SonicWall® SonicOS 6.5 Security Configuration

Administration

SONICWALL®

Contents

Part 1. Security Config | Firewall Settings

Configuring Advanced Firewall Settings 11
Firewall Settings > Advanced 12
Detection Prevention 12
Dynamic Ports 13
Source Routed Packets 16
Internal VLAN 16
Connections 16
Dynamic Connection Sizing 18
Access Rule Service Options 19
IP and UDP Checksum Enforcement 19
Jumbo Frame 20
IPv6 Advanced Configuration 20
Control Plane Flood Protection 22

Configuring Bandwidth Management 23
Understanding Bandwidth Management 23
Glossary 24
Firewall Settings > BWM 25
Configuring Global Bandwidth Management Settings 25
Action Objects 28
Configuring Global Bandwidth Management 28
Configuring Global Bandwidth Management 29
Configuring Global BWM on an Interface 29
Configuring Global BWM in an Access Rule 31
Configuring Global BWM in an Action Object 31
Configuring Application Rules 32
Configuring App Flow Monitor 34
Advanced Bandwidth Management 36
Elemental Bandwidth Settings 38
Zone-Free Bandwidth Management 39
Weighted Fair Queuing 39
Configuring Advanced Bandwidth Management 40
Enabling Advanced Bandwidth Management 40
Enabling Global Bandwidth Management 41
Configuring Bandwidth Policies 41
Setting Interface Bandwidth Limitations with Advanced BWM 47
Setting Interface Bandwidth Limitations with Global BWM 49

Configuring Flood Protection 51
Firewall Settings > Flood Protection 52
TCP View 53
UDP View 63

ICMP View 66

Configuring Firewall Multicast 69
Firewall Settings > Multicast 69
Multicast Snooping 70
Multicast Policies 70
IGMP State Table 71
Enabling Multicast 72
Enabling Multicast on LAN-Dedicated Interfaces 72
Enabling Multicast Through a VPN 73

Managing Quality of Service 76
Firewall Settings > Quality of Service Mapping 76
Classification 76
Marking 77
Conditioning 78
802.1p and DSCP QoS 79
Bandwidth Management 90
Glossary 90

Configuring SSL Control 94
About SSL Control 94
Firewall Settings > SSL Control 102
SSL Control Configuration 102
Enabling SSL Control on Zones 106
SSL Control Events 106

Configuring Cipher Control 108
About Cipher Control 108
Firewall Settings > Cipher Control 108
TLS Ciphers 109
SSH Ciphers 115

Part 2. Security Config | Security Services

Managing SonicWall Security Services 117
About SonicWall Security Services 117
Configuring Security Services 118
Viewing and Managing Licenses 119
Synchronize Licenses 119
Security Services Settings 119
Signature Downloads Through a Proxy Server 120
Security Services Information 121
Update Signature Manually 121

Configuring Content Filtering Service 122
Security Services > Content Filter: SonicWall CFS 123
About CFS 124

Enabling CFS . . . 126
Enabling the Local CFS Server . . . 127
Configuring CFS Policies . . . 128
Configuring CFS Custom Categories . . . 128
Security Services > Content Filter: Websense Enterprise . . . 135
Selecting Websense Enterprise Content Filter Type . . . 136
Viewing Websense Server Status . . . 136
Configuring General Settings . . . 137
Configuring Web Features to Block . . . 138
Configuring CFS Exclusions . . . 138
Creating a Use Policy Blocking Page . . . 139

Configuring Client AV Enforcement . . . 141
Security Services > Client AV Enforcement . . . 141
About SonicWall Capture Client . . . 141
About Client AV Enforcement . . . 141
Configuring Client Anti-Virus Enforcement . . . 142
Capture Client Status . . . 144
McAfee / SentinelOne Client AV Status . . . 145
Client Anti-Virus Policies . . . 146
Client Anti-Virus Enforcement . . . 147

DPI-SSL Enforcement . . . 152
About DPI-SSL Enforcement . . . 152
Links . . . 153
DPI-SSL Enforcement Status . . . 153
DPI-SSL Enforcement . . . 153
Managing DPI-SSL Enforcement . . . 153
Editing a DPI-SSL Enforcement List . . . 154
Adding a Policy to a DPI-SSL Enforcement List . . . 154
Editing a DPI-SSL Enforcement Policy . . . 154
Managing Zones for DPI-SSL Enforcement . . . 155

Configuring Client CF Enforcement . . . 156
Security Services > Client CF Enforcement . . . 156
Enabling and Configuring Client CF Enforcement . . . 156
Enabling Client CFS in Network Zones . . . 158

Managing SonicWall Gateway Anti-Virus Service . . . 160
About SonicWall Gateway Anti-Virus Service . . . 160
SonicWall GAV Multi-Layered Approach . . . 161
SonicWall GAV Architecture . . . 164
Activating the Gateway Anti-Virus, Anti-Spyware, and IPS License . . . 165
Setting Up SonicWall Gateway Anti-Virus Protection . . . 165
Viewing SonicWall GAV Signatures . . . 175

Activating Intrusion Prevention Service . . . 178
About Intrusion Prevention Service . . . 178
SonicWall Deep Packet Inspection . . . 178

How SonicWall's Deep Packet Inspection Works . . . 179
Glossary . . . 179
Configuring Intrusion Prevention Service . . . 180
IPS Status . . . 181
IPS Global Settings . . . 181
Enabling IPS . . . 182
Configuring an IPS Exclusion List . . . 183
Resetting the IPS Settings and Policies . . . 184
Configuring IPS Protection on Zones . . . 184
IPS Policies . . . 185

Configuring Capture ATP . . . 189
Security Services > Capture ATP . . . 189
About Capture ATP . . . 190
Files are Preprocessed . . . 190
Files Blocked Until Completely Analyzed . . . 191
Files are Sent over an Encrypted Connection . . . 191
Capture ATP Friendly Filename Display . . . 191
Activating the Capture ATP License . . . 191
Enabling Capture ATP . . . 192
About the Security Services > Capture ATP Page . . . 192
Configuring Capture ATP . . . 198
Disabling GAV or Cloud Anti-Virus . . . 201

Activating Anti-Spyware Service . . . 202
About Anti-Spyware . . . 202
Security Services > Anti-Spyware . . . 203
Anti-Spyware Status . . . 204
Anti-Spyware Global Settings . . . 204
Enabling Inspection of Outbound Spyware Communication . . . 205
Applying Anti-Spyware Protection on Zones . . . 205
Anti-Spyware Policies . . . 206
Configuring Anti-Spyware Policies . . . 207
Configuring Category Policies . . . 207
Configuring Signature Policies . . . 209

Configuring SonicWall Real-Time Blacklist . . . 211
Security Services > RBL Filter . . . 211
About Real-Time Black List Filtering . . . 212
Configuring the RBL Filter . . . 212
Enabling RBL Blocking . . . 212
Adding RBL Services . . . 213
Configuring User-Defined SMTP Server Lists . . . 214
Testing SMTP IP Addresses . . . 216

Configuring Geo-IP Filters . . . 217
Security Services > Geo-IP Filter . . . 217
Configuring Geo-IP Filtering . . . 218

Creating a Custom Country List 220
Creating a Custom List 221
Editing a Custom List Entry 223
Deleting Custom List Entries 223
Customizing Web Block Page Settings 224
Using Geo-IP Filter Diagnostics 226
Show Resolved Locations 227
Geo-IP Cache Statistics 227
Custom Countries Statistics 228
Check GEO Location Server Lookup 228
Incorrectly Marked Address 229

Configuring Botnet Filters 230
Security Services > Botnet Filter 230
Configuring Botnet Filtering 231
Creating a Custom Botnet List 232
Creating a Custom Botnet List 233
Editing a Custom Botnet List Entry 234
Deleting Custom Botnet List Entries 235
Configuring Dynamic HTTP Authentication 236
Customizing Web Block Page Settings 237
Using Botnet Filter Diagnostics 239
Show Resolved Botnet Locations 240
Botnet Cache Statistics 240
Botnets Statistics 241
Check Botnet Server Lookup 241
Incorrectly Marked Address 242
Displaying the Status of the Botnet Feature and Database 242

Part 3. Security Config | Decryption Services

About DPI-SSL 244
About DPI-SSL 244
Supported Features 244
Security Services 246
Deployment Scenarios 247
Proxy Deployment 247
Customizing DPI-SSL 247
Connections per Appliance Model 248

Configuring the DPI-SSL/TLS Client 249
Decryption Services > DPI-SSL/TLS Client 249
Viewing DPI-SSL Status 250
Configuring the DPI-SSL/TLS Client 250
Configuring General Settings 250
Selecting the Re-Signing Certificate Authority 253
Configuring Exclusions and Inclusions 254
Excluding/Including by Common Name 256

Specifying CFS Category-based Exclusions/Inclusions . . . 262
Client DPI-SSL Examples . . . 264

Configuring DPI-SSL/TLS Server Settings . . . 267
Decryption Services > DPI-SSL/TLS Server . . . 267
Configuring DPI-SSL/TLS Server Settings . . . 268

Configuring DPI-SSH . . . 271
About DPI-SSH . . . 271
Supported Clients/Servers and Connections . . . 272
Supported Key Exchange Algorithms . . . 272
Caveats . . . 272
Activating Your DPI-SSH License . . . 273
Configuring DPI-SSH . . . 273
Viewing Connection Status . . . 274
Configuring Client DPI-SSH Inspection . . . 274
DPI-SSH Blocking of Port Forwarding . . . 275
Customizing Client DPI-SSH Inspection . . . 276

