP and RIPng:

IPv6 NAT Policies

NAT policies can be configured for IPv6 or NAT64 on **MANAGE | Policies | Rules |NAT Policies**. When configuring IPv6 NAT policies, the source and destination objects can only be IPv6 address objects unless an IP version of NAT64 is specified. For more information about configuring NAT policies, see *SonicOS Policies*.

NOTE: IPv6 probing for NAT policies is not currently supported.

NAT64 Stateful Inspection Network Streams Support

Stateful inspection network streams (usually including application layer data) need to create cache entries on the fly. These cache entries usually are illegal based on the packet filter's rule table, but they are allowed due to specific directives in the application layer data (for instance, the addition of an inbound cache entry for an FTP data connection).

In SonicOS, these network streams are handled differently from general application layer protocol streams like HTTPS or SNMP. These stateful inspection network streams include FTP, TFTP, H.323, MSN, Oracle, PPTP, RTSP, and RealAudio. Stateful inspection network streams need to anticipate the creation of data cache when client and server communicate with each other through a control channel.

Our system supports FTP (including active and passive mode) and TFTP protocol well for NAT64.

Neighbor Discovery Protocol

The Neighbor Discovery Protocol (NDP) is a new messaging protocol that was created as part of IPv6 to perform a number of the tasks that ICMP and ARP accomplish in IPv4. Just like ARP, Neighbor Discovery builds a cache of dynamic entries, and the administrator can configure static Neighbor Discovery entries. The following table shows the IPv6 neighbor messages and functions that are analogous to the traditional IPv4 neighbor messages.

IPv4 vs. IPv6 neighbor messages

IPv4 neighbor message	IPv6 neighbor message
ARP request message	Neighbor solicitation message
ARP relay message	Neighbor advertisement message
ARP cache	Neighbor cache
Gratuitous ARP	Duplicate address detection
Router solicitation message (optional)	Router solicitation (required)
Router advertisement message (optional)	Router advertisement (required)
Redirect message	Redirect Message

The Static NDP feature allows for static mappings to be created between a Layer 3 IPv6 address and a Layer 2 MAC address.

To configure a Static NDP entry:

1 Navigate to **MANAGE | System Setup | Network > Neighbor Discovery**.

2 Click **Add**.

3 In the **IP Address** field, enter the IPv6 address for the remote device.

4 In the **Interface** drop-down menu, select the interface on the firewall that is used for the entry.

5 In the **MAC Address** field, enter the MAC address of the remote device.

6 Click **OK**. The static NDP entry is added.

The NDP Cache table displays all current IPv6 neighbors. The follow types of neighbors are displayed:

- REACHABLE - The neighbor is known to have been reachable within 30 seconds.
- STALE - The neighbor is no longer known to be reachable, and traffic has been sent to the neighbor within 1200 seconds.
- STATIC - The neighbor was manually configured as a static neighbor.

DHCPv6 Configuration

DHCPv6 server can be configured similar to IPv4 after selecting **IPv6** in **View IP Version** of **MANAGE | System Setup | Network > DNS**.

IPv6 Access Rules Configuration

IPv6 access rules can be configured in the same manner as IPv4 access rules by choosing IPv6 address objects instead of IPv4 address objects. For more information about firewall access rules, see *SonicOS Policies*.

When adding an IPv6 access rule, the source and destination can only be IPv6 address objects.

IPv6 Advanced Firewall Settings

You can configure advanced firewall settings for IPv6, including packet limitations and traffic restrictions on **MANAGE | Security Configuration | Firewall Settings | Advanced Settings**. For information about configuring advanced firewall settings, see *SonicOS Security Configuration*.

IPv6 IPSec VPN Configuration

IPSec VPNs can be configured for IPv6 in a similar manner to IPv4 VPNs after selecting **IPv6** in the **View IP Version** at the top left of **MANAGE | Connectivity | VPN | Settings**. For information about configuring VPN, see *SonicOS Connectivity*.

There are certain VPN features that are currently not supported for IPv6, including:

- IKEv2 is supported, while IKE is currently not supported
- GroupVPN is not supported
- DHCP Over VPN is not supported.

When configuring an IPv6 VPN policy, in **General** of the dialog, the gateways must be configured using IPv6 addresses. FQDN is not supported. When configuring IKE authentication, IPV6 addresses can be used for the local and peer IKE IDs.

On **Network** of the VPN policy, IPV6 address objects (or address groups that contain only IPv6 address objects) must be selected for the **Local Network** and **Remote Network**.

DHCP Over VPN is not supported, thus the DHCP options for protected network are not available.

The **Any address** option for **Local Networks** and the **Tunnel All** option for **Remote Networks** are removed. Select an all zero IPv6 Network address object could be selected for the same functionality and behavior.

On **Proposals**, the configuration is identical for IPv6 and IPv4, except for the fact that IPv6 only support **IKEv2 mode**.

On **Advanced**, only **Enable Keep Alive** and the **IKEv2 Settings** can be configured for IPv6 VPN policies.

NOTE: Because an interface may have multiple IPv6 address, sometimes the local address of the tunnel may vary periodically. If the user needs a consistent IP address, configure the VPN policy to be bound to an interface instead of Zone, and specify the address manually. The address must be one of the IPv6 addresses for that interface.

SSL VPN Configuration for IPv6

SonicOS supports NetExtender connections for users with IPv6 addresses. On **MANAGE | Connectivity | SSL VPN | Client Settings**, first configure the traditional IPv6 IP address pool, and then configure an IPv6 IP Pool. Clients are assigned two internal addresses: one IPv4 and one IPv6. For more information about configuring SSL VPN, see *SonicOS 6.5 Connectivity*.

On the **Edit Device Profile** dialog of **MANAGE | Connectivity | SSL VPN | Client Settings**, you can select a client routes from all address objects, including all the predefined IPv6 address objects.

NOTE: IPv6 FQDN is supported.

IPv6 Visualization

IPv6 Visualization for the App Flow Reports and Live Monitor is an extension of IPv4 Visualization, providing real-time monitoring of interface/application rates and visibility of sessions in the management interface. You can see what websites you employees are accessing, what applications and services are being used in their networks and to what extent, to police content transmitted in and out of your organization. For more information about these visualization tools, refer to *SonicOS 6.5 Investigate* and *SonicOS 6.5 Monitoring*, respectively.

IPv6 Visualization Feature Limitations

Visualization for IPv6 has these feature limitations:

- The IPv6 URL Rating is not supported, because CFS does not support all aspects of IPv6.
- IPv6 Country information is not supported.
- IPv6 External Reporting is not supported.

Configuring IPv6 Visualization

App Flow Reports and Live Monitor visualization are configured the same in IPv6 and IPv4. For more information about these visualization tools, refer to *SonicOS Investigation* and *SonicOS Monitoring*, respectively.

IPv6 High Availability Monitoring

IPv6 High Availability (HA) Monitoring is implemented as an extension of HA Monitoring in IPv4. After configuring HA Monitoring for IPv6, both the primary and backup appliances can be managed from the IPv6 monitoring address, and IPv6 Probing is capable of detecting the network status of HA pairs.

On **MANAGE | System Setup | High Availability > Monitoring Settings**, toggle between the IPv6 and IPv4 views for easy configuration of both IP versions.

Topics:

- IPv6 High Availability Monitoring Feature Limitations on page 962
- IPv6 High Availability Probing on page 962
- Configuring IPv6 High Availability Monitoring on page 962

IPv6 High Availability Monitoring Feature Limitations

The IPv6 HA Monitoring feature limitations are:

- Physical/Link Monitoring property cannot be changed in the IPv6 HA Monitoring configuration page. Set the property in the IPv4 HA Monitoring configuration page.
- Override Virtual MAC property cannot be changed in IPv6 HA Monitoring configuration page. Set the property in the IPv4 HA Monitoring configuration page.
- HA Probing cannot be enabled on both IPv4 and IP{v6 at the same time. That is, if IPv4 probing is enabled, then IPv6 probing must be disabled, and vice versa.

IPv6 High Availability Probing

An ICMPv6 packet is periodically sent out from the primary and backup appliances to probe the IPv6 address, and the response from the probed IPv6 address is monitored. If the active Security Appliance cannot reach the probed IPv6 address, but the idle Security Appliance can, the backup Security Appliance has a better network status and failover initiates.

In IPv6 HA Probing, the IPv6 addresses, ICMPv6 echo requests, and ICMPv6 echo replies are used. The logic used to judge network status of the primary and backup appliance is the same for IPv4 and IPv6.

Configuring IPv6 High Availability Monitoring

The IPv6 HA Monitoring configuration page is inherited from IPv4, so the configuration procedures are almost identical. Just select IPv6 and refer to IPv6 on page 932 for configuration details.

Consider the following when configuring IPv6 HA Monitoring:

- **Physical/Link Monitoring** and **Virtual MAC** are dimmed because they are layer two properties. That is, as the properties are used by both IPv4 and IPv6, you have to configure them in the IPv4 monitoring page.
- The primary/backup IPv6 address must be in the same subnet of the interface, and it can not be same as the global IP or Link-Local-IP of the primary/backup Security Appliance.
- If the primary/backup monitoring IP is set to (not ::), then they cannot be the same.
- If **Management** is enabled, then primary/backup monitoring IP cannot be unspecified (that is, ::).
- If the probe checkbox is enabled, then the probe IP cannot be unspecified.

IPv6 Diagnostics and Monitoring

SonicOS provides a full compliment of diagnostic tools for IPv6, including:

- Packet Monitor on page 963
- IPv6 Ping on page 963
- IPv6 DNS Name Lookup and Reverse Name Resolution on page 963

Packet Monitor

INVESTIGATE | Tools | Packet Monitor fully supports IPv6. In addition, IPv6 keywords can be used to filter the packet capture. For more information about Packet Monitor, see *SonicOS Investigation*.

IPv6 Ping

When pinging a domain name, the tool uses the first IP address that is returned and shows the actual pinging address. If both an IPv4 and IPv6 address are returned, by default, the Security Appliance pings the IPv4 address. The ping tool includes a **Prefer IPv6 networking** option, that when enabled, makes the Security Appliance ping the IPv6 address. For more information about IPv6 Ping, see *SonicOS Investigation*.

IPv6 DNS Name Lookup and Reverse Name Resolution

When performing IPv6 DNS Name Lookup or IPv6 Reverse Name Resolution, you must enter the DNS server address. Either an IPv6 or IPv4 address can be used. For more information about these tools, see *SonicOS Investigation*.

C

BGP Advanced Routing

Topics:

- BGP Advanced Routing on page 964
  - About BGP on page 964
  - Caveats on page 971
  - Configuring BGP on page 971
  - Verifying BGP Configuration on page 981
  - IPv6 BGP on page 984

BGP Advanced Routing

This appendix provides an overview of SonicWall's implementation of Border Gateway protocol (BGP), how BGP operates, and how to configure BGP for your network.

NOTE: BGP is supported on the TZ series appliances with the purchase of a SonicOS Expanded License. BGP is not supported on the SOHO Wireless appliance.

Topics:

- About BGP on page 964
- Caveats on page 971
- Configuring BGP on page 971
- Verifying BGP Configuration on page 981
- IPv6 BGP on page 984

About BGP

Topics:

- What is BGP? on page 965
- Background Information on page 965
- Autonomous Systems on page 966
- BGP over VPN Tunnel Interface on page 967
- Why Use BGP? on page 967
- How Does BGP Work? on page 967
- BGP Terms on page 970

What is BGP?

BGP is a large-scale routing protocol used to communicate routing information between Autonomous Systems (ASs), which are well-defined, separately administered network domains. BGP support allows for SonicWall Security Appliances to replace a traditional BGP router on the edge of a network's AS. The current SonicWall implementation of BGP is most appropriate for single-provider/single-homed environments, where the network uses one ISP as their Internet provider and has a single connection to that provider. SonicWall BGP is also capable of supporting single-provider/multi-homed environments, where the network uses a single ISP but has a small number of separate routes to the provider. BGP is enabled on the **Network > Routing** page of the SonicOS Management Interface, and then it is fully configured through the SonicOS Command Line Interface (CLI).

BGP licensing requirements are shown in BGP Licensing Requirements.

BGP Licensing Requirements

Platform	Additional License Required
SM 9600 / NS*a* 9650	None; BGP is included.
SM 9400 / NS*a* 9450	None; BGP is included.
SM 9200 / NS*a* 9250	None; BGP is included.
NSA 6600 / NS*a* 6650	None; BGP is included.
NSA 5600 / NS*a* 5650	None; BGP is included.
NSA 4600 / NS*a* 4650	None; BGP is included.
NSA 3600 / NS*a* 3650	SonicOS Expanded 01-SSC-7091
NSA 2600 / NS*a* 2650	SonicOS Expanded License
TZ600 / TZ600P	SonicOS Expanded License
TZ500/TZ500W	SonicOS Expanded License
TZ400/TZ400W	SonicOS Expanded License
TZ350/TZ350W	N/A
TZ300/TZ300W / TZ300P	N/A
SOHO 250/250W / SOHO W	N/A

NOTE: Licenses can be purchased on www.MySonicWall.com.

Background Information

Routing protocols are not just packets transmitted over a network, but comprise all the mechanisms by which individual routers, and groups of routers, discover, organize, and communicate network topologies. Routing protocols use distributed algorithms that depend on each participant following the protocol as it is specified, and are most useful when routes within a network domain dynamically change as links between network nodes change state.

Routing protocols typically interact with two databases:

- **Routing Information Base (RIB)** - Used to store all the route information required by the routing protocols themselves.
- **Forward Information Base (FIB)** - Used for actual packet forwarding.

The best routes chosen from the RIB are used to populate the FIB. Both the RIB and FIB change dynamically as routing updates are received by each routing protocol, or connectivity on the device changes.