Part 4. Security Config | Anti-Spam

About Anti-Spam . . . 279
About Anti-Spam . . . 279
What is Anti-Spam? . . . 279
Benefits . . . 280
How Does the Anti-Spam Service Work? . . . 280
GRID Network . . . 280
Address and Service Objects . . . 282
Purchasing an Anti-Spam License . . . 285

Enabling and Activating Anti-Spam . . . 288
Anti-Spam > Base Setup . . . 289
Activating Anti-Spam . . . 289
Installing the Junk Store . . . 291
Configuring Email Threat Categories . . . 292
Configuring Access Lists . . . 293
Configuring the Access Lists . . . 294
Adding a Host to the Access Lists . . . 294
Configuring Advanced Options . . . 295

Configuring Anti-Spam Logging . . . 298
Anti-Spam > Advanced . . . 298
Downloading System/Log Files . . . 299
Selecting the Amount and Level of Log Information . . . 300

Configuring the RBL Filter . . . 302
Anti-spam > Real-Time Blacklist Filter . . . 303
About RBL Lists . . . 303

Enabling the RBL Filter 304
Managing RBL Services 305
Clearing Statistics 306
Adding an RBL Service 306
Editing an RBL Service 307
Deleting an RBL Service 308
User-Defined SMTP Server Lists 308
Testing the Real-time Black List 310

Specifying Relay Domains 311
Anti-Spam > Relay Domains 311
About Open Relay 312
Listing Allowed Relay Domains 312

Configuring Junk Box Settings 313
Anti-Spam > Junk Box Settings 313

Managing the Junk Summary 315
Anti-Spam > Junkbox Summary 315
Managing the Junk Summary 317
Reverting to Defaults 319

Configuring the Junk Box View 320
Anti-Spam > Junk Box 321
About the Junk Box Tabs 322
Searching the Messages 323
Performing a Simple Search 323
Performing an Advanced Search 324
Managing Messages in the Junk Store 327

Configuring User-Visible Settings 329
Anti-Spam > User View Setup 329
Configuring User View Setup 330
Reverting to Default Settings 330

Configuring Corporate Allowed and Blocked Lists 331
Anti-Spam > Address Books 331
About the Tabs 332
Allowed Lists 332
Blocked Lists 332
Adding Items to the Allowed or Blocked List 333
Deleting Items from the Allowed or Blocked List 334
Importing Address Book Entries 334
Exporting Address Book Entries 335
Searching the Allowed and Blocked Lists 335

Managing Users 336
Anti-Spam > Manage Users 337

Updating the User Table 338
Enabling Non-LDAP User Authentication 338
Viewing Users 339
Selecting the Type of User to View 339
Selecting a Server's Users to View 339
Finding a User 340
Adding Users 340
Adding Users Manually to the User Table 341
Importing Users to the User Table 341
Signing In as a User 342

Configuring the LDAP Server 343
Anti-Spam > LDAP Configuration 343
Available LDAP Servers 344
Adding an LDAP Server 345
Configuring LDAP Queries 349
Adding LDAP Mappings 351
Configuring Global LDAP Settings 353
Editing an LDAP Server Configuration 354
Deleting an LDAP Server 355

Downloading Anti-Spam Desktop Buttons 356
Anti-Spam > Downloads 356

Part 5. Security Config | Appendix

SonicWall Support 358
About This Document 359

Part 1

Security Config | Firewall Settings

- Configuring Advanced Firewall Settings
- Configuring Bandwidth Management
- Configuring Flood Protection
- Configuring Firewall Multicast
- Managing Quality of Service
- Configuring SSL Control
- Configuring Cipher Control

1

Configuring Advanced Firewall Settings

- Firewall Settings > Advanced on page 12
- Detection Prevention on page 12
- Dynamic Ports on page 13
- Source Routed Packets on page 16
- Internal VLAN on page 16
- Connections on page 16
- Dynamic Connection Sizing on page 18
- Access Rule Service Options on page 19
- IP and UDP Checksum Enforcement on page 19
- Jumbo Frame on page 20
- IPv6 Advanced Configuration on page 20
- Control Plane Flood Protection on page 22

Firewall Settings > Advanced

This section provides advanced firewall settings for configuring detection prevention, dynamic ports, source routed packets, connection selection, and access rule options. To configure advanced access rule options, select **MANAGE | Security Configuration > Firewall Settings > Advanced Settings**.

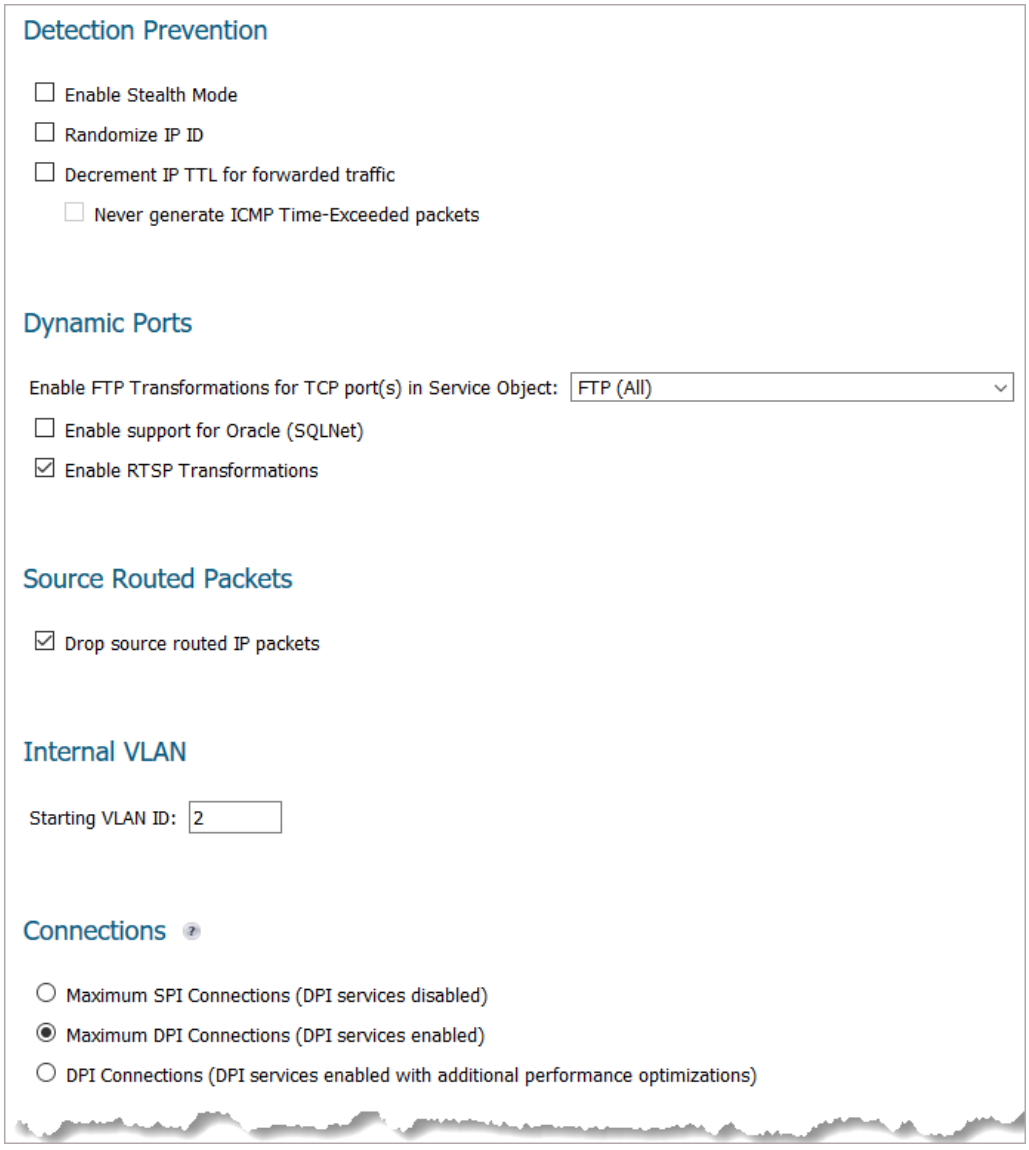

Detection Prevention

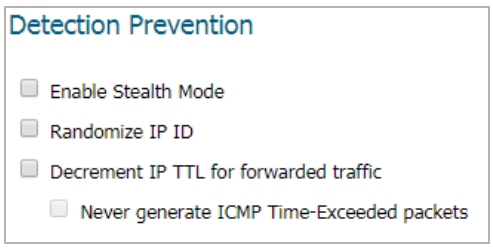

To enable detection prevention:

1 Navigate to **MANAGE | Security Configuration > Firewall Settings > Advanced Settings**.

2 Scroll to **Detection Prevention**.

3 By default, the security appliance responds to incoming connection requests as either blocked or open. To ensure your security appliance does not respond to blocked inbound connection requests, select **Enable Stealth Mode**. Stealth Mode makes your security appliance essentially invisible to hackers. This option is not selected by default.