There are two basic classes of routing protocols:

- **Interior Gateway Protocols (IGPs)** - Interior Gateway Protocols are routing protocols designed to communicate routes within the networks that exist inside of an AS. There are two generations of IGPs. The first generation consists of distance-vector protocols. The second generation consists of link-state protocols. The distance-vector protocols are relatively simple, but have issues when scaled to a large number of routers. The link-state protocols are more complex, but have better scaling capability. The existing distance-vector protocols are Interior Gateway Routing Protocol (IGRP), Enhanced Interior Gateway Routing Protocol (EIGRP), Routing Information Protocol (RIP), and RIPv2, an enhanced version of RIP. IGRP and EIGRP are proprietary Cisco protocols. The link-state protocols currently in use are Open Shortest Path First (OSPF) protocol and the little-used Intermediate System to Intermediate System (IS-IS) protocol.

  SonicOS supports OSPFv2 and RIPv1/v2 protocols, the two most common routing Interior Gateway Protocols, allowing our customers to use our products in their IGP networks and avoid the additional cost of a separate traditional router.

- **Exterior Gateway Protocols (EGPs)** - The standard, ubiquitous Exterior Gateway Protocol is BGP (BGP4, to be exact). BGP is large-scale routing protocol that communicates routing information and policy between well-defined network domains called Autonomous Systems (ASs). An Autonomous System is a separately administered network domain, independent of other Autonomous Systems. BGP is used to convey routes and route policy between Autonomous Systems. ISPs commonly use BGP to convey routes and route policy with their customers as well as with other ISPs.

  Each Autonomous System has a 16-bit number assigned. Like IP addresses, an AS number may be public or private. Public AS numbers are a limited resource and are provisioned based on a number of factors. ISP customers with large networks multi-homed to two or more ISPs usually have a public AS, whereas smaller customers are given a private AS administered by their ISP provider.

  As our products evolve in support of enterprise-level requirements, some customers may want to place our products on the edge of their AS in place of a traditional BGP router.

Autonomous Systems

Each Autonomous System has a 16-bit number assigned. Like IP addresses, an AS number may be public or private. Public AS numbers are a limited resource and are provisioned based on a number of factors. ISP customers with large networks multi-homed to two or more ISPs usually have a public AS, whereas smaller customers are given a private AS administered by their ISP provider.

Types of BGP Topologies

BGP is a very flexible and complex routing protocol. As such, BGP routers may be placed in a large variety of topology settings, such as Internet core routers, intermediary ISP routers, ISP Customer Premises Equipment (CPE), or routers in small private BGP networks. The number of BGP routes required for different topologies varies from greater than 300,000 for core routers, to 0 for ISP customers that use a single ISP and use default routing for all destinations outside of their AS. ISP customers are often required to run BGP from their edge router (the CPE) to the ISP regardless of the number of routes they receive from the ISP. This allows ISP customers to control which networks to advertise to the outside world. There's always the fear that a customer might advertise a network, or network aggregate, not owned by the customer, black-holing Internet traffic to those networks. In reality, ISP providers are careful to filter invalid advertisements from their customers (one of BGP's strengths), so this rarely happens.

There are three basic scales of BGP networks:

- **Single-Provider/Single-Homed** - The network receives a single route (single-homed) from a single ISP (single-provider). The number of routes an ISP customer receives from its ISP depends on the nature of its AS. An ISP customer that uses only one ISP as their Internet provider, and has a single connection to that provider (single-provider / single-homed) has no need to receive any routes - all traffic destined

outside of the AS goes to their ISP. These customers might still advertise some or all of their inside network to the ISP.

- **Single-Provider/Multi-Homed** - The network receives multiple routes (multi-homed) from a single ISP (single-provider). ISP customers that use a single ISP, but have multiple connections to their ISP may only receive the default route (0.0.0.0/0) at each ISP gateway. If an ISP connection goes down, the advertised default route sent from the connected CPE router to internal routers would be withdrawn, and Internet traffic would then flow to a CPE router that has connectivity to the ISP. The customer's inside network would also be advertised to the ISP at each CPE router gateway, allowing the ISP to use alternate paths should a particular connection to a customer go down.
- **Multi-Provider/Multi-Homed** - ISP customers that use more than one ISP (multi-provider / multi-homed) have one or more separate gateway routers for each ISP. In this case, the customer's AS must be a public AS, and may either be a transit or non-transit AS. A transit AS receives and forwards traffic from one ISP destined for a network reachable through another ISP (the traffic destination is not in the customer's AS). A non-transit AS should only receive traffic destined for its AS - all other traffic would be dropped. BGP routers in a transit AS would often receive a large portion (in many cases, all) of the full BGP route table from each ISP.

BGP over VPN Tunnel Interface

BGP interfaces support both numbered and unnumbered tunnel interfaces. This feature is supported on all platforms where BGP and unnumbered tunnel interfaces can be set up.

Why Use BGP?

- Even if you are not a large network on the Internet, BGP is the standard for multi-homing, load-balancing, and redundancy:
  - **Single-provider/Single-homed** – Not typically a strong candidate for BGP, but may still use it to advertise networks to the ISP. single-homed networks are not eligible for a public AS from RIRs.
  - **Single-provider/Multi-homed** – Common to follow RFC2270 suggestion to use a single private AS (64512 to 65535) to get the benefit of BGP while preserving public ASN.
  - **Multi-provider/Multi-homed** – Highly redundant, typically with dedicated routers to each ISP. Requires public ASN. Large memory footprint
- Route summarization makes routing scalable.

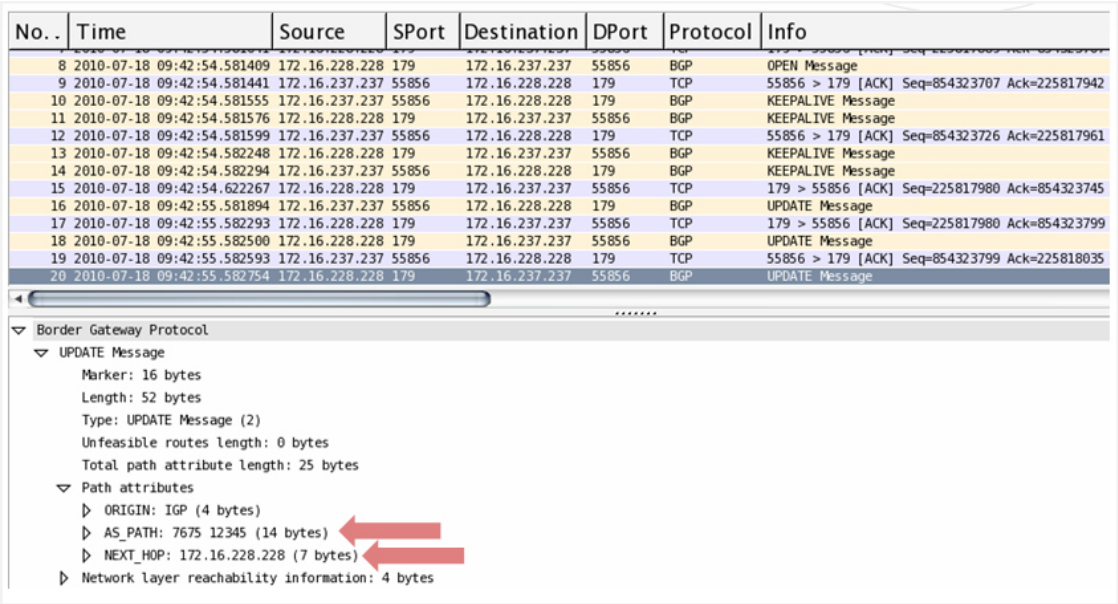
How Does BGP Work?

BGP uses TCP port 179 for communication. BGP is considered a path-vector protocol, containing end-to-end path descriptions for destinations. BGP neighbors can either be internal (iBGP) or external (eBGP):

- **iBGP** – Neighbor is in the same AS.
- **eBGP** – Neighbor is in a different AS.

Paths are advertised in UPDATE messages that are tagged with various path attributes. AS_PATH and NEXT_HOP are the two most important attributes that describe the path of a route in a BGP update message.

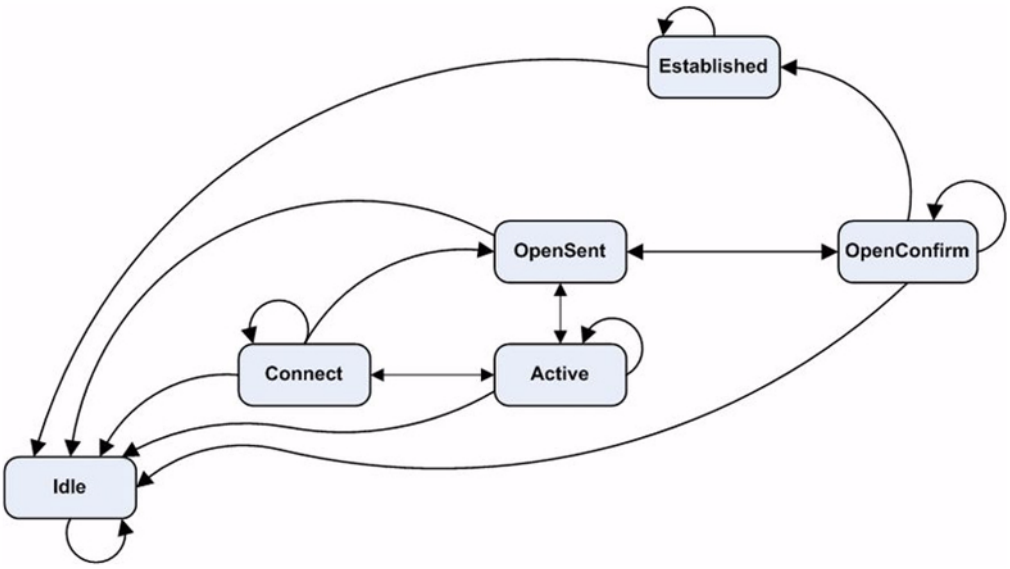
- **AS_PATH**: Indicates the ASs that the route is traveling from and two. In the example below, the AS_PATH is from AS 7675 to AS 12345. For internal BGP, the AS_PATH specifies the same AS for both the source and destination.
- **NEXT_HOP**: Indicates the IP address of the next router the path travels to. Paths advertised across AS boundaries inherit the NEXT_HOP address of the boundary router. BGP relies on interior routing protocols to reach NEXT_HOP addresses.



BGP Finite State Machine

RFC 1771, which defines BGP, describes the operation of BGP in terms of the following state machine. The table following the diagram provides additional information on the various states.

BGP finite state machine



BGP finite state descriptions

State	Description
Idle	Waiting for Start event, after establishing new BGP session or resetting an existing session. In the event of errors, falls back to the Idle state. After a Start event, BGP initializes, resets connect retry timer, initiates TCP transport connection, and listens for connections
Connect	After the TCP layer is up, transition to OpenSent, and send OPEN. If no TCP, transition to Active. If the connect retry timer expires, remain in Connect, reset the timer, and initiate a transport connection. Otherwise, transition back to Idle.

BGP finite state descriptions (Continued)

State	Description
Active	Try to establish TCP connection with peer. If successful, transition to OpenSent and send OPEN. If connect retry expires, restart the timer and fall back to the Connect state. Also actively listen for connection by another peer. Go back to Idle in case of other events. Connect to Active flapping indicates a TCP transport problem, for example, TCP retransmissions or unreachability of a peer.
OpenSent	Waiting for OPEN message from peer. Validate on receipt. On validation failure, send NOTIFICATION and go to Idle. On success, send KEEPALIVE and reset the keepalive timer. Negotiate hold time, smaller value wins. If zero, hold timer and keepalive timer are not restarted.
OpenConfirm	Wait for KEEPALIVE or NOTIFICATION. If KEEPALIVE is received, transition to Established. If UPDATE or KEEPALIVE is received, restart the hold timer (unless the negotiated hold time is zero). If NOTIFICATION is received, transition to Idle. Periodic KEEPALIVE messages are sent. If TCP layer breaks, transition to Idle. If an error occurs, send a NOTIFICATION with error code, transition to Idle.
Established	Session up, exchange updates with peers. If a NOTIFICATION is received, transition to Idle. Updates are checked for errors. On error, send NOTIFICATION, and transition to Idle. In case of hold time expiration, disconnect TCP.

BGP Messages

BGP communication includes the following types of messages:

- **Open** – The first message between BGP peers after TCP session establishment. Contains the necessary information to establish a peering session, for example, ASN, hold time, and capabilities such as multi-product extensions and route-refresh.
- **Update** – These messages contain path information, such as route announcements or withdrawals.
- **Keepalive** – Periodic messages to keep TCP layer up, and to advertise liveliness.
- **Notification** – A request to terminate the BGP session. Non-fatal notifications contain the error code “cease”. Subcodes provide further detail, as shown in Notification subcodes.

  Notification subcodes

Subcode	Description
1 – Maximum number of prefixes reached	The configured “neighbor maximum-prefix” value was exceeded
2 – Administratively shutdown	Session was administratively shutdown
3 – Peer unconfigured	Peer configuration has been removed
4 – Administratively reset	Session was administratively reset
5 – Connection rejected	Rejection (sometimes temporary) of BGP session
6 – Other configuration change	Session was administratively reset for some reason

- **Route-refresh** – A request for the peer to resend its routes.