4 To prevent hackers using various detection tools from detecting the presence of a security appliance, select **Randomize IP ID**. IP packets are given random IP IDs, which makes it more difficult for hackers to "fingerprint" the security appliance. This option is not selected by default.

5 Time-to-live (TTL) is a value in an IP packet that tells a network router whether or not the packet has been in the network too long and should be discarded. To decrease the TTL value for packets that have been forwarded and, therefore, have already been in the network for some time, select **Decrement IP TTL for forwarded traffic**. This option is not selected by default.

When you select this option, the following option becomes available.

6 The firewall generates Time-Exceeded packets to report when a packet its dropped because its TTL value has decreased to zero. To prevent the firewall from generate these reporting packets, select **Never generate ICMP Time-Exceeded packets**. This option is not selected by default.

7 Click **ACCEPT**.

Dynamic Ports



To configure dynamic ports:

1 Navigate to **MANAGE | Security Configuration > Firewall Settings > Advanced Settings**.

2 Scroll to **Dynamic Ports**.

3 From **Enable FTP Transformations for TCP port(s) in Service Object**, select the service group to enable FTP transformations for a particular service object. By default, service group **FTP (All)** is selected.

FTP operates on TCP ports 20 and 21, where port 21 is the Control Port and 20 is Data Port. When using non-standard ports (for example, 2020, 2121), however, SonicWall drops the packets by default as it is not able to identify it as FTP traffic. The **Enable FTP Transformations for TCP port(s) in Service Object** option allows you to select a Service Object to specify a custom control port for FTP traffic.

To illustrate how this feature works, consider the following example of an FTP server behind the SonicWall listening on port 2121:

a On the **MANAGE | Policies > Objects > Address Objects** page, create an **Address Object** for the private IP address of the FTP server with the following values:

- **Name**: FTP Server Private
- **Zone**: LAN

- **Type**: Host
- **IP Address**: 192.168.168.2

b On the **MANAGE | Policies > Objects > Services Objects** page, create a custom Service for the FTP Server with the following values:

- **Name**: FTP Custom Port Control
- **Protocol**: TCP(6)
- **Port Range**: 2121 - 2121

c On the **MANAGE | Policies > Rules > NAT Policies** page, create this NAT Policy:

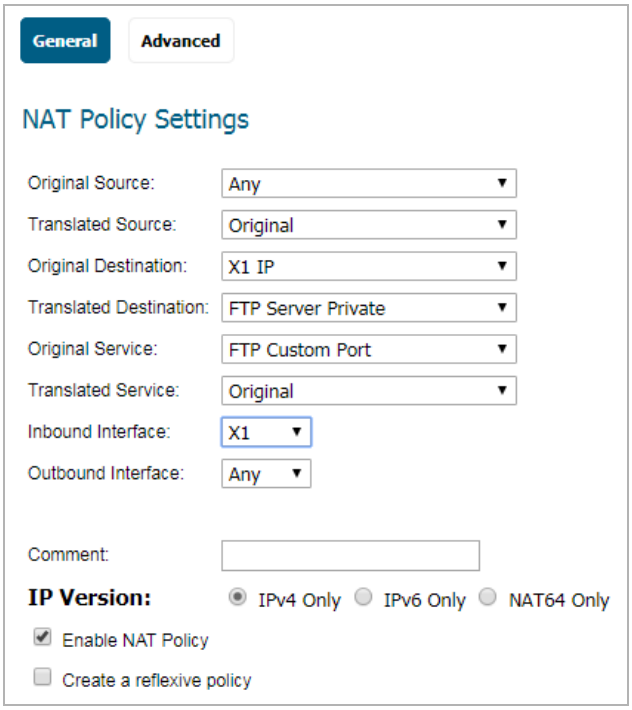

d On the **MANAGE | Policies > Rules > Access Rules** page, create this Access Rule:

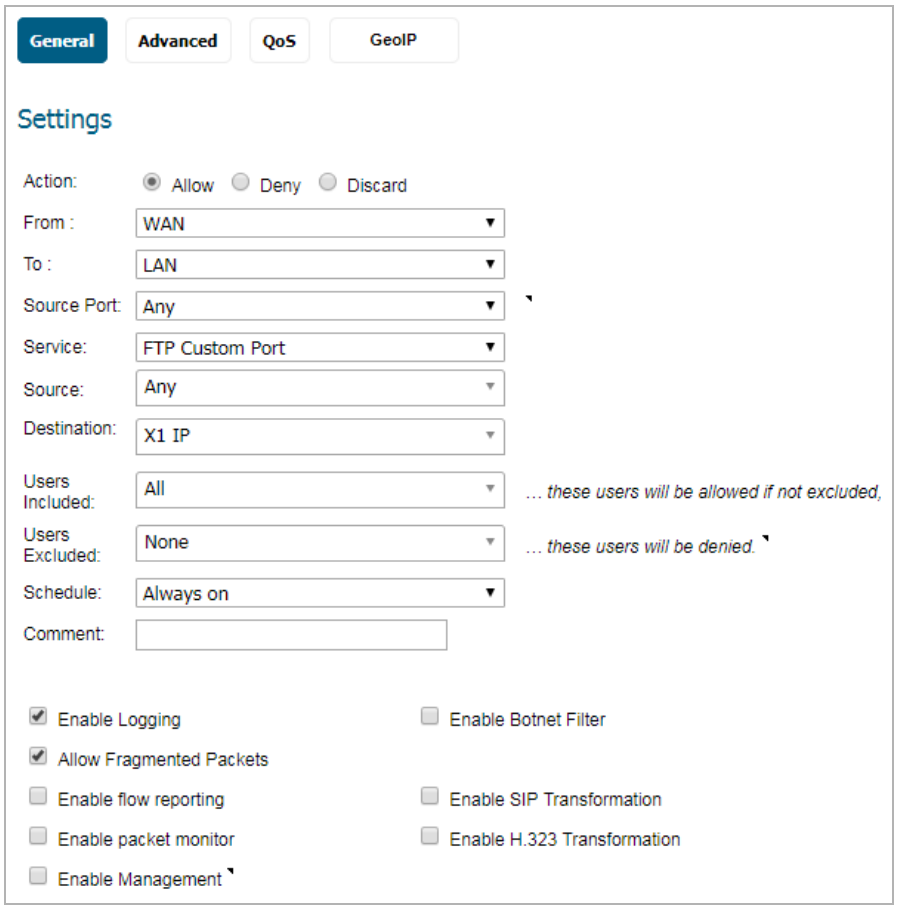

e On the **MANAGE | Security Configuration > Firewall Settings > Advanced Settings** page, from **Enable FTP Transformations for TCP port(s) in Service Object**, select the **FTP Custom Port Control** Service Object.

NOTE: For more information on configuring service groups and service objects, refer to *SonicWall SonicOS 6.5 System Setup*.

4 If you have Oracle9i or earlier applications on your network, select **Enable support for Oracle (SQLNet)**. This option is not selected by default.

IMPORTANT: For Oracle10g or later applications, it is recommended that this option not be selected.

For Oracle9i and earlier applications, the data channel port is different from the control connection port. When this option is enabled, a SQLNet control connection is scanned for a data channel being negotiated. When a negotiation is found, a connection entry for the data channel is created dynamically, with NAT applied if necessary. Within SonicOS, the SQLNet and data channel are associated with each other and treated as a session.

For Oracle10g and later applications, the two ports are the same, so the data channel port does not need to be tracked separately; thus, the option does not need to be enabled.

5 To support on-demand delivery of real-time data, such as audio and video, select **Enable RTSP Transformations**. RTSP (Real Time Streaming Protocol) is an application-level protocol for control over delivery of data with real-time properties. This option is selected by default.

6 Click **ACCEPT**.

Source Routed Packets

IP Source Routing is a standard option in IP that allows the sender of a packet to specify some or all of the routers that should be used to get the packet to its destination.

This IP option is typically blocked from use as it can be used by an eavesdropper to receive packets by inserting an option to send packets from A to B via router C. The routing table should control the path that a packet takes, so that it is not overridden by the sender or a downstream router.

To configure source-routed packets:

1 Navigate to **MANAGE | Security Configuration > Firewall Settings > Advanced Settings**.

2 Scroll to **Source Routed Packets**.



3 Ensure the **Drop Source Routed IP Packets** option is selected. This option is selected by default.

TIP: If you are testing traffic between two specific hosts and you are using source routing, deselect this option.

4 Click **ACCEPT**.

Internal VLAN

The Internal VLAN section allows you to specify the starting VLAN ID.

To change the internal VLAN ID:

1 Navigate to **MANAGE | Security Configuration > Firewall Settings > Advanced Settings**.

2 Scroll to the **Internal VLAN** section.

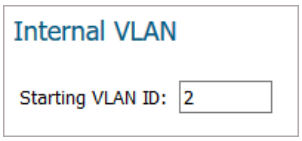

3 Enter the VLAN ID in the **Starting VLAN ID** field. The default ID is **2**.

4 Click **ACCEPT**.

Connections

IMPORTANT: Any change to the **Connections** setting requires the SonicWall security appliance be restarted for the change to be implemented.

The **Connections** section provides the ability to fine-tune the firewall to prioritize for either optimal throughput or an increased number of simultaneous connections that are inspected by Deep-Packet Inspection (DPI) services.

TIP: A hardware platform may differ from another in the amount of memory available, which corresponds to the number of connections. For maximum DPI-SSL connections, see Connections per Appliance Model on page 248.