BGP Attributes

BGP update messages can include the attributes shown in BGP update message attributes:

BGP update message attributes

Value	Code
1	ORIGIN
2	AS_PATH
3	NEXT_HOP
4	MULTI_EXIT_DISC
5	LOCAL_PREF
6	ATOMIC_AGGREGATE
7	AGGREGATOR
8	COMMUNITY
9	ORIGINATOR_ID
10	CLUSTER_LIST
11	DPA
12	ADVERTISER (Historic)
13	RCID_PATH / CLUSTER_ID (Historic)
14	MP_REACH_NLRI
15	MP_UNREACH_NLRI
16	EXTENDED COMMUNITIES
17	AS4_PATH
18	AS4_AGGREGATOR
19	SAFI Specific Attribute (SSA) (deprecated)
20	Connector Attribute (deprecated)
21	AS_PATHLIMIT (deprecated)
22	PMSI_TUNNEL
23	Tunnel Encapsulation Attribute
24	Traffic Engineering
25	IPv6 Address Specific Extended Community
26	AIGP (TEMPORARY - expires 2011-02-23)
27-254	Unassigned
255	Reserved for development

For more information on BGP attributes, see: http://www.iana.org/assignments/bgp-parameters/bgp-parameters.xml

BGP Terms

ARD	Autonomous Routing Domain – A collection of networks/routers that have a common administrative routing policy.
AS	Autonomous System – An ARD that has been assigned an identifying number, typically running BGP4 at its border router(s).
BGP4	Border Gateway Protocol 4: The most prevalent EGP.
CIDR	Classless inter-domain routing, enables efficient route advertisement through route aggregation.

CPE	Customer Premise Equipment - The equipment at the edge of a customer's network used to interface with the ISP.
EGP	Exterior Gateway Protocol – Any protocol (in practice, BGP4) used to communicate routing information between Autonomous Systems.
Full-Routes	The entire global BGP route table.
FIB	Forwarding Information Base – Our existing route table, used to find the egress interface and next hop when forwarding packets.
Looking Glass*	A Looking Glass (LG) server is a read-only view of routers of organizations running the LG servers. Typically, publicly accessible looking glass servers are run by ISPs or NOCs.
Multi-Homed	An ISP customer that has multiple connections to one or more ISPs.
Multi-Provider	An ISP customer that uses multiple ISPs to connect to the Internet.
NSM	Network Services Module - The ZebOS component that centralizes the interface to the FIB and RIB. The separate routing protocol daemons interface with the NSM for all RIB updates. NSM alone updates the FIB with best-route information from the RIB.
Partial Routes	A subset of the full BGP route table, usually specific to destinations that are part of an ISP's domain.
RIB	Route Information Base – A run-time database owned by the NSM, and used to store all route information gathered and used by the routing protocols.

Caveats

Scale	Currently, SonicOS supports from 512 to 2,048 policy-based routes (PBRs). This is not sufficient for full or even partial routing tables. The number of routes that exist in the RIB may be greater than the number installed into PBR (which is the FIB). This occurs when multiple competing routes have been received through the routing protocols. For each case in which the RIB contains competing routes to a particular network destination, only one of these routes is chosen to be installed in the FIB. Currently, our implementation is most appropriate for the single-provider/single-homed customers. Single-provider/multi-homed installations may also be appropriate when either the default route is being received from the ISP, or a very small number of ISP-specific routes are received by the customer. The latter allows inside routers to take the optimal path to destinations outside of the AS, but still within the ISP's network domain (this is called partial-routes).
Load balancing	here is currently no multi-path support in SonicOS or Zebos (the maximum-paths capability). This precludes load-balancing without splitting networks.
Loopback	There is currently no loopback interface support.
NAT	BGP is for routing. It does not co-exist well with NAT.
Asymmetric paths	Stateful Security Appliance does not currently handle asymmetric paths, especially not across multiple Security Appliances.

Configuring BGP

Topics:

- IPSec Configuration for BGP on page 972
- Basic BGP Configuration on page 973

- BGP Path Selection Process on page 974
- AS_Path Prepending on page 977
- Multiple Exit Discriminator (MED) on page 977
- BGP Communities on page 978
- Synchronization and Auto-Summary on page 979
- Preventing an Accidental Transit AS on page 979
- Using Multi-Homed BGP for Load Sharing on page 980

IPSec Configuration for BGP

BGP transmits packets in the clear. Therefore for strong security, SonicWall recommends configuring an IPSec tunnel to use for BGP sessions. The IPSec tunnel and BGP are configured independently of each other. For information about configuring an IPSec tunnel for BGP, see *SonicOS Connectivity*.

To configure an IPSec tunnel for BGP:

1 The IPSec tunnel is configured completely within the **Manage | Connectivity | VPN** configuration section of the SonicOS Management Interface. When configuring the IPSec tunnel, ensure that these options are set:

Option	Value
Policy Type	**Site to Site** **NOTE:** A site-to-site VPN tunnel must be used for BGP over IPSec.
IPsec Primary Gateway Name or Address	IP address of the remote peer
Local IKE ID	IP address of the SonicWall Security Appliance
Peer IKE ID	IP address of the remote peer
Network | Choose destination network from list	Remote peer's IP address
Advanced | Enable Keep Alive	Enable

IMPORTANT: When configuring BGP over IPSec:
1 Configure the IPSec tunnel.
2 Verify connectivity over the tunnel before configuring BGP.

NOTE: For how to configure VPN policies, see *SonicOS Connectivity*.

2 Enable BGP on **MANAGE | System Setup | Network > Routing** page by selecting **BGP** for the **Service** option when adding a route policy. For how to add a route policy, see Configuring BGP Advanced Routing on page 490; for basic BGP configuration, see Basic BGP Configuration on page 973.

3 Finish configuring the route through the SonicOS Command Line Interface.

4 When the VPN policy is configured on the Security Appliance, complete the corresponding IPSec configuration on the remote peer.

5 When the IPSec configuration on the remote peer is complete, return to **MANAGE | Connectivity | VPN | Base Settings**, and enable the VPN policy to initiate the IPSec tunnel.

6 Use the ping diagnostic on the SonicWall Security Appliance to ping the BGP peer IP address. For more information about the ping diagnostic, see *SonicOS Investigation*.

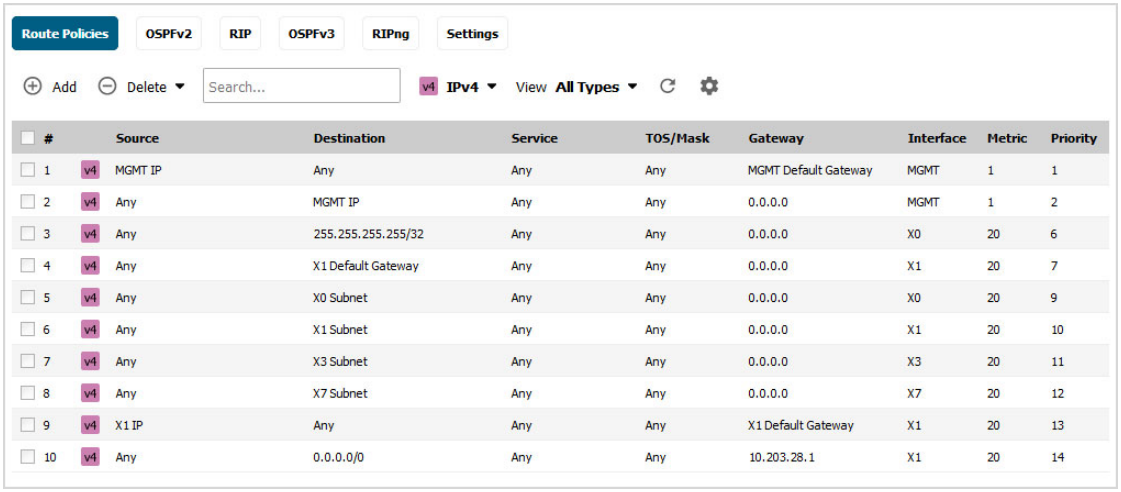
7 Use Wireshark to ensure that the request and response are being encapsulated in ESP packets.

NOTE: As configured in this example, routed traffic does not go through the IPSEC tunnel used for BGP. That traffic is sent and received in the clear, which is most likely the desired behavior as the goal is to secure BGP, not all the routed network traffic.

Basic BGP Configuration

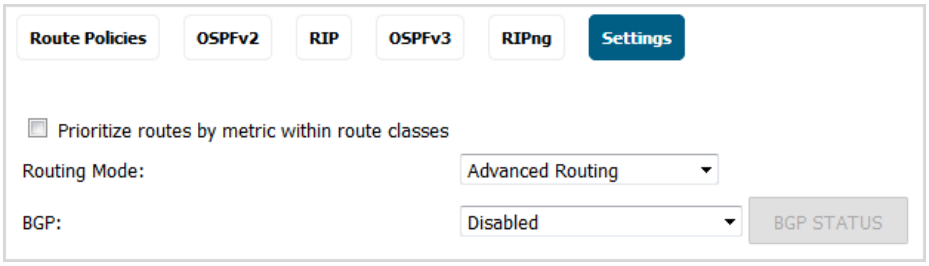
To configure BGP on a SonicWall Security Appliance:

1 Navigate to **MANAGE | System Setup | Network > Routing**.



#		Source	Destination	Service	TOS/Mask	Gateway	Interface	Metric	Priority
1	v4	MGMT IP	Any	Any	Any	MGMT Default Gateway	MGMT	1	1
2	v4	Any	MGMT IP	Any	Any	0.0.0.0	MGMT	1	2
3	v4	Any	255.255.255.255/32	Any	Any	0.0.0.0	X0	20	6
4	v4	Any	X1 Default Gateway	Any	Any	0.0.0.0	X1	20	7
5	v4	Any	X0 Subnet	Any	Any	0.0.0.0	X0	20	9
6	v4	Any	X1 Subnet	Any	Any	0.0.0.0	X1	20	10
7	v4	Any	X3 Subnet	Any	Any	0.0.0.0	X3	20	11
8	v4	Any	X7 Subnet	Any	Any	0.0.0.0	X7	20	12
9	v4	X1 IP	Any	Any	Any	X1 Default Gateway	X1	20	13
10	v4	Any	0.0.0.0/0	Any	Any	10.203.28.1	X1	20	14

2 Click **Settings**.



3 From **Routing Mode**, select **Advanced Routing**.

4 From **BGP**, select **Enabled (Configure with CLI)**. A confirmation message displays.



NOTE: After BGP has been enabled through the Management Interface, the specifics of the BGP configuration are performed using the SonicOS command line interface (CLI).

5 Log in to the SonicOS CLI through the console interface.

6 Enter configuration mode by typing the **`configure`** command.

7 Enter the BGP CLI by typing the **`configure routing bgp`** command. This prompt displays:

```
ZebOS version 7.7.0 IPIRouter 7/2009

ARS BGP>
```

8 You are now in BGP Non-Config Mode. Type **?** to see a list of non-config commands.

9 Type **show running-config** to see the current BGP running configuration.

10 To enter BGP Configuration Mode, type the **configure terminal** command. Type **?** to see a list of configuration commands.

11 When you have completed your configuration, type the **write file** command. If the unit is part of a High Availability pair or cluster, the configuration changes are automatically conveyed to the other unit or units.

BGP Path Selection Process

BGP path selection process attributes describes the attributes used to configure the BGP path selection process.

BGP path selection process attributes

Attribute	Description
Weight	Prefer routes learned from neighbors with the highest weight set. Only relevant to the local router.
Local Preference	Administratively prefer routes learned from a neighbor. Shared with the whole AS.
Network or **Aggregate** paths	Prefer paths that were locally originated from the **network** and **aggregate-address** commands.
AS_PATH	Prefer the path with the shortest AS_PATH.
Origin	Prefer the path with the lowest origin type (as advertised in UPDATE messages): `IGP < EGP < Incomplete`.
Multi Exit Discriminator (MED)	Provides path preference information to neighbors for paths into originating AS.
Recency	Prefer the most recently received path.
Router ID	Prefer the path from the router with the lower router ID.

Weight

The weight command assigns a weight value, per address-family, to all routes learned from a neighbor. The route with the highest weight gets preference when the same prefix is learned from more than one peer. The weight is relevant only to the local router.

The weights assigned using the **set weight** command override the weights assigned using this command.

When the weight is set for a peer-group, all members of the peer-group have the same weight. The command can also be used to assign a different weight to a particular peer-group member.

This example shows weight configuration:

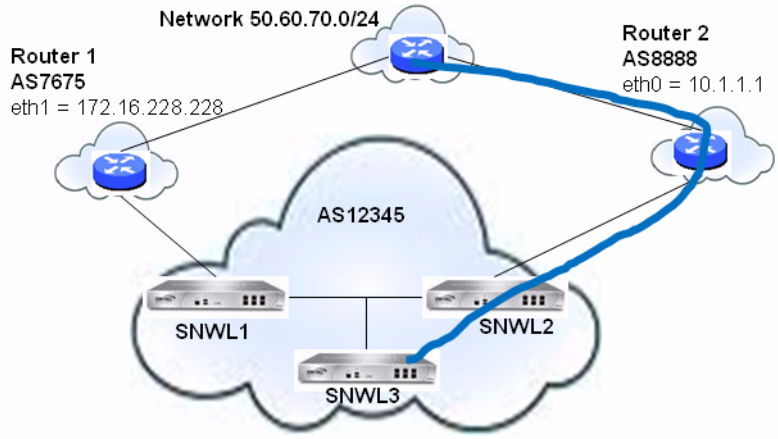
```
router bgp 12345
   neighbor 12.34.5.237 remote-as 12345
   neighbor 12.34.5.237 weight 60


router bgp 12345
   neighbor group1 peer-group
   neighbor 12.34.5.237 peer-group group1
   neighbor 67.78.9.237 peer-group group1
   neighbor group1 weight 60
```

Local Preference

The Local Preference attribute is used to indicate the degree of preference for each external route in an appliance's routing table. The Local Preference attribute is included in all update messages sent to devices in the same AS. Local Preference is not communicated to outside AS. BGP local preference topology shows a sample topology illustrating how Local Preference affects routes between neighboring ASs.