For specific SPI and DPI connection count maximums, refer to the latest SonicWall datasheet for your firewall platform:

- NS*a* Series - Datasheet at SonicWall NS*a* Series
- TZ Series - Datasheet at SonicWall TZ Series
- SuperMassive Series - Datasheet at SonicWall SuperMassive Series

Refer to the SonicWall resources page for more information about our Product Series. Search for high-end, mid-range, entry level, and virtual firewall details, such as Maximum connections (DPI SSL), from the **By Product Series** drop-down menu.

The maximum number of connections depends on the physical capabilities of the particular model of SonicWall security appliance. Flow Reporting does not reduce the connection count on NS*a* Series, NSA Series, and SuperMassive Series firewalls.

Mousing over the **Information** icon next to the **Connections** heading displays a popup table of the maximum number of connections for your specific SonicWall security appliance for the various configuration permutations. The table entry for your current configuration is indicated in the popup table.

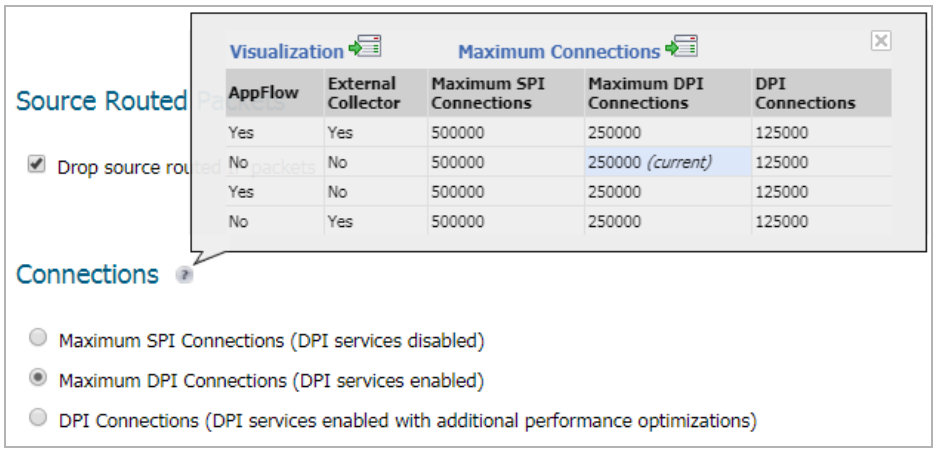

AppFlow	External Collector	Maximum SPI Connections	Maximum DPI Connections	DPI Connections
Yes	Yes	500000	250000	125000
No	No	500000	250000 *(current)*	125000
Yes	No	500000	250000	125000
No	Yes	500000	250000	125000

To configure connection services:

1 Navigate to **MANAGE | Security Configuration > Firewall Settings > Advanced Settings**.

2 Scroll to **Source Routed Packets**.

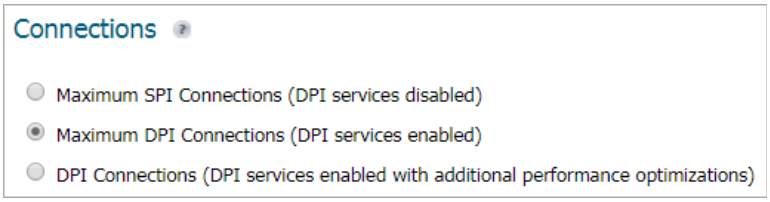

3 To display the connections for the firewall, click the **Information** icon.

4 Choose the type services to be enabled/disabled. There is no change in the level of security protection provided by the DPI Connections settings.

- **Maximum SPI Connections (DPI services disabled)** - This option (Stateful Packet Inspection) does not provide SonicWall DPI Security Services protection and optimizes the firewall for maximum number of connections with only stateful packet inspection enabled. This option should be used

by networks that require **only** stateful packet inspection, which is not recommended for most SonicWall network security appliance deployments.

- **Maximum DPI Connections (DPI services enabled)** - This is the recommended setting for most SonicWall network security appliance deployments. This option is selected by default.
- **DPI Connections (DPI services enabled with additional performance optimization)** - This option is intended for performance critical deployments. This option trades off the number of maximum DPI connections for an increased firewall DPI inspection throughput.

NOTE: If either DPI Connections option is chosen and the DPI connection count is greater than 250,000, you can have the firewall resize the DPI connection and DPI-SSL counts dynamically. For more information, see Dynamic Connection Sizing on page 18.

Dynamic Connection Sizing

NOTE: Dynamic connection sizing is supported on NS*a* Series, NSA 3600 (and higher) and SuperMassive Series network security appliances.

TIP: For the maximum number of DPI-SSL connections per platform, see Connections per Appliance Model on page 248.

If either **Maximum DPI Connections (DPI services enabled)** or **DPI Connections (DPI services enabled with additional performance optimization)** is selected for **Connections** and the DPI connection count is greater than 250,000, the **Dynamic Connection Sizing** section displays. Configuring this option allows you to have the firewall increase the number of DPI-SSL connections by 750 by reducing the number of DPI connections by 1250000 dynamically.

To configure dynamic connection sizing:

1. Navigate to **MANAGE | Security Configuration > Firewall Settings > Advanced Settings**.
2. Scroll to **Dynamic Connection Sizing**.

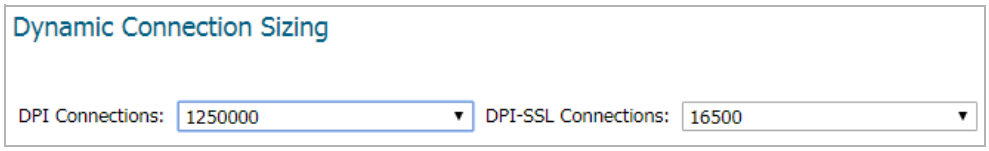

3. Do one of these:

   TIP: Changing the count in one option changes the value in the other automatically.

   - Select the maximum number of DPI connections from **DPI Connections**, in increments of 125,000.
   - Select the maximum number of DPI-SSL connections from **DPI-SSL Connections**, in increments of 750.

   For example, if the number of DPI connections selected in **DPI Connections** is **1250000**, the number of DPI-SSL connections in **DPI-SSL Connections** is **165000**. If you select **1000000** from **DPI Connections**, the number of DPI-SSL connections changes to **18000**. If you select **12000** from **DPI-SSL Connections**, the number of DPI connections changes to **2000000**.

4. Click **ACCEPT**.

Access Rule Service Options

To configure Access Rule options:

1 Navigate to **MANAGE | Security Configuration > Firewall Settings > Advanced Settings**.

2 Scroll to **Access Rule Options**.

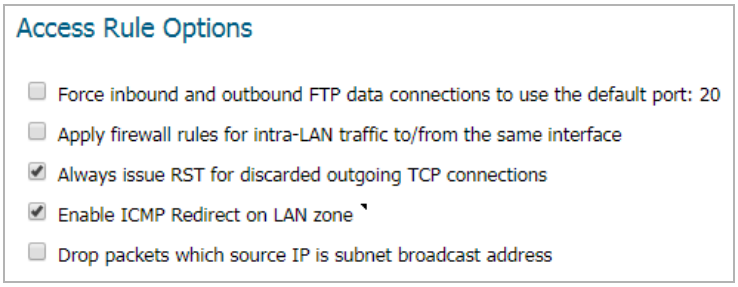

3 The default configuration allows FTP connections from port 20, but remaps outbound traffic to a port such as 1024. To enforce any FTP data connection through the security appliance must come from port 20 or the connection is dropped, select **Force inbound and outbound FTP data connections to use default port 20**. If the option is selected, the event is then logged as a log event on the security appliance. This option is not selected by default.

4 To apply firewall rules received on a LAN interface and destined for the same LAN interface, select **Apply firewall rules for intra-LAN traffic to/from the same interface**.Typically, this is only necessary when secondary LAN subnets are configured. This option is not selected by default.

5 To send an RST (reset) packet to drop the connection for discarded outgoing TCP connections, select **Always issue RST for discarded outgoing TCP connections**. This option is selected by default.

6 To redirect ICMP packets on LAN zone interfaces, select **Enable ICMP Redirect on LAN zone**. This option is selected by default.

7 To drop packets when the detected IP address is recognized as the one by the subnet, select **Drop packets which source IP is subnet broadcast address**. This option is not selected by default.

8 Click **ACCEPT**.

IP and UDP Checksum Enforcement

To configure IP and UDP checksum enforcement:

1 Navigate to **MANAGE | Security Configuration > Firewall Settings > Advanced Settings**.

2 Scroll to **IP and UDP Checksum Enforcement**.

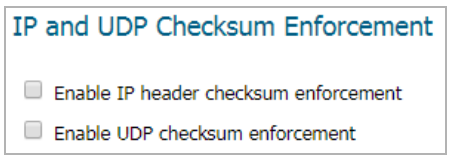

3 To drop packets with incorrect checksums in the IP header by enforcing IP header checksums, select **Enable IP header checksum enforcement**. This option is not selected by default.

4 To drop packets with incorrect checksums in the UDP header by enforcing UDP header checksums, select **enable UDP checksum enforcement** - This option is not selected by default.

5 Click **ACCEPT**.

Jumbo Frame

NOTE: Jumbo frames are supported on NSA 3600 and higher appliances.

To configure jumbo frame support:

1 Navigate to **MANAGE | Security Configuration > Firewall Settings > Advanced Settings**.

2 Scroll to **Jumbo Frame**.

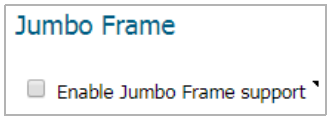

3 To enable jumbo frame support, select **Enable Jumbo Frame support.** This option is not selected by default.

Enabling this option increases throughput and reduces the number of Ethernet frames to be processed. Throughput increase may not be seen in some cases although there will be some improvement in throughput if the packets traversing are really jumbo size.

NOTE: Jumbo frame packets are 9000 kilobytes in size and increase memory requirements by a factor of 4. Interface MTUs must be changed to 9000 bytes after enabling jumbo frame support, as described in *SonicWall SonicOS 6.5 System Setup*.

4 Click **ACCEPT**.

IPv6 Advanced Configuration

To configure advanced IPv6:

1 Navigate to **MANAGE | Security Configuration > Firewall Settings > Advanced Settings**.

2 Scroll to **IPv6 Advanced Configurations**.

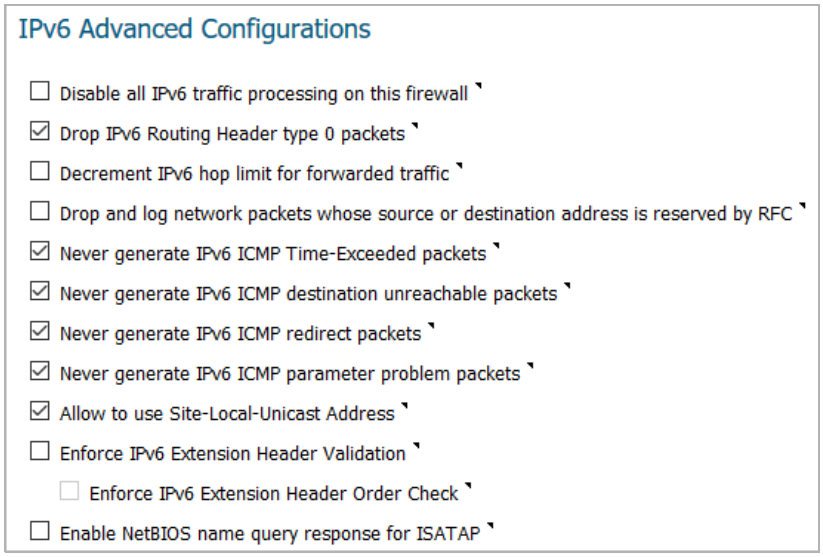

3 To disable IPv6 completely on the firewall, select **Disable all IPv6 traffic processing on this firewall**. When enabled, this option takes precedence over the other IPv6 options in this section. This option is not selected by default.

4 To prevent a potential DoS attack that exploits IPv6 Routing Header type 0 (RH0) packets, select **Drop IPv6 Routing Header type 0 packets**. When this setting is enabled, RH0 packets are dropped unless their destination is the SonicWall security appliance and their Segments Left value is 0. Segments Left specifies the number of route segments remaining before reaching the final destination. This option is selected by default. For more information, see http://tools.ietf.org/html/rfc5095.

5 To drop a packet when the hop limit has been decremented to 0, select **Decrement IPv6 hop limit for forwarded traffic**; this is similar to IPv4 TTL This option is not selected by default.

6 To reject and log network packets that have a source or destination address of the network packet defined as an address reserved for future definition and use as specified in RFC 4921 for IPv6, select **Drop and log network packets whose source or destination address is reserved by RFC**. This option is not selected by default.

7 By default, the SonicWall appliance generates IPv6 ICMP Time-Exceeded Packets that report when the appliance drops packets due to the hop limit decrementing to 0. To disable this function so the SonicWall appliance does not generate these packets, select **Never generate IPv6 ICMP Time-Exceeded packets**. This option is selected by default.

8 By default, the SonicWall appliance generates IPv6 ICMP destination unreachable packets. To disable this function so the SonicWall appliance does not generate these packets, select **Never generate IPv6 ICMP destination unreachable packets**. This option is selected by default.

9 By default, the SonicWall appliance generates redirect packets. To disable this function so the SonicWall appliance does not generate redirect packets, select **Never generate IPv6 ICMP redirect packets**. This option is selected by default.

10 By default, the SonicWall appliance generates IPv6 ICMP parameter problem packets. To disable this function; so the SonicWall appliance does not generate these packets, select **Never generate IPv6 ICMP parameter problem packets**. This option is selected by default.

11 To allow Site-Local Unicast (SLU) address, the default SonicWall appliance behavior, select **Allow to use Site-Local-Unicast Address**. This option is selected by default.

As currently defined, SLU addresses are ambiguous and can represent multiple sites. The use of SLU addresses may adversely affect network security through leaks, ambiguity, and potential misrouting. To avoid the issue, deselect the option to prevent the appliance from using SLU addresses.

12 To have the SonicWall appliance check the validity of IPv6 extension headers, select **Enforce IPv6 Extension Header Validation**. This option is selected by default.

When this option is selected, the **Enforce IPv6 Extension Header Order Check** option becomes available. (You may need to refresh the page.)

- To have the SonicWall appliance check the order of IPv6 Extension Headers, select **Enforce IPv6 Extension Header Order Check**. This option is not selected by default.

13 To have the SonicWall appliance generate a NetBIOS name in response to a broadcast ISATAP query, select **Enable NetBIOS name query response for ISATAP**. This option is not selected by default.

IMPORTANT: Select this option only when one ISATAP tunnel interface is configured.

14 Click **ACCEPT**.

Control Plane Flood Protection

To configure control plane flood protection:

1 Navigate to **MANAGE | Security Configuration > Firewall Settings > Advanced Settings**.

2 Scroll to **Control Plan Flood Protection**.



3 To have the firewall forward only control traffic destined to the firewall to the system Control Plane core (Core 0) if traffic on the Control Plane exceeds the specified threshold, select **Enable Control Plane Food Protection**, and then specify the threshold in now available **Control Flood Protection Threshold (CPU %)**. This option is not enabled by default.

To give precedence to legitimate control traffic, excess data traffic is dropped. This restriction prevents too much data traffic from reaching the Control Plane core, which can cause slow system response and potential network connection drops. The percentage configured for control traffic is guaranteed.

- Enter the flood protection threshold as a percentage in **Control Flood Protection Threshold (CPU %)**. The minimum is 5 (%), the maximum is 95, and the default is **75**.

4 Click **ACCEPT**.

2

Configuring Bandwidth Management

- Understanding Bandwidth Management on page 23
- Firewall Settings > BWM on page 25
- Configuring Global Bandwidth Management on page 28
- Advanced Bandwidth Management on page 36
- Configuring Advanced Bandwidth Management on page 40

Understanding Bandwidth Management

Bandwidth management (BWM) is a means of allocating bandwidth resources to critical applications on a network.

SonicOS offers an integrated traffic shaping mechanism through its outbound (Egress) and inbound (Ingress) BWM interfaces. Egress BWM can be applied to traffic sourced from Trusted and Public zones traveling to Untrusted and Encrypted zones. Ingress BWM can be applied to traffic sourced from Untrusted and Encrypted zones traveling to Trusted and Public zones.

The SonicWall security appliance uses BWM to control ingress and egress traffic. BWM allows network administrators to guarantee minimum bandwidth and prioritize traffic based on access rules created in the **MANAGE | Policies > Rules > Access Rules** page of the management interface. By controlling the amount of bandwidth to an application or user, you can prevent a small number of applications or users to consume all available bandwidth. Balancing the bandwidth allocated to different network traffic and then assigning priorities to traffic can improve network performance.

NOTE: Although BWM is a fully integrated Quality of Service (QoS) system, wherein classification and shaping is performed on the single SonicWall appliance, effectively eliminating the dependency on external systems and thus obviating the need for marking, it is possible to concurrently configure **BWM** and **QoS** (layer 2 and/or layer 3 marking) settings on a single Access Rule. This allows those external systems to benefit from the classification performed on the firewall even after it has already shaped the traffic. For BWM QoS details, refer to Managing Quality of Service on page 76.

BWM priority queues lists the SonicOS priority queues.

BWM priority queues

0 – Realtime	3 – Medium High	6 – Low
1 – Highest	4 – Medium	7 – Lowest
2 – High	5 – Medium Low	

Various types of bandwidth management are available and can be selected on the **MANAGE | Security Configuration > Firewall Settings > Bandwidth Management** page.

Bandwidth management types

BWM Type	Description
Advanced	Enables Advanced Bandwidth Management. Maximum egress and ingress bandwidth limitations can be configured on any interface, per interface, by configuring bandwidth objects, access rules, and application policies and attaching them to the interface.
Global	All zones can have assigned guaranteed and maximum bandwidth to services and have prioritized traffic. When global BWM is enabled on an interface, all of the traffic to and from that interface is bandwidth managed according to the priority queue. Default Global BWM queues: **2** — High **4** — Medium **6** — Low **4 Medium** is the default priority for all traffic that is not managed by an Access rule or an Application Control Policy that is BWM enabled. For traffic more than 1 Gbps, maximum bandwidth is limited to 1 Gbps because of queuing, which may limit the number of packets processed.
None	(Default) Disables BWM.