BGP local preference topology



The BGP configurations shown in SNWL1 and SNWL2 configurations are entered on SNWL1 and SNWL2. The higher Local Preference on SNWL2 leads to SNWL2 being the preferred route advertised by AS 12345 (the SonicWall AS) to outside ASs.

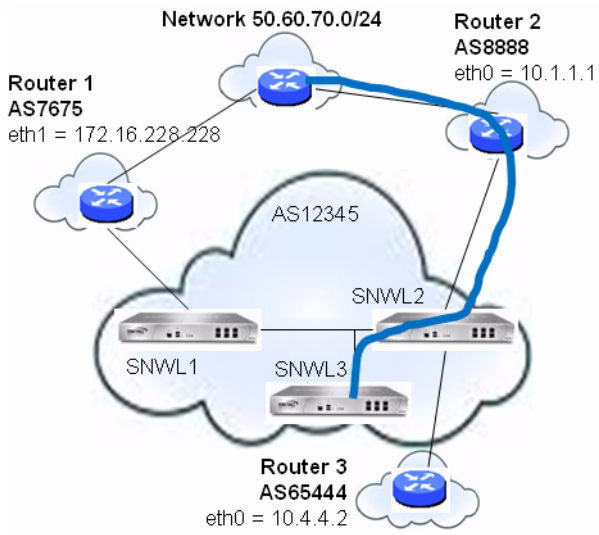
SNWL1 and SNWL2 configurations

SNWL1 Configuration	SNWL2 Configuration
x0 = 12.34.5.228	x0 = 12.34.5.237
x1 = 172.16.228.45	x1 = 10.1.1.2
----------	----------
router bgp 12345	router bgp 12345
neighbor 172.16.228.228 remote-as 7675	neighbor 10.1.1.1 remote-as 8888
neighbor 12.34.5.237 remote-as 12345	neighbor 12.34.5.228 remote-as 12345
bgp default local-preference 150	bgp default local-preference 200

Local Preference used with Route Maps

Route Maps are similar to Access Control Lists. They consist of a series of Permit and/or Deny statements that determine how the appliance processes the routes. Route maps are applied to inbound traffic—not outbound traffic. BGP local preference topology with route maps shows a sample topology that uses a route map to configure local preference.

BGP local preference topology with route maps



The BGP configurations shown in SNWL1 and SNWL2 configurations with route maps are entered on SNWL1 and SNWL2.

SNWL1 and SNWL2 configurations with route maps

SNWL1 Configuration	SNWL2 Configuration
x1 = 172.16.228.45	x0 = 12.34.5.237 x1 = 10.1.1.2 x4 = 10.4.4.1
----------	----------
router bgp 12345 neighbor 172.16.228.228 remote-as 7675 neighbor 12.34.5.237 remote-as 12345 bgp default local-preference 150	router bgp 12345 neighbor 10.1.1.1 remote-as 9999 neighbor 10.1.1.1 route-map rmap1 in neighbor 12.34.5.237 remote-as 12345 .... ip as-path access-list 100 permit ^8888$ ... route-map rmap1 permit 10 match as-path 100 set local-preference 200 route-map rmap1 permit 20 set local-preference 150

The Route Map configured on SNWL2 (rmap1) is configured to apply to inbound routes from neighbor `10.1.1.1`. It has two permit conditions:

- **route-map rmap1 permit 10**: This permit condition matches access list 100 that is configured to permit traffic from AS 8888 and set routes from AS 8888 to a Local Preference of 200.
- **route-map rmap1 permit 10**: This permit condition sets all other traffic that does not match access list 100 (that is, traffic coming from ASs other than 8888) to a Local Preference of 150.

AS_Path Prepending

AS_Path Prepending is the practice of adding additional AS numbers at the beginning of a path update. This makes the path for this route longer, and thus decreases its preference.

AS_Path Prepending can be applied on either outbound or inbound paths. AS_Path Prepending may not be honored if it is over-ruled by a neighbor.

Outbound and Inbound path configurations

Outbound Path Configuration	Inbound Path Configuration
router bgp 12345	router bgp 7675
bgp router-id 10.50.165.233	bgp router-id 10.50.165.228
network 12.34.5.0/24	network 7.6.7.0/24
neighbor 10.50.165.228 remote-as 7675	neighbor 10.50.165.233 remote-as 12345
neighbor 10.50.165.228 route-map long out	neighbor 10.50.165.233 route-map prepend in
!	!
route-map long permit 10	route-map prepend permit 10
set as-path prepend 12345 12345	set as-path prepend 12345 12345

This configuration leads to a route being installed to the neighbor `10.50.165.233` with the AS_Path Prepended as `12345 12345`. This can be viewed by entering the **show ip bgp** command.

```
ARS BGP>show ip bgp
BGP table version is 98, local router ID is 10.50.165.228
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
              S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
*> 12.34.5.0/24     10.50.165.233            0             0 12345 12345 12345 i
*> 7.6.7.0/24       0.0.0.0                       100  32768 i
Total number of prefixes 2
```

Multiple Exit Discriminator (MED)

The **set metric** command can be used in a route map to make paths more or less preferable:

```
router bgp 7675
network 7.6.7.0/24
   neighbor 10.50.165.233 remote-as 12345
   neighbor 10.50.165.233 route-map highmetric out
!
route-map highmetric permit 10
   set metric 300
```

The Multi Exit Discriminator (MED) is an optional attribute that can be used to influence path preference. It is non-transitive, meaning it is configured on a single appliance and not advertised to neighbors in update messages. In this section, consider the uses of the bgp always-compare-med command on page 977 and bgp deterministic-med command on page 978.

bgp always-compare-med command

The **bgp always-compare-med** command allows comparison of the MED values for paths from different ASs for path selection. A path with lower MED is preferred.

As an example, consider the following routes in the BGP table and the **always-compare-med** command is enabled:

```
Route1: as-path 7675, med 300
Route2: as-path 200, med 200
Route3: as-path 7675, med 250
```

Route2 would be the chosen path because it has the lowest MED.

If the **always-compare-med** command was disabled, MED would not be considered when comparing Route1 and Route2 because they have different AS paths. MED would be compared for only Route1 and Route3.

bgp deterministic-med command

The selected route is also affected by the **bgp deterministic-med** command, which compares MED when choosing among routes advertised by different peers in the same autonomous system.

When the **bgp deterministic-med** command is enabled, routes from the same AS are grouped together, and the best routes of each group are compared. If the BGP table showed:

```
Route1: as-path 200, med 300, internal
Route2: as-path 400, med 200, internal
Route3: as-path 400, med 250, external
```

BGP would have a group of Route1 and a second group of Route2 and Route3 (the same AS).

The best of each group is compared. Route1 is the best of its group because it is the only route from AS 200.

Route1 is compared to the Route2, the best of group AS 400 (the lower MED).

As the two routes are not from the same AS, the MED is not considered in the comparison. The external BGP route is preferred over the internal BGP route, making Route3 the best route.

BGP Communities

A community is a group of prefixes that share some common property and can be configured with the transitive BGP community attribute. A prefix can have more than one community attribute. Routers can act on one, some or all the attributes. BGP communities can be thought of as a form of tagging. The following is an example of a BGP communities configuration.

```
router bgp 12345
   bgp router-id 10.50.165.233
   network 12.34.5.0/24
   network 23.45.6.0/24
   neighbor 10.50.165.228 remote-as 7675
   neighbor 10.50.165.228 send-community
   neighbor 10.50.165.228 route-map comm out
!
access-list 105 permit 12.34.5.0/24
access-list 110 permit 23.45.6.0/24
!
route-map comm permit 10
   match ip address 105
   set community 7675:300
!
route-map comm permit 20
   match ip address 110
   set community 7675:500
!
router bgp 7675
   bgp router-id 10.50.165.228
   network 7.6.7.0/24
   neighbor 10.50.165.233 remote-as 12345
```

```
   neighbor 10.50.165.233 route-map shape in
!
ip community-list 1 permit 7675:300
ip community-list 2 permit 7675:500
!
route-map shape permit 10
   match community 1
   set local preference 120
route-map shape permit 20
   match community 2
   set local preference 130
```

Synchronization and Auto-Summary

The synchronization setting controls whether the router advertises routes learned from an iBGP neighbor based on the presence of those routes in its IGP. When synchronization is enabled, BGP only advertises routes that are reachable through OSPF or RIP (the Exterior Gateway Protocols as opposed to BGP, the Exterior Gateway Protocol). Synchronization is a common cause of BGP route advertisement problems.

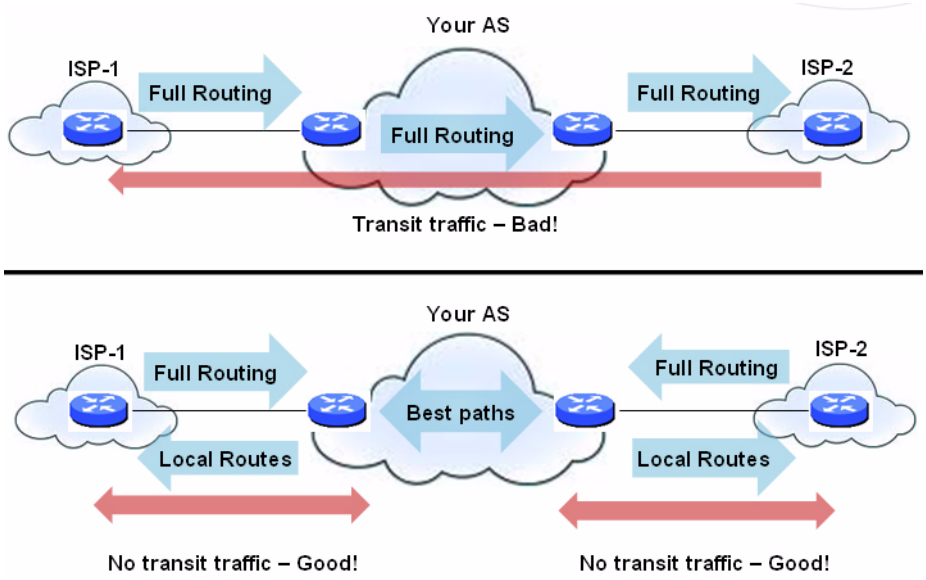
The auto-summary setting controls whether or not routes are advertised classfully. Auto-summary is another common cause of BGP configuration problems

By default, auto-summary and synchronization are disabled on Zebos.

Preventing an Accidental Transit AS

As discussed earlier, an AS peer can either be a transit peer (allowing traffic from an outside AS to another outside AS) or a non-transit peer (requiring all traffic to either originate or terminate on its AS). See Transit peers vs. Non-transit peers. Transit peers have dramatically larger routing tables. Typically, you do not want to configure a SonicWall Security Appliance as a transit peer.

Transit peers vs. Non-transit peers



To prevent your Security Appliance from inadvertently becoming a transit peer, configure inbound and outbound filters:

- Outbound Filters on page 980
- Inbound Filters on page 980

Outbound Filters

Permit only routes originated from the local AS out:

```
ip as-path access-list 1 permit ^$


router bgp 12345
   bgp router-id 10.50.165.233
   network 12.34.5.0/24
   neighbor 10.50.165.228 remote-as 7675
   neighbor 10.50.165.228 filter-list 1 out
   neighbor 172.1.1.2 remote-as 9999
   neighbor 10.50.165.228 filter list 1 out
```

Permit only owned prefixes out:

```
ip prefix-list myPrefixes seq 5 permit 12.34.5.0/24
ip prefix-list myPrefixes seq 10 permit 23.45.6.0/24


router bgp 12345
   bgp router-id 10.50.165.233
   network 12.34.5.0/24
   network 23.45.6.0/24
   neighbor 10.50.165.228 remote-as 7675
   neighbor 172.1.1.2 remote-as 9999
   neighbor 10.50.165.228 prefix-list myPrefixes out
   neighbor 172.1.1.2 prefix-list myPrefixes out
```

Inbound Filters

Drop all owned and private inbound prefixes.

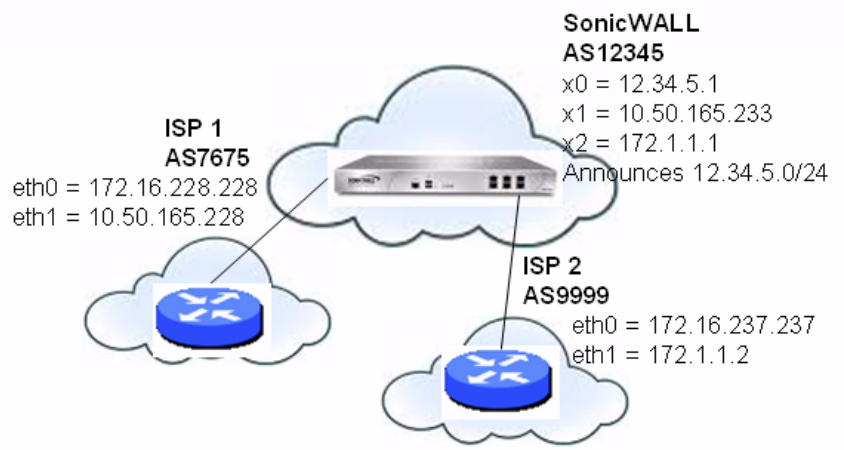
```
ip prefix-list unwantedPrefixes seq 5 deny 12.34.5.0/24 le 32
ip prefix-list unwantedPrefixes seq 10 deny 23.45.6.0/24 le 32
ip prefix-list unwantedPrefixes seq 20 deny 10.0.0.0/8 le 32
ip prefix-list unwantedPrefixes seq 21 deny 172.16.0.0/12 le 32
ip prefix-list unwantedPrefixes seq 22 deny 192.168.0.0/16 le 32
ip prefix-list unwantedPrefixes seq 30 permit 0.0.0.0/0 le 32


router bgp 12345
   bgp router-id 10.50.165.233
   network 12.34.5.0/24
   network 23.45.6.0/24
   neighbor 10.50.165.228 remote-as 7675
   neighbor 172.1.1.2 remote-as 9999
   neighbor 10.50.165.228 prefix-list unwantedPrefixes in
   neighbor 172.1.1.2 prefix-list unwantedPrefixes in
```

Using Multi-Homed BGP for Load Sharing

The topology shown in Multi-homed BGP for load sharing topology is an example where a SonicWall Security Appliance uses a multi-homed BGP network to load share between two ISPs.