If the bandwidth management type is **None**, and there are three traffic types that are using an interface, if the link capacity of the interface is 100 Mbps, the cumulative capacity for all three types of traffic is 100 Mbps.

When **Global** bandwidth management is enabled on an interface, all traffic to and from that interface is bandwidth managed. If the available ingress and egress traffic is configured at 10 Mbps, then by default, all three traffic types are sent to the medium priority queue. The medium priority queue, by default, has a guaranteed bandwidth of 50 percent and a maximum bandwidth of 100 percent. If no **Global** bandwidth management policies are configured, the cumulative link capacity for each traffic type is 10 Mbps.

NOTE: BWM rules each consume memory for packet queuing, so the number of allowed queued packets and rules on SonicOS is limited by platform (values are subject to change).

Global uses the unused guaranteed bandwidth from other queues for maximum bandwidth. If there is only default or single-queue traffic and all the queues have a total of 100% allocated as guaranteed, **Global** uses the unused global bandwidth from other queues to give you up to maximum bandwidth for the default/single queue

Glossary

Bandwidth Management (BWM)	Any of a variety of algorithms or methods used to shape traffic or police traffic. Shaping often refers to the management of outbound traffic, while policing often refers to the management of inbound traffic (also known as admission control). There are many different methods of bandwidth management, including various queuing and discarding techniques, each with their own design strengths. SonicWall employs a Token Based Class Based Queuing method for inbound and outbound BWM, as well as a discard mechanism for certain types of inbound traffic.
Guaranteed Bandwidth	A declared percentage of the total available bandwidth on an interface which is always granted to a certain class of traffic. Applicable to both inbound and outbound BWM. The total Guaranteed Bandwidth across all BWM rules cannot exceed 100% of the total available bandwidth. SonicOS 5.0 and higher enhances the Bandwidth Management feature to provide rate limiting functionality. You can create traffic policies that specify maximum rates for Layer 2, 3, or 4 network traffic. The Guaranteed Bandwidth can also be set to 0%.

Ingress BWM	The ability to shape the rate at which traffic enters a particular interface. For TCP traffic, actual shaping occurs when the rate of the ingress flow is adjusted by the TCP Window Adjustment mechanism. For UDP traffic, a discard mechanism is used as UDP has no native feedback controls.
Maximum Bandwidth:	A declared percentage of the total available bandwidth on an interface defining the maximum bandwidth to be allowed to a certain class of traffic. Applicable to both inbound and outbound BWM. Used as a throttling mechanism to specify a bandwidth rate limit. The Bandwidth Management feature is enhanced to provide rate-limiting functionality. You can create traffic policies that specify maximum rates for Layer 2, 3, or 4 network traffic. This enables bandwidth management in cases where the primary WAN link fails over to a secondary connection that cannot handle as much traffic. The Maximum Bandwidth can be set to 0%, which prevents all traffic.
Egress BWM	Conditioning the rate at which traffic is sent out an interface. Outbound BWM uses a credit (or token) based queuing system with 8 priority rings to service different types of traffic, as classified by Access Rules.
Priority	An additional dimension used in the classification of traffic. SonicOS uses eight priority values (0 = highest, 7 = lowest) for the queue structure used for BWM. Queues are serviced in the order of their priority.
Queuing	To effectively make use of the available bandwidth on a link. Queues are commonly employed to sort and separately manage traffic after it has been classified.

Firewall Settings > BWM

Topics:

- Configuring Global Bandwidth Management Settings on page 25
- Action Objects on page 28

Configuring Global Bandwidth Management Settings

BWM works by first enabling bandwidth management in the **MANAGE | Security Configuration > Firewall Settings > Bandwidth Management** page, enabling BWM on an interface/firewall/app rule, and then allocating the available bandwidth for that interface on the ingress and egress traffic. It then assigns individual limits for each class of network traffic. By assigning priorities to network traffic, applications requiring a quick response time, such as Telnet, can take precedence over traffic requiring less response time, such as FTP.

To view the BWM configuration, navigate to the **MANAGE | Security Configuration > Firewall Settings > Bandwidth Management** page.

NOTE: The default settings for this page consists of three priorities with preconfigured guaranteed and maximum bandwidth. The medium priority has the highest guaranteed value as this priority queue is used by default for all traffic not governed by a BWM-enabled policy.

NOTE: The defaults are set by SonicWall to provide BWM ease-of-use. It is recommended that you review your specific bandwidth needs and enter the values on this page accordingly.

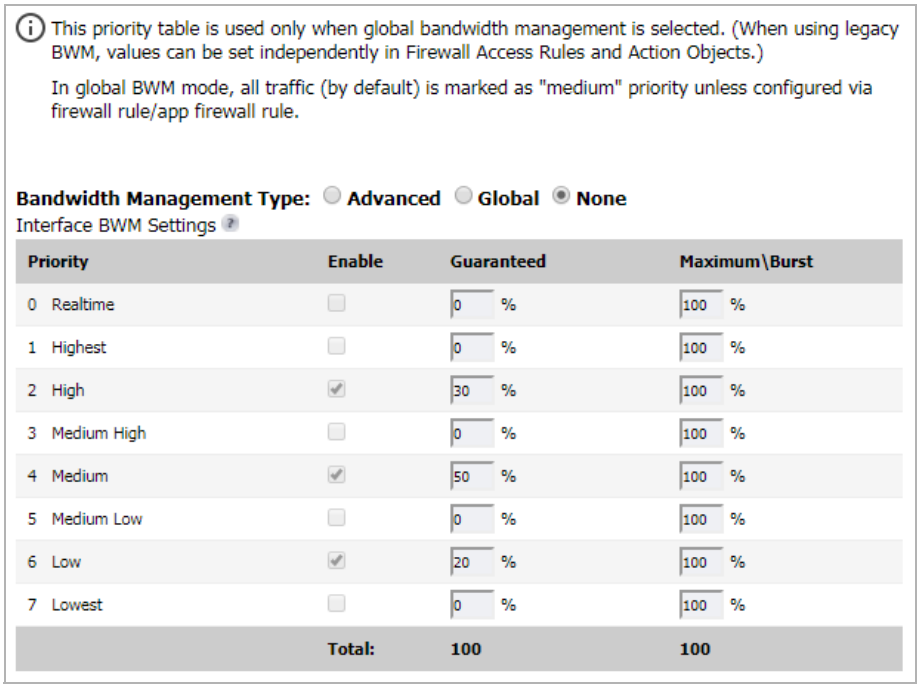
This priority table is used only when global bandwidth management is selected. (When using legacy BWM, values can be set independently in Firewall Access Rules and Action Objects.)

In global BWM mode, all traffic (by default) is marked as "medium" priority unless configured via firewall rule/app firewall rule.

Bandwidth Management Type: Advanced Global None

Interface BWM Settings

Priority	Enable	Guaranteed	Maximum\Burst
0 Realtime	☐	0 %	100 %
1 Highest	☐	0 %	100 %
2 High	☑	30 %	100 %
3 Medium High	☐	0 %	100 %
4 Medium	☑	50 %	100 %
5 Medium Low	☐	0 %	100 %
6 Low	☑	20 %	100 %
7 Lowest	☐	0 %	100 %
	Total:	**100**	**100**

- **Bandwidth Management Type** option:

  IMPORTANT: When you change the **Bandwidth Management Type** from:

  - **Global** to **Advanced**, the default BWM actions that are in use in any App Rules policies are automatically converted to **Advanced BWM** settings.
  - **Advanced** to **Global**, the default BWM actions are converted to **BWM Global-Medium**.

  The firewall does not store your previous action priority levels when you switch the **Type** back and forth. You can view the conversions on the **MANAGE | Policies > Rules > Application Control** page.

  - **Advanced** — Any zone can have guaranteed and maximum bandwidth and prioritized traffic assigned per interface.
  - **Global** — All zones can have assigned guaranteed and maximum bandwidth to services and have prioritized traffic. For traffic more than 1 Gbps, maximum bandwidth is limited to 1 Gbps.
  - **None** — Disables BWM. This is the default.