Multi-homed BGP for load sharing topology



The SonicWall Security Appliance is configured as follows:

```
router bgp 12345
   bgp router-id 10.50.165.233
   network 12.34.5.0/24
   neighbor 10.50.165.228 remote-as 7675
   neighbor 10.50.165.228 route-map ISP1 out
   neighbor 172.1.1.2 remote-as 9999
   neighbor 10.50.165.228 route-map ISP2 out
!
route-map ISP1 permit 10
match ip address 1
set weight 100

route-map ISP1 permit 20
match ip address 2

route-map ISP2 permit 10
match ip address 1


route-map ISP2 permit 20
match ip address 2
set weight 100


access-list 1 permit 12.34.5.0/25
access-list 2 deny 12.34.5.0/25
access-list 2 permit any
```

Verifying BGP Configuration

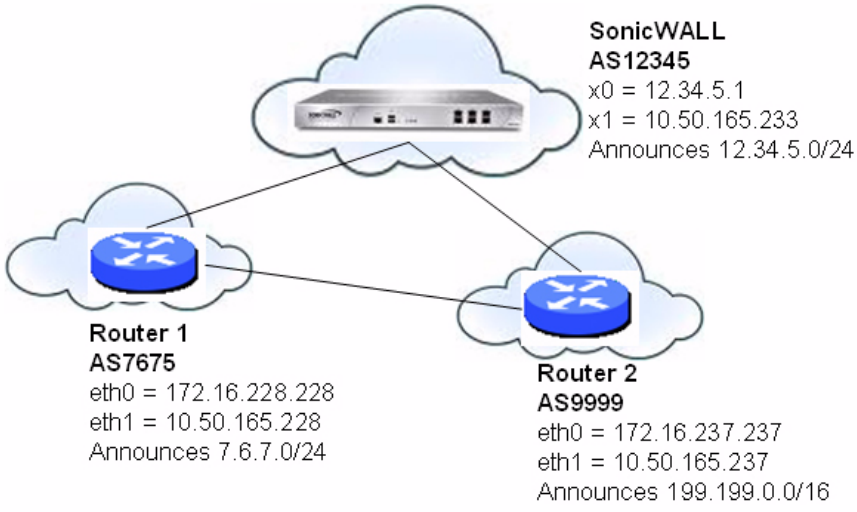
Topics:

- Viewing BGP Routes on page 981
- Configuring BGP Debug and Log on page 983

Viewing BGP Routes

BGP topology shows a basic BGP topology where a SonicWall Security Appliance is configured for BGP to connect to two routers on two different ASs.

BGP topology



The routes in the FIB for this network can be viewed either in the SonicOS Management Interface or by using the CLI.

Topics:

- Viewing FIB routes in the Management Interface on page 982
- Viewing FIB Routes in the CLI on page 982
- Viewing RIB Routes in the CLI on page 983

Viewing FIB routes in the Management Interface

A summary of the BGP configuration can be viewed on the SonicOS Management Interface through **MANAGE | System Setup | Network > Routing > Settings** by clicking **BGP STATUS**. The **BGP Status** dialog displays the output of the **`show ip bgp summary`** and **`show ip bgp neighbor`** commands.

The BGP routes in the FIB can also be viewed through the CLI as described in Viewing FIB Routes in the CLI on page 982.

Viewing FIB Routes in the CLI

To view the FIB routes in the CLI:

```
SonicWall> configure
(config[SonicWall])> route ars-nsm

ZebOS version 7.7.0 IPIRouter 7/2009
ARS NSM>show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default

B       7.6.7.0/24 [20/0] via 10.50.165.228, X1, 05:08:31
B       199.199.0/16 [20/0] via 10.50.165.237, X1, 05:08:31
C       10.50.165.192/26 is directly connected, X1
```

```
C       127.0.0.0/8 is directly connected, lo0
C       12.34.5.0/24 is directly connected, X0
```

Viewing RIB Routes in the CLI

To view the RIB routes in the CLI:

```
ARS BGP>show ip bgp
BGP table version is 98, local router ID is 10.50.165.233
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
              S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
*> 7.6.7.0/24       10.50.165.228            0             0 7675 i
*> 12.34.5.0/24     0.0.0.0                        100  32768 i
*> 199.199.0.0/16   10.50.165.228            0             0 7675 9999 i
Total number of prefixes 3
```

NOTE: The last route is the path to AS9999 that was learned through AS7675.

Configuring BGP Debug and Log

SonicWall BGP offers a comprehensive selection of debug commands to display log events related to BGP traffic. BGP logging can be configured on the CLI by using the **`debug bgp`** command followed by one of the keywords shown in BGP debug keywords.

BGP debug keywords

BGP Debug Keywords	Enables
all	All BGP debugging.
dampening	Debugging for BGP dampening.
events	Debugging for BGP events.
filters	Debugging for BGP filters.
fsm	Debugging for BGP Finite State Machine (FSM).
keepalives	Debugging for BGP keepalives.
nht	Debugging for NHT messages.
nsm	Debugging for NSM messages.
updates	Debugging for inbound/outbound BGP updates.

To disable BGP debugging, enter the “no” form of the command. For example, to disable event debugging, type the **`no debug events`** command.

BGP log messages can also be viewed on the SonicOS GUI on **MANAGE | Investigate | Logs | Event Logs**. BGP messages are displayed as part of the **Advanced Routing** category of log messages. For more information about logs and logging, see *SonicOS Logs and Reporting*.

To allow for BGP peers that are not directly connected, use the **`ebgp-multihop`** keyword with the **`neighbor`** command. For example:

```
neighbor 10.50.165.228 ebgp-multihop
```

IPv6 BGP

IPv6 Border Gateway protocol (BGP) communicates IPv6 routing information between Autonomous Systems (ASs). A SonicWall Security Appliance with IPv6 BGP support can replace a traditional BGP router on the edge of a network's AS.

IPv6 BGP is enabled on **MANAGE | System Setup | Network > Routing**, but must be configured on the SonicOS Command Line Interface (CLI).

The following restrictions apply:

- IPv6 BGP is supported only on NSA platforms.
- IPv6 BGP depends on IPv6 functions and ZebOS (Zebra OS).
- MPLS/VPN and multicast are not supported in IPv6 BGP.

Topics:

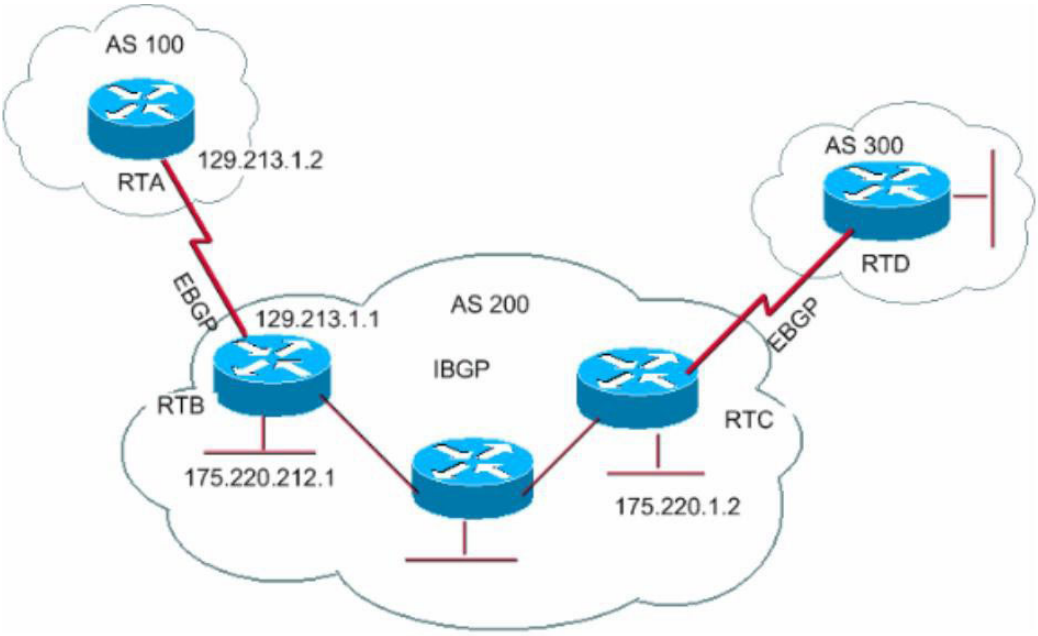
- Configuring Multiple Autonomous Systems on page 984
- Configuring Basic BGP over IPv6 on page 985
- Configuring EBGP Multihop on page 986
- Configuring IPv6 BGP Outbound Route Filter on page 987
- Configuring IPv6 BGP Distribute List on page 988
- IPv6 BGP Route-Map on page 988
- Configuring an AS Regular Expression on page 989
- EBGP Route Selection on page 991
- IPv6 BGP Synchronization on page 993
- BGP Route Reflection on page 995
- IPv6 BGP Local Preference on page 998
- BGP Peer Group Update Policies on page 1001
- BGP Confederation on page 1002

Configuring Multiple Autonomous Systems

If an Autonomous System (AS) has multiple BGP routers, the AS can serve as a transit service for other ASs. When BGP runs between routers in different ASs, it uses exterior BGP (eBGP). When BGP runs between routers in the same AS, it uses interior BGP (iBGP).

In Autonomous System with multiple BGP routers configuration, AS 200 is a transit AS for AS 100 and AS 300.

Autonomous System with multiple BGP routers configuration



To configure multiple ASs as shown in Autonomous System with multiple BGP routers configuration, configure routers RTA, RTB, and RTC as follows:

On RTA:

```
router bgp 100
   neighbor 129.213.1.1 remote-as 200

address-family ipv6
   redistribute connected
   neighbor 129.213.1.1 activate
```

On RTB:

```
router bgp 200
   neighbor 129.213.1.2 remote-as 100
   neighbor 175.220.1.2 remote-as 200

address-family ipv6
   redistribute connected
   neighbor 129.213.1.2 activate
   neighbor 175.220.1.2 activate
```

On RTC:

```
router bgp 200
   neighbor 175.220.212.1 remote-as 200

address-family ipv6
   neighbor 175.220.212.1 activate
   neighbor 175.220.212.1 activate
```

Configuring Basic BGP over IPv6

A IPv6 BGP peer router can be configured to carry either IPv4 or IPv6 route information over either an IPv6 address family or an IPv4 address family. See Basic BGP over IPv6 configuration.

Basic BGP over IPv6 configuration



To configure basic BGP over IPv6:

1 Configure routers R1 and R2:

On R1:

```
router bgp 6501
   bgp router-id 1.1.1.1
   neighbor 2011:11:11:11::2 remote-as 6502

address-family ipv6
   neighbor 2011:11:11:11::2 activate

exit-address-family
```

On R2:

```
router bgp 6502
   bgp router-id 2.2.2.2
   neighbor 2011:11:11:11::1 remote-as 6501


address-family ipv6
   network 1010::1/128
   network 2020::1/128
   neighbor 2011:11:11:11::1 activate
```

Configuring EBGP Multihop

EBGP Multihop enables you to establish a neighbor connection between two external peers that are not directly connected. Multihop is available only for eBGP and is not available in for iBGP. When the Security Appliance has an external neighbor that does not have a direct connection, you can use the **`ebgp-multihop`** command to establish a neighbor connection.

To configure EBGP Multihop:

1 Configure routers R1 and R2:

On R1:

```
router bgp 6501
   bgp router-id 1.1.1.1
   neighbor 2011:11:11:11::2 remote-as 6502
   neighbor 2011:11:11:11::2 ebgp-multihop


address-family ipv6
   neighbor 2011:11:11:11::2 activate
exit-address-family
```

On R2:

```
router bgp 6502
```

```
   bgp router-id 2.2.2.2
   neighbor 2011:11:11:11::1 remote-as 6501
   neighbor 2011:11:11:11::1 ebgp-multihop

address-family ipv6
   network 1010::1/128
   network 2020::1/128
   neighbor 2011:11:11:11::1 activate
```

Configuring IPv6 BGP Outbound Route Filter

IPv6 BGP Outbound Route Filter (ORF) can be used to minimize the number of BGP updates sent between peer routers by filtering out unwanted routing updates at the source.

To configure IPv6 BGP Outbound Route Filter (ORF):

1 Configure routers R1 and R2:

On R1:

```
router bgp 6501
   bgp router-id 1.1.1.1
   neighbor 2011:11:11:11::2 remote-as 6502

address-family ipv6
   redistribute connected
   neighbor 2011:11:11:11::2 activate
   neighbor 2011:11:11:11::2 prefix-list pref1 in
   neighbor 2011:11:11:11::2 prefix-list pref2 out
exit-address-family

ipv6 prefix-list pref1 seq 10 deny 1010::1/128
ipv6 prefix-list pref1 seq 20 permit any
ipv6 prefix-list pref2 seq 10 deny 1111::1/128
ipv6 prefix-list pref2 seq 20 permit any
```

On R2:

```
router bgp 6502
   bgp router-id 2.2.2.2
   neighbor 2011:11:11:11::1 remote-as 6501

address-family ipv6
   redistribute connected
   neighbor 2011:11:11:11::1 activate
```

To check the routes on R1 and R2, use the **show bgp ipv6 unicast** command.

The route on R1 should have IPv6 address 1010::1/128.

The route on R2 should have IPv6 address 1111::1/128.