- **Interface BWM Settings** — Mousing over the **Question Mark** icon displays a table showing whether the BWM settings are disabled or enabled for ingress and egress on the various interfaces:

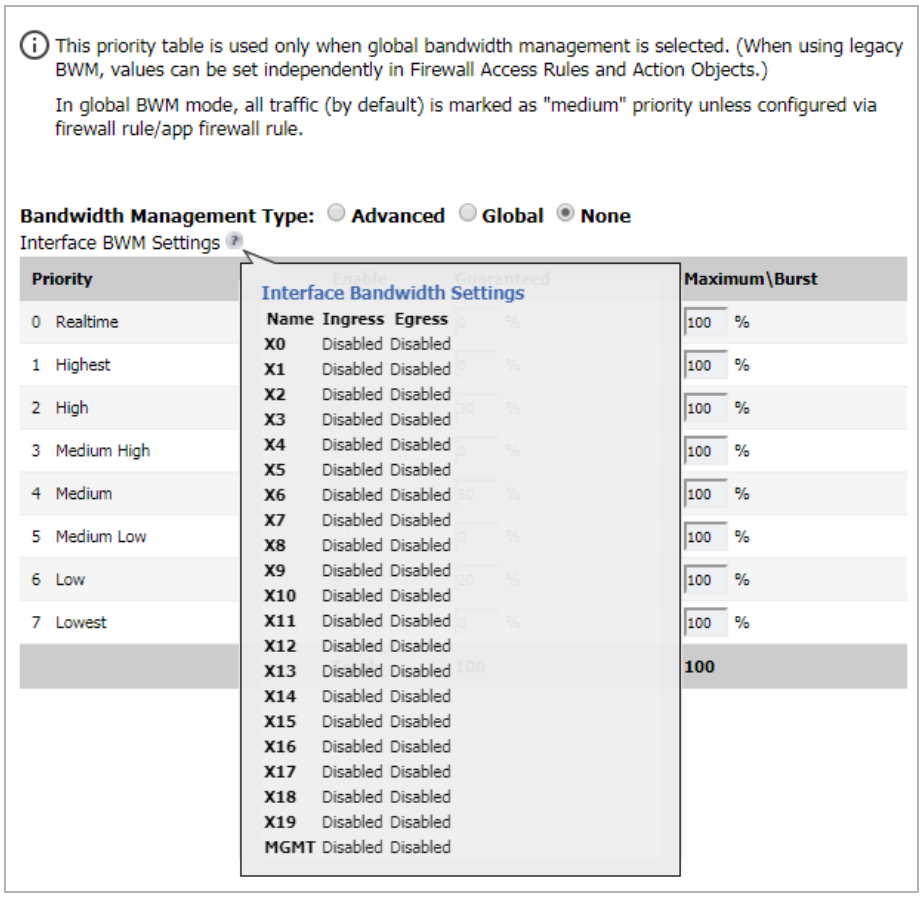

- **Global Priority** Bandwidth table — Displays this information about the priorities:

  NOTE: This table is used only when **Global** BWM is selected. The table is dimmed when **Advanced** or **None** is selected.

  - **Priority** — Displays the priority number and name, from **0 Realtime** through **7 Lowest**.
  - **Enable** — When a priority is selected, the priority queue is enabled for that priority.
  - **Guaranteed** — Enables the guaranteed rate, as a percentage, for the enabled priority. The configured bandwidth on an interface is used in calculating the absolute value.

    The corresponding **Enable** checkbox must be checked for the rate to take effect. By default, only these priorities and their guaranteed percentages are enabled:

    - **2 High** 30%
    - **4 Medium** 50%
    - **6 Low** 20%

    TIP: You cannot disable priority **4 Medium**, but you can change its percentage.

    The sum of all guaranteed bandwidth must not exceed 100%. If the bandwidth exceeds 100%, the **Total** number becomes red. Also, the guaranteed bandwidth must not be greater than the maximum bandwidth per queue.

- **Maximum\Burst** — Enables the maximum/burst rate, as a percentage, for the enabled priority. The corresponding **Enable** checkbox must be checked for the rate to take effect.

Action Objects

Action Objects define how the App Rules policy reacts to matching events. You can customize an action or select one of the predefined default actions. The predefined actions are displayed in the App Control Policy Settings page when you add or edit a policy from the App Rules page.

Custom BWM actions behave differently than the default BWM actions. Custom BWM actions are configured by adding a new action object from the **MANAGE | Policies > Objects** > **Action Objects** page and then selecting the Bandwidth Management action type. Custom BWM actions and policies using them retain their priority level setting when the Bandwidth Management Type is changed from **Global** to **Advanced**, and from **Advanced** to **Global**.

A number of BWM action options are also available in the predefined, default action list. The BWM action options change depending on the Bandwidth Management Type setting on the **Firewall Settings** > **Bandwidth Management** page. If the **Bandwidth Management Type** is set to:

- **Global**, all eight levels of BWM are available.
- **Advanced**, no priorities are set. The priorities are set by configuring a bandwidth object under **MANAGE | Policies > Objects** > **Bandwidth Objects**.

Adding a policy: Default actions lists the predefined default actions that are available when adding a policy.

Adding a policy: Default actions

	If BWM Type =
Global	**Advanced**
BWM Global-Realtime	Advanced BWM High
BWM Global-Highest	Advanced BWM Medium
BWM Global-High	Advanced BWM Low
BWM Global-Medium High	
BWM Global-Medium	
BWM Global-Medium Low	
BWM Global-Low	
BWM Global-Lowest	

Configuring Global Bandwidth Management

IMPORTANT: BWM must be enabled on **MANAGE | Security Configuration > Firewall Settings > Bandwidth Management** first.

Global Bandwidth Management can be configured using the following methods:

- Configuring Global Bandwidth Management on page 29
- Configuring Global BWM on an Interface on page 29
- Configuring Global BWM in an Access Rule on page 31
- Configuring Global BWM in an Action Object on page 31

- Configuring Application Rules on page 32
- Configuring App Flow Monitor on page 34

Configuring Global Bandwidth Management

To set the Bandwidth Management type to Global:

1 Navigate to **MANAGE | Security Configuration > Firewall Settings > Bandwidth Management**.

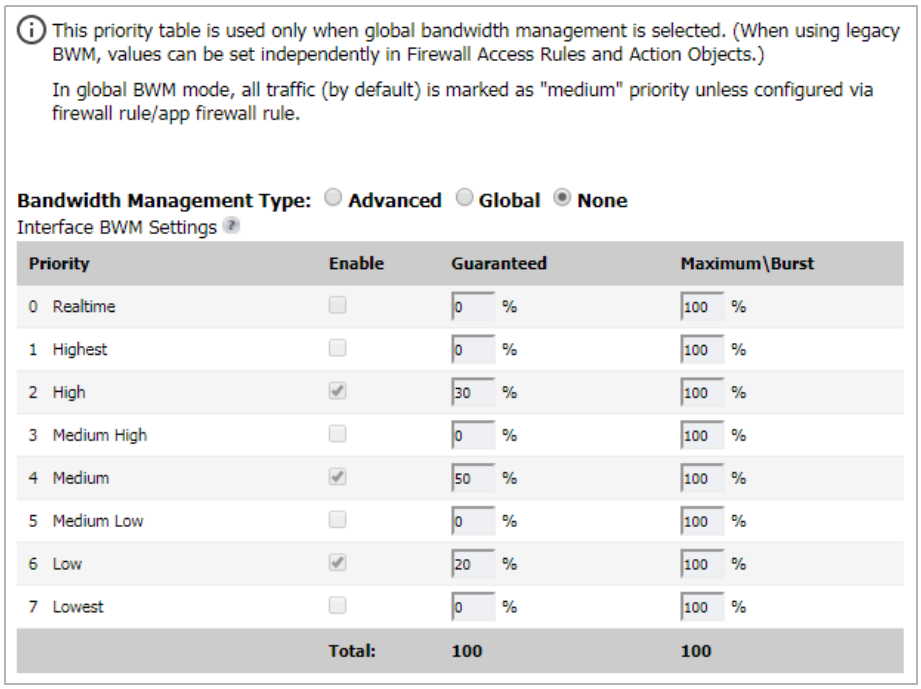

2 Set the **Bandwidth Management Type** option to **Global**.

3 Enable the priorities that you want by selecting the appropriate checkboxes in the **Enable** column.

NOTE: You must enable the priorities on this page to be able to configure these priorities in Access Rules, App Rules, and Action Objects.

4 Enter the **Guaranteed** bandwidth percentage that you want for each selected priority. The total amount cannot exceed 100%.

5 Enter the **Maximum\Burst** bandwidth percentage that you want for each selected priority.

6 CLICK **ACCEPT**.

Configuring Global BWM on an Interface

IMPORTANT: Global BWM must be enabled on **Firewall Settings > Bandwidth Management** first, as described in Configuring Global Bandwidth Management on page 29.

To configure BWM on an interface:

1 Navigate to **MANAGE | System Setup > Network > Interfaces**.

2 In the **Interface Settings** table, click the **Edit** button under the **Configure** column for the appropriate interface. The **Edit Interface** dialog displays.

3 Click **Advanced**.

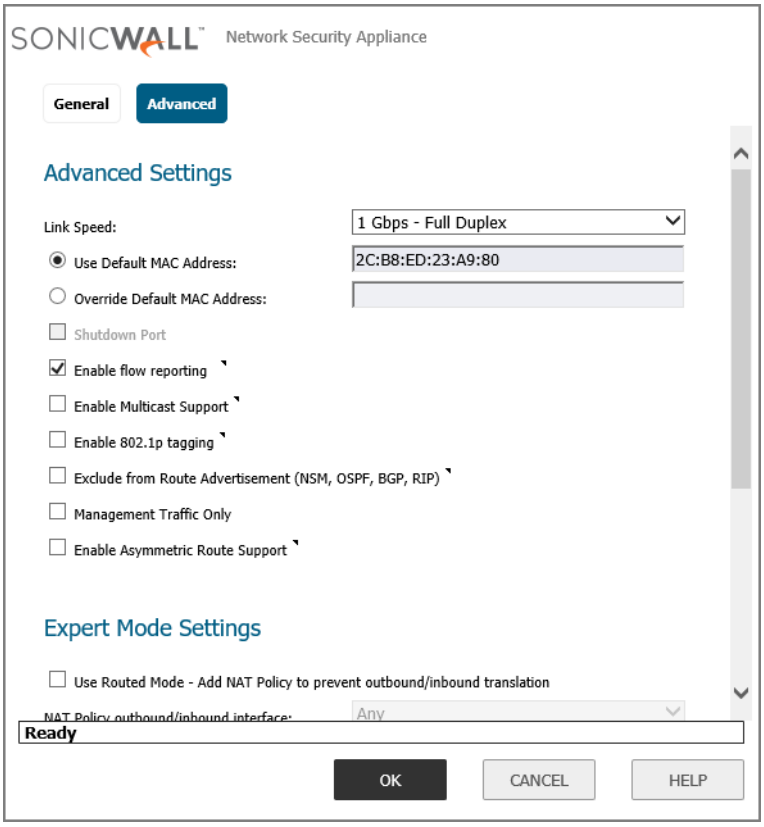

NOTE: Displayed options may differ depending on how the interface is configured.