On R1:

```
R1> show bgp ipv6 unicast
```

On R2:

```
R2> show bgp ipv6 unicast
```

Configuring IPv6 BGP Distribute List

IPv6 BGP Distribute List can be used to minimize the number of BGP updates sent between peer routers by filtering out unwanted routing updates at the source.

To configure IPv6 BGP Distribute List:

1 Configure routers R1 and R2:

On R1:

```
router bgp 6501
   bgp router-id 1.1.1.1
   neighbor 2011:11:11:11::2 remote-as 6502


address-family ipv6
   redistribute connected
   neighbor 2011:11:11:11::2 activate
   neighbor 2011:11:11:11::2 distribute-list acl1 in
   neighbor 2011:11:11:11::2 distribute-list acl2 out
exit-address-family


ipv6 access-list acl1 deny 1010::1/128
ipv6 access-list acl1 permit any
ipv6 access-list acl2 deny 1111::1/128
ipv6 access-list acl2 permit any
```

On R2:

```
router bgp 6502
   bgp router-id 2.2.2.2
   neighbor 2011:11:11:11::1 remote-as 6501


address-family ipv6
   redistribute connected
   neighbor 2011:11:11:11::1 activate
```

To check the routes on R1 and R2, use the **show bgp ipv6 unicast** command.

The route on R1 should have IPv6 address 1010::1/128.

The route on R2 should have IPv6 address 1111::1/128.

On R1:

```
R1> show bgp ipv6 unicast
```

On R2:

```
R2> show bgp ipv6 unicast
```

IPv6 BGP Route-Map

IPv6 BGP Route-Map can be used to minimize the number of BGP updates sent between peer routers by filtering out unwanted routing updates at the source.

To configure IPv6 BGP Route-Map:

1 Configure routers R1 and R2:

On R1:

```
router bgp 6501
   bgp router-id 1.1.1.1
   neighbor 2011:11:11:11::2 remote-as 6502


address-family ipv6
   redistribute connected
   neighbor 2011:11:11:11::2 activate
   neighbor 2011:11:11:11::2 route-map map1 in
   neighbor 2011:11:11:11::2 route-map map2 out
exit-address-family


ipv6 access-list acl1 deny 1010::1/128
ipv6 access-list acl1 permit any
ipv6 access-list acl2 deny 1111::1/128
ipv6 access-list acl2 permit any
!
route-map map1 permit 1 match ipv6 address acl1
!
route-map map2 permit 1 match ipv6 address acl2
!
```

On R2:

```
router bgp 6502
   bgp router-id 2.2.2.2
   neighbor 2011:11:11:11::1 remote-as 6501


address-family ipv6
   redistribute connected
   neighbor 2011:11:11:11::1 activate
```

To check the routes on R1 and R2, use the **show bgp ipv6 unicast** command.

On R1:

```
R1> show bgp ipv6 unicast
```

The route on R1 should have IPv6 address 1010::1/128.

On R2:

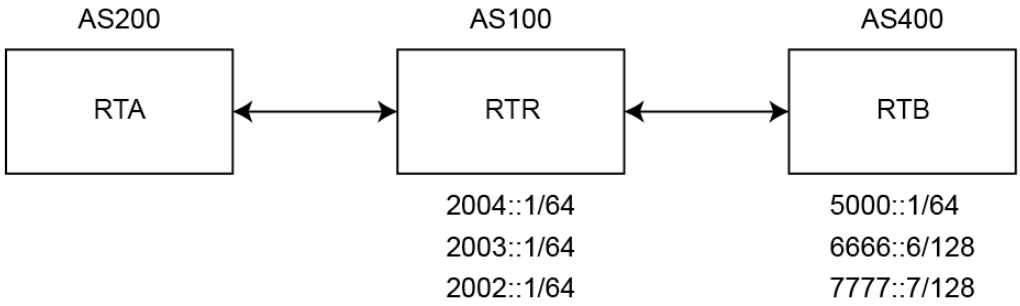
```
R2> show bgp ipv6 unicast
```

The route on R2 should have IPv6 address 1111::1/128.

Configuring an AS Regular Expression

You can configure regular expressions that can be matched and used to deny or allow addresses from an AS. See Autonomous System regular expression configuration.

Autonomous System regular expression configuration



RTB advertises these routes:

- 2004::/64
- 2003::/64
- 2002::/64

RTC advertises these routes:

- 5000::/64
- 6666::6/128
- 7777::7/128

To check the routes on router RTA:

1 Use the **`show bgp ipv6 unicast`** command:

On RTA:

```
RTA> show bgp ipv6 unicast

BGP table version is 4, local router ID is 10.0.1.2
Status codes: s suppressed, d damped, h history, * valid, > best,
i - internal, l - labeled
S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

Network          Next Hop            Metric     LocPrf     Weight     Path
*> 2002::/64     ::ffff:a00:101      0          0          100        i
*> 2003::/64     ::ffff:a00:101      0          0          100        i
*> 2004::/64     ::ffff:a00:101      0          0          100        i
*> 5000::/64     ::ffff:a00:101      0          0          100        400i
*> 6666::6/128   ::ffff:a00:101      0          0          100        400
*> 7777::7/128   ::ffff:a00:101      0          0          100        400
```

To configure AS regular expressions on RTA and deny all routes originated in AS100:

```
router bgp 200
   neighbor 10.0.1.1 remote-as 100
   neighbor 10.0.1.1 update-source X2
   neighbor 2004::1 remote-as 100
   neighbor 2004::1 update-source X2
!
address-family ipv6
   neighbor 10.0.1.1 activate
   neighbor 10.0.1.1 filter-list 1 in
```

```
   neighbor 2004::1 activate
   exit-address-family

ip as-path access-list 1 deny ^100$
ip as-path access-list 1 permit .*
```

To check the routes on router RTA:

1 Use the **`show bgp ipv6 unicast`** command.

On RTA:

```
RTA> show bgp ipv6 unicast
BGP table version is 4, local router ID is 10.0.1.2
Status codes: s suppressed, d damped, h history, * valid, > best,
i - internal, l - labeled
S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 5000::/64	::ffff:a00:101	0	0	100	400i
*> 6666::6/128	::ffff:a00:101	0	0	100	400i
*> 7777::7/128	::ffff:a00:101	0	0	100	400i

```
Total number of prefixes 3
```

To modify the AS path to deny all routes learned from the AS100:

On RTA:

```
router bgp 200
   neighbor 10.0.1.1 remote-as 100
   neighbor 10.0.1.1 update-source X2
   neighbor 2004::1 remote-as 100
   neighbor 2004::1 update-source X2
!
address-family ipv6
   neighbor 10.0.1.1 activate
   neighbor 10.0.1.1 filter-list 1 in
   neighbor 2004::1 activate
exit-address-family

ip as-path access-list 1 deny _100_
ip as-path access-list 1 permit .*
```

To check the routes on router RTA:

1 Use the **`show bgp ipv6 unicast`** command.

On RTA:

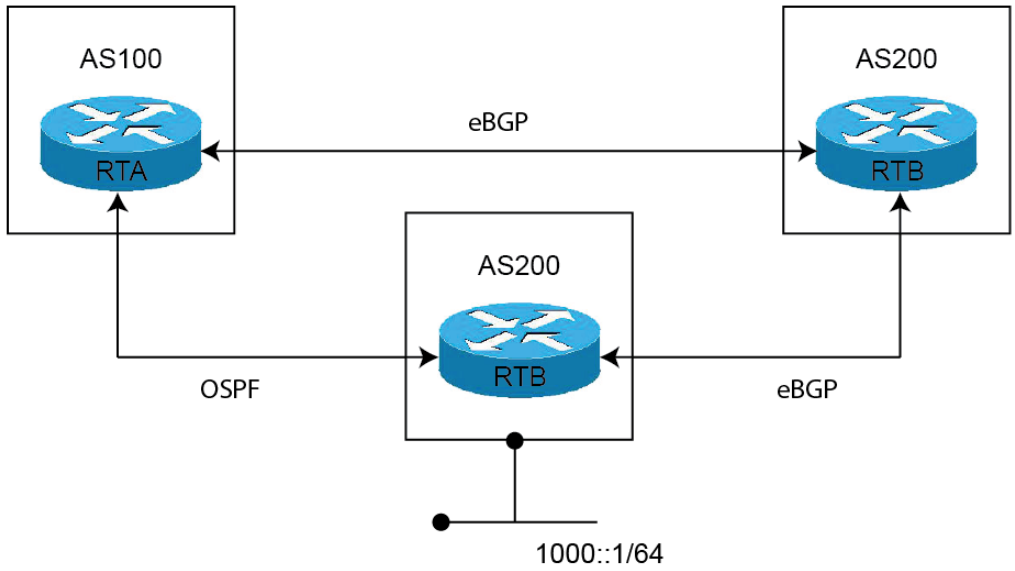
```
RTA> show bgp ipv6 unicast
```

EBGP Route Selection

Routes are selected based on the administrative distance of the routing protocol running on that route. Routing protocols with lower administrative distances are given priority over routing protocols with higher administrative distances. EBGP has an administrative distance of 20. OSPF has an administrative distance of 110.

Autonomous systems EBGP route selection configuration shows three ASs and the routing protocols used by the BGP routers.

Autonomous systems EBGP route selection configuration



The RTC router in AS300 advertises route 1000::/64 to both AS100 and to AS200.

The route from RTC (AS300) to RTA (AS100) runs OSPF.

The route from RTC (AS300) to RTB (AS200) runs eBGP.

The route from RTA (AS100) to RTB (AS200) runs eBGP.

RTA (AS100) receives updates about route 1000::/64 from both OSPF and eBGP. The route learned from eBGP is selected and added to RTA's routing table, because the administrative distance of eBGP is less than the administrative distance of OSPF.

On RTA:

```
router bgp 100
   neighbor 3001::1 remote-as 200
!
address-family ipv6
   distance bgp 150 150 150
   neighbor 3001::1 activate
exit-address-family
```

On RTB:

```
router bgp 200
   bgp log-neighbor-changes
   neighbor 1001::1 remote-as 300
   neighbor 2003::1 remote-as 100

address-family ipv6
   network 6666::6/128
   neighbor 1001::1 activate
   neighbor 2003::1 activate
exit-address-family
```

On RTC:

```
router bgp 300
```

```
   neighbor 3002::1 remote-as 200
!
address-family ipv6 network 1000::/64
   neighbor 3002::1 activate
exit-address-family
```

To check the routes on router RTA, use the **show ipv6 route** command.

```
RTA> show ipv6 route

IPv6 Routing Table

Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF, I - IS-IS, B - BGP
Timers: Uptime


B   1000::/64 [20/0] via fe80::204:27ff:fe0c:b006, X1, 00:01:07
C   2003::/64 via ::, X1, 00:30:50
B   6666::6/128 [20/0] via fe80::204:27ff:fe0c:b006, X1, 00:01:07
C   fe80::/64 via ::, X1, 00:30:53
```

As RTC is directly connected to RTA, the route from OSPF is actually a better route than the route learned by BGP. To ensure that the route between RTA and RTC is selected for the routing table, you can use the **distance** command to change the default administrative distance of the BGP route to a higher administrative distance than the OSPF route. For example:

```
   distance bgp 150 150 150
```

You can also use the **backdoor neighbor** command to set the BGP route as the preferred route. For example:

On RTA:

```
router bgp 100
   neighbor 3001::1 remote-as 200
!
address-family ipv6
   network 1000::/64
   backdoor neighbor 3001::1 activate
exit-address-family
```

To check the routes on router RTA:

1. Use the **show ipv6 route** command.

```
RTA> show ipv6 route

IPv6 Routing Table


Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF, I - IS-IS, B - BGP
Timers: Uptime


O   1000::/64 [110/2] via fe80::217:c5ff:feb4:57f2, X4, 00:30:53
C   2003::/64 via ::, X1, 00:31:18
B   6666::6/128 [20/0] via fe80::204:27ff:fe0c:b006, X1, 00:00:03
C   fe80::/64 via ::, X1, 00:31:21
```

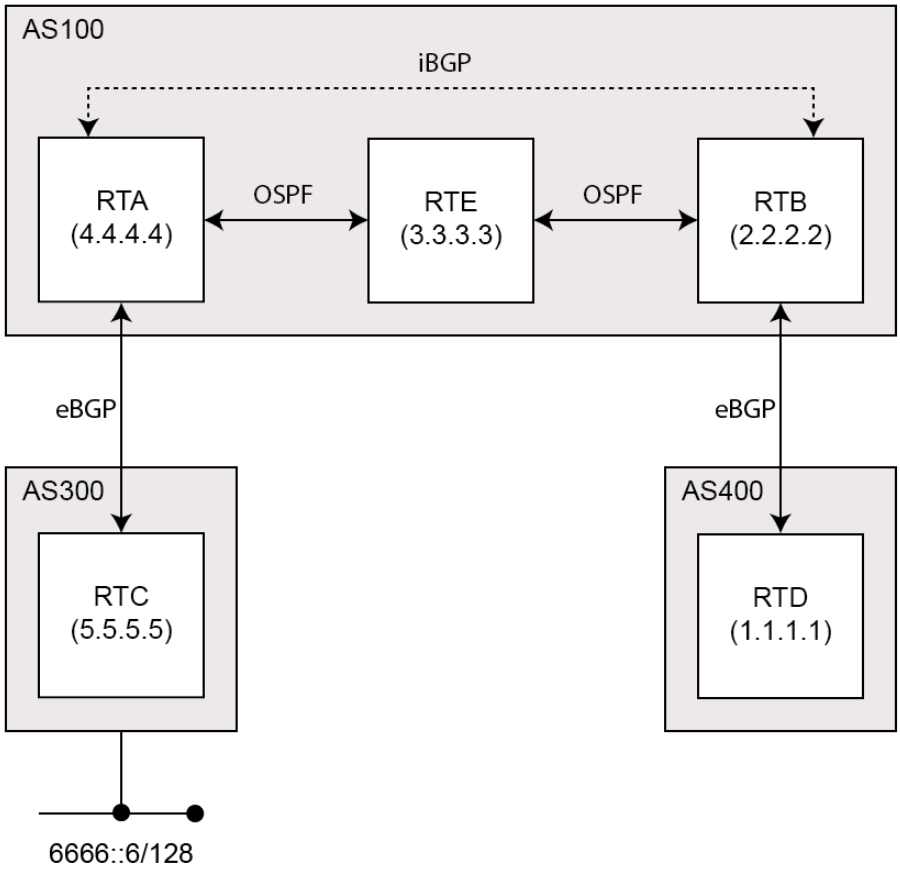
IPv6 BGP Synchronization

IPv6 BGP Synchronization keeps all BGP routers updated with the IPv6 addresses of all available routes and networks.

In BGP Synchronization, if an AS (AS100) passes traffic from another AS (AS300) to a third AS (AS400), BGP does not advertise that route until all the routers in AS100 have learned that route from the IGP. In this case, the IGP is iBGP. AS100 must wait until iBGP has propagated that route to all routers within AS100. Then, eBGP advertises the route to external ASs.

In this example, after RTB learns address 6666::6/128 via iBGP. it then advertises the address to RTD.

IPv6 BGP synchronization example



NOTE: You can make RTB think that IGP has already propagated the route information by adding a static route to 6666::6/128 on RTB and making sure that the other routers can reach 6666::6/128.

In this example, RTC (AS2) advertises address 6666::6/128 to RTA (AS100). In AS100, RTA and RTB are running iBGP, so RTB learns address 6666::6/128 and is able to reach it through next hop 5.5.5.5 (RTC). Next hop is carried through iBGP. However, to reach the next hop (RTC), RTB must send traffic through RTE, but RTE does not know IP address 6666::6/128.

If RTB advertises 6666::6/128 to RTD (AS400), traffic that tries to reach 6666::6/128 from RTD must pass through RTB and RTE in AS100. However, as RTE has not learned 6666::6/128, all packets are dropped at RTE.

To configure BGP Synchronization on RTB in AS100:

On RTB:

```
router bgp 100
   neighbor 10.103.10.129 remote-as 100
   neighbor 3001::1 remote-as 100
   neighbor 3001::1 update-source X4
```

```
   neighbor 5000::1 remote-as 400
   neighbor 5000::1 update-source X2
!
address-family ipv6
   synchronization
   neighbor 10.103.10.129 activate
   neighbor 3001::1 activate
   neighbor 5000::1 activate
exit-address-family
```

You can disable synchronization if you do not pass traffic from one AS to another AS through an intermediate AS. You can also disable synchronization if all routers in the intermediate AS run BGP. Disabling synchronization lets you to carry fewer routes in your IGP and allows BGP to converge more quickly.

To disable BGP Synchronization on RTB in AS100:

On RTB:

```
router bgp 100
   neighbor 10.103.10.129 remote-as 100
   neighbor 3001::1 remote-as 100
   neighbor 3001::1 update-source X4
   neighbor 5000::1 remote-as 400
   neighbor 5000::1 update-source X2
!
address-family ipv6
   neighbor 10.103.10.129 activate
   neighbor 3001::1 activate
   neighbor 5000::1 activate
exit-address-family
```

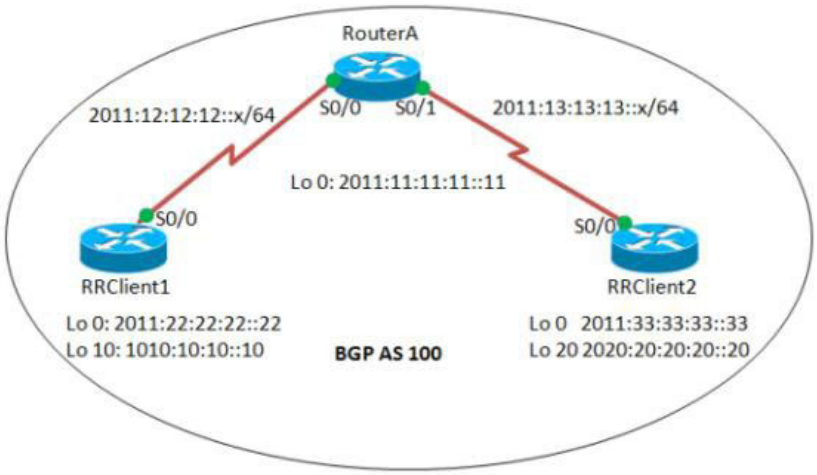
BGP Route Reflection

By default, all iBGP routers in an AS must be in a full mesh configuration. Each router must be configured as a peer to every other router.

With route reflection, all iBGP routers do not need to be fully meshed. Route reflection eliminates the need for each iBGP router to communicate with every other iBGP router in the AS. An iBGP router can be designated as a route reflector and can pass iBGP learned routes to multiple iBGP clients.

When a router is configured as a route reflector, it acts as a single point where all the other iBGP routers can get the iBGP learned routes. The route reflector acts like a server, rather than a peer, for every other router in the AS. All the other IBGP routers become route reflector clients. A router is a route reflector as long as it has at least one route reflector client.

BGP route reflection configuration



To configure route reflection in an AS:

On RouterA:

```
interface Serial0/0
   ipv6 address 2011:12:12:12::1/64
   ipv6 ospf 10 area 0

interface Serial0/1
   ipv6 address 2011:13:13:13::1/64
   ipv6 ospf 10 area 0

router bgp 100

bgp router-id 1.1.1.1
no bgp default ipv4-unicast
bgp log-neighbor-changes
   neighbor 2011:22:22:22::22 remote-as 100
   neighbor 2011:22:22:22::22 update-source Loopback0
   neighbor 2011:33:33:33::33 remote-as 100
   neighbor 2011:33:33:33::33 update-source Loopback0
!
address-family ipv6
   neighbor 2011:22:22:22::22 activate
   neighbor 2011:22:22:22::22 route-reflector-client
   neighbor 2011:33:33:33::33 activate
   neighbor 2011:33:33:33::33 route-reflector-client
exit-address-family
!
ipv6 router ospf 10
   router-id 1.1.1.1
```

On RRClient1:

```
interface Loopback0
   ipv6 address 2011:22:22:22::22/128
   ipv6 ospf 10 area 0
!
interface Loopback10
   ipv6 address 1010:10:10:10::10/128
```

```
interface Serial0/0
   ipv6 address 2011:12:12:12::2/64
   ipv6 ospf 10 area 0
!
router bgp 100
   bgp router-id 2.2.2.2
   bgp log-neighbor-changes
      neighbor 2011:11:11:11::11 remote-as 100
      neighbor 2011:11:11:11::11 update-source Loopback0
!
address-family ipv6
   neighbor 2011:11:11:11::11 activate
   network 1010:10:10:10::10/128
exit-address-family
!
ipv6 router ospf 10
   router-id 2.2.2.2
```

RRClient2:

```
interface Loopback0
   ipv6 address 2011:33:33:33::33/128
   ipv6 ospf 10 area 0
!
interface Loopback20
   ipv6 address 2020:20:20:20::20/128
!
interface Serial0/0
   no ip address
   ipv6 address 2011:13:13:13::2/64
   ipv6 ospf 10 area 0
!
router bgp 100
   bgp router-id 3.3.3.3
   bgp log-neighbor-changes
      neighbor 2011:11:11:11::11 remote-as 100
      neighbor 2011:11:11:11::11 update-source Loopback0
!
address-family ipv6
   neighbor 2011:11:11:11::11 activate
   network 2020:20:20:20::20/128
exit-address-family
!
ipv6 router ospf 10
   router-id 3.3.3.3
   log-adjacency-changes
```

To check the routes:

1. Use the **`show bgp ipv6 unicast`** command:

On RRClient1:

```
RRClient1> show bgp ipv6 unicast
```

You should see route `2020:20:20:20::20/128`.

On RRClient2:

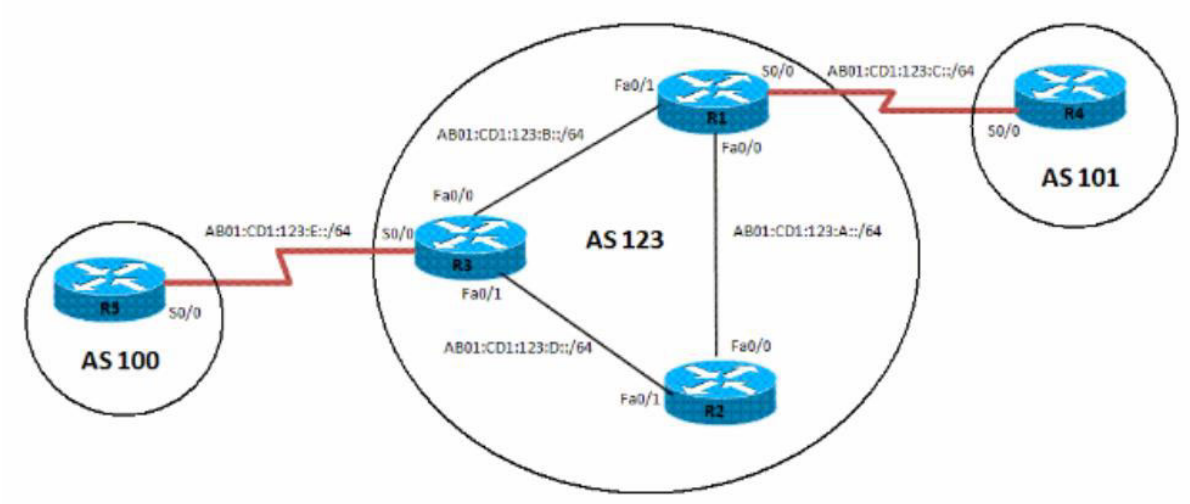
```
RRClient2> show bgp ipv6 unicast
```

You should see route `1010:10:10:10::10/128`.

IPv6 BGP Local Preference

The local preference designates a route to a certain network as the preferred exit route to that network from the AS. The route with a highest local preference is the preferred route. The default value of the local preference is 100, but this can be changed using the **`set local-preference`** command.

IPv6 BGP local preference configuration



To configure the local preference of a preferred route in an AS:

On R1:

```
interface Loopback0
   ipv6 address 1111:111:111:A::/64 eui-64
   ipv6 ospf 10 area 0

interface FastEthernet0/0
   ipv6 address AB01:CD1:123:A::/64 eui-64
   ipv6 ospf 10 area 0
!
interface Serial0/0
   ipv6 address AB01:CD1:123:C::/64 eui-64
!
interface FastEthernet0/1
   ipv6 address AB01:CD1:123:B::/64 eui-64
   ipv6 ospf 10 area 0
!
   ipv6 router ospf 10 router-id 1.1.1.1 log-adjacency-changes
   redistribute connected route-map CONNECTED
!
route-map CONNECTED permit 10
   match interface Serial0/0
!
router bgp 123
bgp router-id 1.1.1.1
   neighbor 2222:222:222:A:C602:3FF:FEF0:0 remote-as 123
   neighbor 2222:222:222:A:C602:3FF:FEF0:0 update-source Loopback0
   neighbor 3333:333:333:A:C603:3FF:FEF0:0 remote-as 123
   neighbor 3333:333:333:A:C603:3FF:FEF0:0 update-source Loopback0
   neighbor AB01:CD1:123:C:C604:16FF:FE98:0 remote-as 101
   neighbor AB01:CD1:123:C:C604:16FF:FE98:0 ebgp-multihop 5
!
address-family ipv6
   neighbor 2222:222:222:A:C602:3FF:FEF0:0 activate
```

```
   neighbor 2222:222:222:A:C602:3FF:FEF0:0 next-hop-self
   neighbor 3333:333:333:A:C603:3FF:FEF0:0 activate
   neighbor 3333:333:333:A:C603:3FF:FEF0:0 next-hop-self
   neighbor AB01:CD1:123:C:C604:16FF:FE98:0 activate exit-address-family
```

On R2:

```
interface Loopback0
   ipv6 address 2222:222:222:A::/64 eui-64
   ipv6 ospf 10 area 0
!
interface FastEthernet0/0
   ipv6 address AB01:CD1:123:A::/64 eui-64
   ipv6 ospf 10 area 0
!
interface FastEthernet0/1
   ipv6 address AB01:CD1:123:D::/64 eui-64
   ipv6 ospf 10 area 0
!
   ipv6 router ospf 10 router-id 2.2.2.2 log-adjacency-changes
!
router bgp 123
bgp router-id 2.2.2.2
   neighbor 1111:111:111:A:C601:3FF:FEF0:0 remote-as 123
   neighbor 1111:111:111:A:C601:3FF:FEF0:0 update-source Loopback0
   neighbor 3333:333:333:A:C603:3FF:FEF0:0 remote-as 123
   neighbor 3333:333:333:A:C603:3FF:FEF0:0 update-source Loopback0

address-family ipv6
   neighbor 1111:111:111:A:C601:3FF:FEF0:0 activate
   neighbor 3333:333:333:A:C603:3FF:FEF0:0 activate
exit-address-family
```

On R3:

```
interface Loopback0
   ipv6 address 3333:333:333:A::/64 eui-64
   ipv6 ospf 10 area 0
!
interface FastEthernet0/0
   ipv6 address AB01:CD1:123:B::/64 eui-64
   ipv6 ospf 10 area 0
!
interface Serial0/0
   ipv6 address AB01:CD1:123:E::/64 eui-64
!
interface FastEthernet0/1
   ipv6 address AB01:CD1:123:D::/64 eui-64
   ipv6 ospf 10 area 0
!
ipv6 router ospf 10
   router-id 3.3.3.3
   redistribute connected route-map CONNECTED
!
router bgp 123
   no synchronization
   bgp router-id 3.3.3.3
   neighbor 1111:111:111:A:C601:3FF:FEF0:0 remote-as 123
   neighbor 1111:111:111:A:C601:3FF:FEF0:0 update-source Loopback0
   neighbor 2222:222:222:A:C602:3FF:FEF0:0 remote-as 123
   neighbor 2222:222:222:A:C602:3FF:FEF0:0 update-source Loopback0
   neighbor AB01:CD1:123:E:C605:16FF:FE98:0 remote-as 202
   neighbor AB01:CD1:123:E:C605:16FF:FE98:0 ebgp-multihop 5
!
```

```
address-family ipv6
   neighbor 1111:111:111:A:C601:3FF:FEF0:0 activate
   neighbor 1111:111:111:A:C601:3FF:FEF0:0 next-hop-self
   neighbor 1111:111:111:A:C601:3FF:FEF0:0 route-map LOCAL_PREF out
   neighbor 2222:222:222:A:C602:3FF:FEF0:0 activate
   neighbor 2222:222:222:A:C602:3FF:FEF0:0 next-hop-self
   neighbor 2222:222:222:A:C602:3FF:FEF0:0 route-map LOCAL_PREF out
   neighbor AB01:CD1:123:E:C605:16FF:FE98:0 activate
exit-address-family
!
ipv6 prefix-list 10 seq 5 permit BC01:BC1:10:A::/64
!
route-map LOCAL_PREF permit 10
   match ipv6 address prefix-list 10
   set local-preference 500
!
route-map LOCAL_PREF permit 20
!
route-map CONNECTED permit 10
   match interface Serial0/0
```

On R4:

```
interface Serial0/0
   ipv6 address AB01:CD1:123:C::/64 eui-64
!
interface Loopback10
   ipv6 address BC01:BC1:10:A::/64 eui-64
!
interface Loopback11
   ipv6 address BC02:BC1:11:A::/64 eui-64
!
interface Loopback12
   ipv6 address BC03:BC1:12:A::/64 eui-64

router bgp 101
bgp router-id 4.4.4.4
   neighbor AB01:CD1:123:C:C601:3FF:FEF0:0 remote-as 123
!
address-family ipv6
   neighbor AB01:CD1:123:C:C601:3FF:FEF0:0 activate
   network BC01:BC1:10:A::/64 network BC02:BC1:11:A::/64
   network BC03:BC1:12:A::/64 exit-address-family
```

On R5:

```
interface Serial0/0
   ipv6 address AB01:CD1:123:E::/64 eui-64
   clock rate 2000000
!
interface Loopback10
   ipv6 address BC01:BC1:10:A::/64 eui-64
!
interface Loopback11
   ipv6 address BC02:BC1:11:A::/64 eui-64
!
interface Loopback12
   ipv6 address BC03:BC1:12:A::/64 eui-64
!
router bgp 202
bgp router-id 5.5.5.5
   neighbor AB01:CD1:123:E:C603:3FF:FEF0:0 remote-as 123
   neighbor AB01:CD1:123:E:C603:3FF:FEF0:0 ebgp-multihop 5
!
address-family ipv6
   neighbor AB01:CD1:123:E:C603:3FF:FEF0:0 activate
   network BC01:BC1:10:A::/64
```

```
   network BC02:BC1:11:A::/64
   network BC03:BC1:12:A::/64
exit-address-family
```

To verify the route:

1 Use the **`show bgp ipv6 unicast`** command:

On R2:

```
R2> show bgp ipv6 unicast
```

Before the local preference is configured, R2 has R1 as its next hop for all learned IPv6 addresses. After configuring the local preference on R3 to 500, R2 has a different preferred exit route for prefix `BC01:BC1:10:A::/64`. R2 can now reach prefix `BC01:BC1:10:A::/64` through the exit path of R3, which is now designated as the local preference.

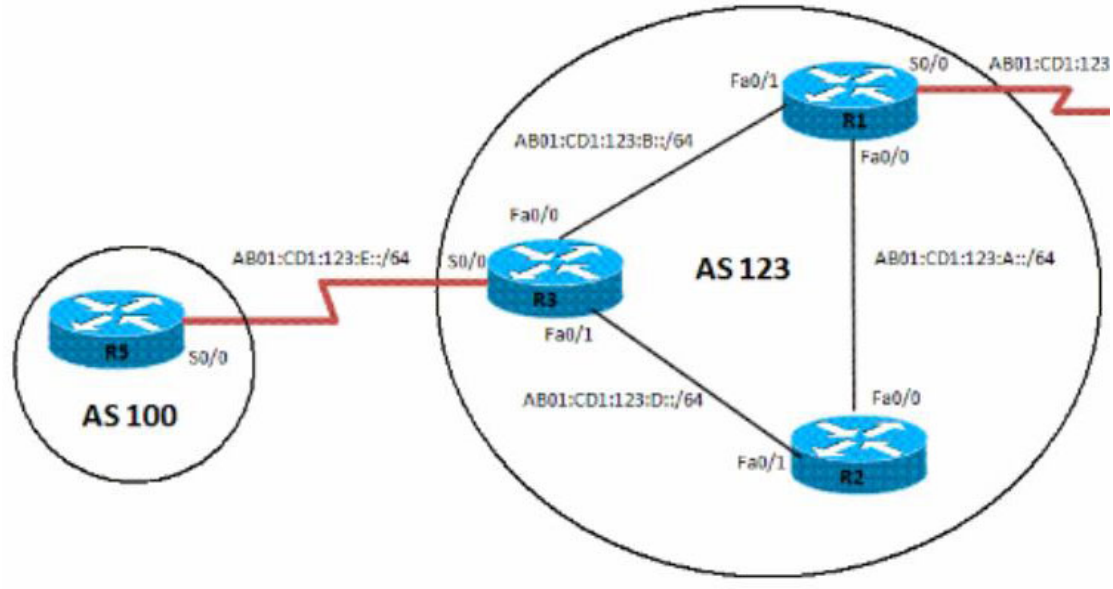
BGP Peer Group Update Policies

A BGP peer group is a group of BGP neighbors that share the same update policies. Update policies are typically set by route maps, distribution lists, and filter lists.

When you define a peer group and add neighbors to it, all of the update policies that you assign to that peer group apply to all of the neighbors in that peer group. You do not need to define a policy for each neighbor.

Members of a peer group inherit all of the configuration settings of that peer group. You can configure certain members to override the update policies, but only if those policies are set for inbound traffic. You cannot configure members to override group policies if the policies apply to outbound traffic.

BGP peer group update policy configuration



To configure an IPv6 BGP peer group and its update policies:

On R3:

```
router bgp 123
   no synchronization
```

```
   bgp router-id 3.3.3.3
neighbor interalmap peer-group
   neighbor interalmap remote-as 123
   neighbor 1111:111:111:A:C601:3FF:FEF0:0 peer-group interalmap
   neighbor 2222:222:222:A:C602:3FF:FEF0:0 peer-group interalmap
   neighbor AB01:CD1:123:E:C605:16FF:FE98:0 remote-as 202
   neighbor AB01:CD1:123:E:C605:16FF:FE98:0 ebgp-multihop 5
!
address-family ipv6
   neighbor interalmap activate
   neighbor interalmap route-map 1 out
   neighbor 1111:111:111:A:C601:3FF:FEF0:0 peer-group interalmap
   neighbor 2222:222:222:A:C602:3FF:FEF0:0 peer-group interalmap
exit-address-family
!
ipv6 prefix-list 10 seq 5 permit BC01:BC1:10:A::/64
!
route-map 1 permit 10
   match ipv6 address prefix-list 1 set tag 333
   set metric 273
   set local-preference 312
```

To verify that the correct local preference route is configured:

1 Use the **`show bgp ipv6 unicast`** command:

On R3:

```
R3> show bgp ipv6 unicast
```

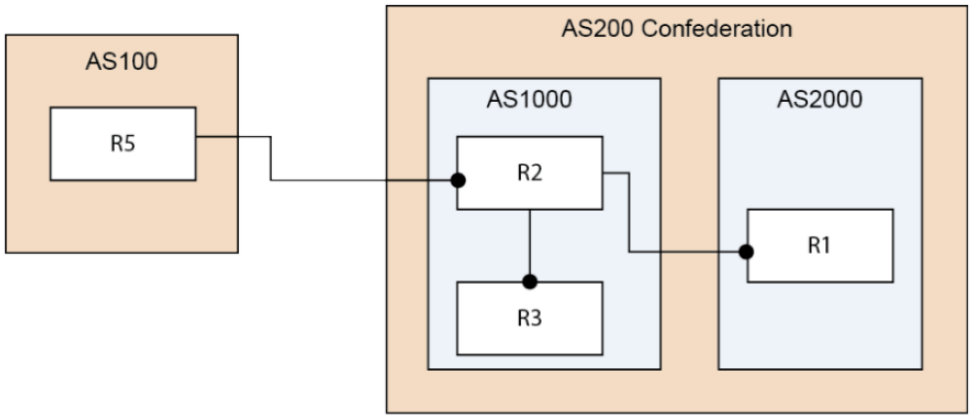
Verify that IPv6 address BC01:BC1:10:A::/64 passes from AS100 to R1 and R2, and that the metric and local preference are set to the corresponding route-map settings.

BGP Confederation

You can divide a single AS into multiple ASs, and then assign these multiple ASs to a single confederation of ASs. The implementation of a BGP confederation reduces the iBGP mesh size of the AS, and the confederation can still advertise as a single AS to external peers.

Each individual AS within a confederation runs fully meshed iBGP, and each individual AS within the confederation also runs eBGP connections to the other ASs inside the confederation. These eBGP peers within the confederation exchange routing information as if they used iBGP. In this way, the confederation preserves next hop, metric, and local preference information. To the outside world, the confederation appears to be a single AS.

BGP confederation configuration



To configure a BGP Confederation:

R1:

```
router bgp 2000
   bgp log-neighbor-changes
   bgp confederation identifier 200
   bgp confederation peers 1000
   neighbor 2003::1 remote-as 1000
!
address-family ipv4
   neighbor 2003::1 activate
exit-address-family
!
address-family ipv6
   network 3002::/64
   network 4000::/64
   neighbor 2003::1 activate
exit-address-family
```

On R2:

```
router bgp 1000
   bgp confederation identifier 200
   neighbor 10.0.1.1 remote-as 1000
!
address-family ipv6
   neighbor 10.0.1.1 activate
exit-address-family
```

On R3:

```
router bgp 1000
   bgp confederation identifier 200
   bgp confederation peers 2000
   neighbor 10.0.1.2 remote-as 1000
   neighbor 3001::1 remote-as 2000
   neighbor 5000::1 remote-as 100
   neighbor 5000::1 update-source X2
!
address-family ipv6
   neighbor 10.0.1.2 activate
   neighbor 3001::1 activate
   neighbor 5000::1 activate
exit-address-family
```

On R5:

```
router bgp 100
   bgp router-id 5.5.5.5
   bgp log-neighbor-changes
   neighbor 2002::1 remote-as 200
!
address-family ipv6
   network 6666::6/128
   network 7777::7/128
   neighbor 2002::1 activate
exit-address-family
```

Verify that R1, R2, and R3 can learn this route that is advertised by R5:

```
      6666::6/128 and 7777::7/128
```

Verify that R2 can learn this route from R1 even though they are not directly connected:

```
3002::/64 and 4000::/64
```

NOTE: The IPv6 BGP configuration data and the IPv6 BGP routes are dumped into a Terminate and Stay Resident (TSR) file.

NOTE: IPv6 BGP uses the ZebOS debug interface. The default setting for all debug switches is closed. Entering the CLI **debug** command on the console opens the debug switch.

D

SonicWall Support

Technical support is available to customers who have purchased SonicWall products with a valid maintenance contract and to customers who have trial versions.

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. To access the Support Portal, go to https://support.sonicwall.com.

The Support Portal enables you to:

- View knowledge base articles and technical documentation
- View and participate in the Community forum discussions: https://community.sonicwall.com/technology-and-support/
- View video tutorials
- Access MySonicWall
- Learn about SonicWall professional services
- Review SonicWall Support services and warranty information
- Register for training and certification
- Request technical support or customer service

To contact SonicWall Support, refer to https://support.sonicwall.com/contact-support.

About This Document

Legend

WARNING: A WARNING icon indicates a potential for property damage, personal injury, or death.

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

IMPORTANT, NOTE, TIP, MOBILE, or VIDEO: An information icon indicates supporting information.

SonicOS 6.5 System Setup
Updated - July 2020
Software Version - 6.5.4
232-001810-04 Rev E

Copyright © 2020 SonicWall Inc. All rights reserved.

SonicWall is a trademark or registered trademark of SonicWall Inc. and/or its affiliates in the U.S.A. and/or other countries. All other trademarks and registered trademarks are property of their respective owners

The information in this document is provided in connection with SonicWall Inc. and/or its affiliates' products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of SonicWall products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, SONICWALL AND/OR ITS AFFILIATES ASSUME NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL SONICWALL AND/OR ITS AFFILIATES BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF SONICWALL AND/OR ITS AFFILIATES HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. SonicWall and/or its affiliates make no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. SonicWall Inc. and/or its affiliates do not make any commitment to update the information contained in this document.

For more information, visit https://www.sonicwall.com/legal.

Open Source Code

SonicWall is able to provide a machine-readable copy of open source code with restrictive licenses such as GPL, LGPL, AGPL when applicable per license requirements. To obtain a complete machine-readable copy, send your written requests, along with certified check or money order in the amount of US 25.00 payable to "SonicWall Inc.", to:

General Public License Source Code Request
SonicWall Inc. Attn: Jennifer Anderson
1033 McCarthy Blvd
Milpitas, CA 95035