4 Scroll to **Bandwidth Management**.

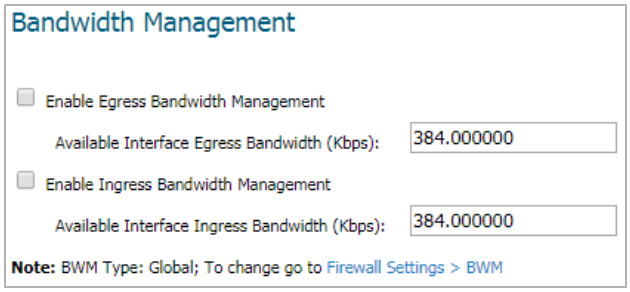

5 Select either or both **Enable Interface Egress Bandwidth Limitation** and **Enable Interface Ingress Bandwidth Limitation.** These options are not selected by default.

When either or both of these options are selected, if a there isn't a corresponding Access Rule or App Rule, the total egress traffic on the interface is limited to the amount specified in the **Enable Interface Ingress Bandwidth Limitation (kbps)** field.

When neither option is selected, no bandwidth limitation is set at the interface level, but egress traffic can still be shaped using other options.

6 In the **Maximum Interface Egress/Ingress Bandwidth (Kbps)** field(s), enter the total bandwidth available for all egress/ingress traffic in Kbps. The default is **384.000000** Kbps.

7 Click **OK**.

Configuring Global BWM in an Access Rule

IMPORTANT: **Global** BWM must be enabled on **Firewall Settings > Bandwidth Management** first, as described in Configuring Global Bandwidth Management on page 29.

You can configure BWM in each Access Rules. This method configures the direction in which to apply BWM and sets the priority queue.

IMPORTANT: Before you can configure any priorities in an Access Rule, you must first enable the priorities that you want to use on the **Firewall Settings > Bandwidth Management** page. Refer to this page to determine which priorities are enabled. If you select a Bandwidth Priority that is not enabled on the **Firewall Settings > Bandwidth Management** page, the traffic is automatically mapped to priority **4 Medium**. See Configuring Global Bandwidth Management on page 29.

Priorities are listed in the **Access Rules** dialog **Bandwidth Priority** table; see BWM priority queues.

To configure Global BWM in an Access Rule:

1 Navigate to the **MANAGE | Policies > Rules > Access Rules** page.

2 Click the **Edit** icon for the rule you want to edit. The **Edit Rule** dialog displays.

3 Click **BWM**.

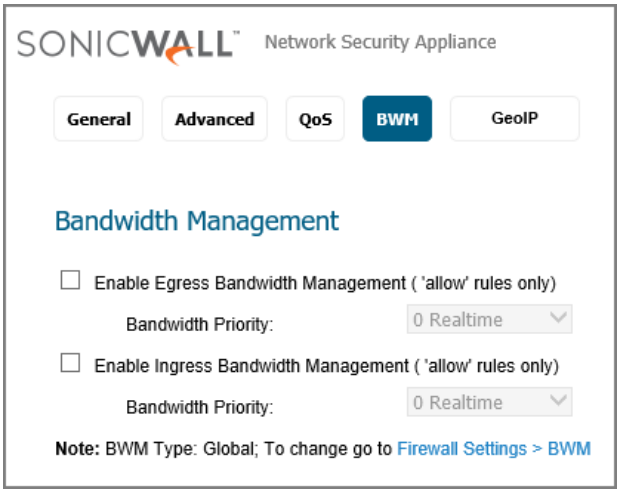

4 Select either or both **Enable Egress Bandwidth Management** and **Enable Ingress Bandwidth Management**. These options are not selected by default.

   a In the appropriate Bandwidth Priority: field(s), enter the priority level from 0 for **Realtime** to 7 for **Lowest**. The priority levels are **1 Highest**, **2 High**, **3 Medium High**, **4 Medium**, **5 Medium Low**, **6 Low**, **and 7 Lowest**.

5 Click **OK**.

Configuring Global BWM in an Action Object

IMPORTANT: **Global** BWM must be enabled on **MANAGE | Security Configuration > Firewall Settings > Bandwidth Management** first, as described in Configuring Global Bandwidth Management.

If you do not want to use the predefined Global BWM actions or policies, you can create new ones that fit your needs.

To create a new Global BWM action object:

1 Navigate to the **MANAGE | Policies > Objects > Action Objects** page.

2 Click the **Add** icon at the top of the **Action Objects** table. The **Add/Edit Action Object** dialog displays.

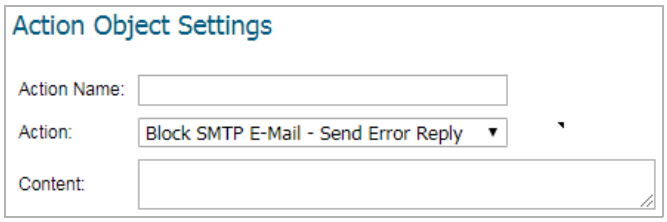

3 In the **Action Name** field, enter a name for the action object.

4 In the **Action** drop-down menu, select **Bandwidth Management** to control and monitor application-level bandwidth usage. The options on the dialog change.

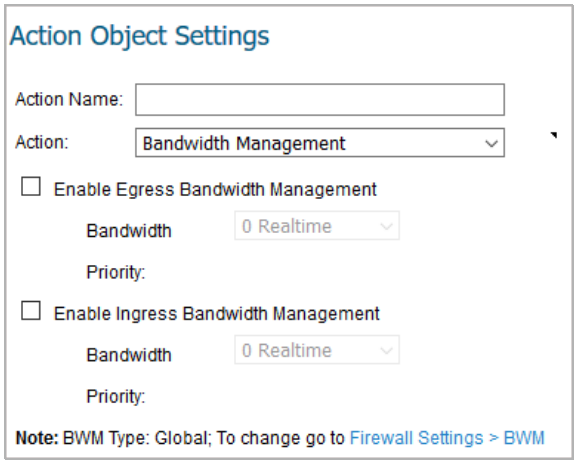

5 To specify BWM by priority, select either or both **Enable Egress Bandwidth Management** and **Enable Ingress Bandwidth Management**. These options are not selected by default.

6 Select the appropriate bandwidth priority from the **Bandwidth Priority** drop-down menu(s). The highest, and default, priority is **0 Realtime**. The lowest priority is **7 Lowest**.

7 Click **OK**.

Configuring Application Rules

Configuring BWM in an Application Rule allows you to create policies that regulate bandwidth consumption by specific file types within a protocol, while allowing other file types to use unlimited bandwidth. This enables you to distinguish between desirable and undesirable traffic within the same protocol.

Application Rule BWM supports the following **Policy Types**:

- SMTP Client
- HTTP client
- HTTP Server
- FTP Client
- FTP Client File Upload
- FTP Client File Download
- FTP Data Transfer
- POP3 Client
- POP3 Server
- Custom Policy
- IPS Content
- App Control Content
- CFS

NOTE: You must first enable BWM before you can configure BWM in an Application Rule.

Before you configure BWM in an App Rule:

1 Enable the priorities you want to use in **MANAGE | Security Configuration > Firewall Settings > Bandwidth Management**. See Configuring Global Bandwidth Management on page 29.

2 Enable BWM in an **Action Object**. See the Configuring Global BWM in an Action Object on page 31.

3 Configure BWM on the **Interface**. See the Configuring Global BWM on an Interface on page 29.

To configure BWM in an Application Rule:

1 Navigate to the **MANAGE | Policies > Rules > App Rules** page.

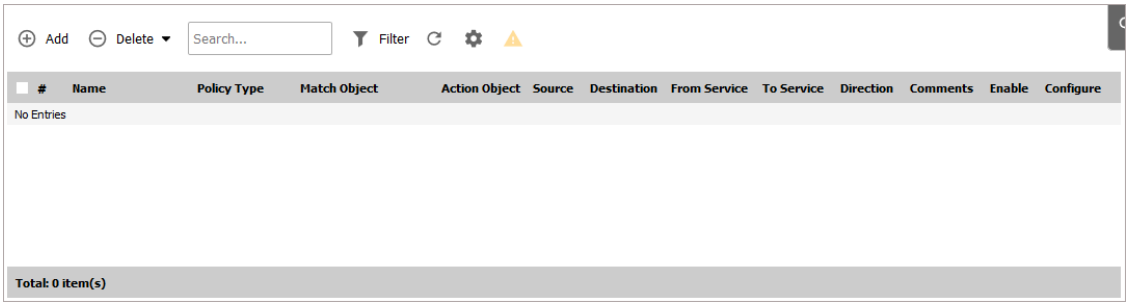

2 Click the **Add** icon. The **Edit App Control Policy** dialog displays.

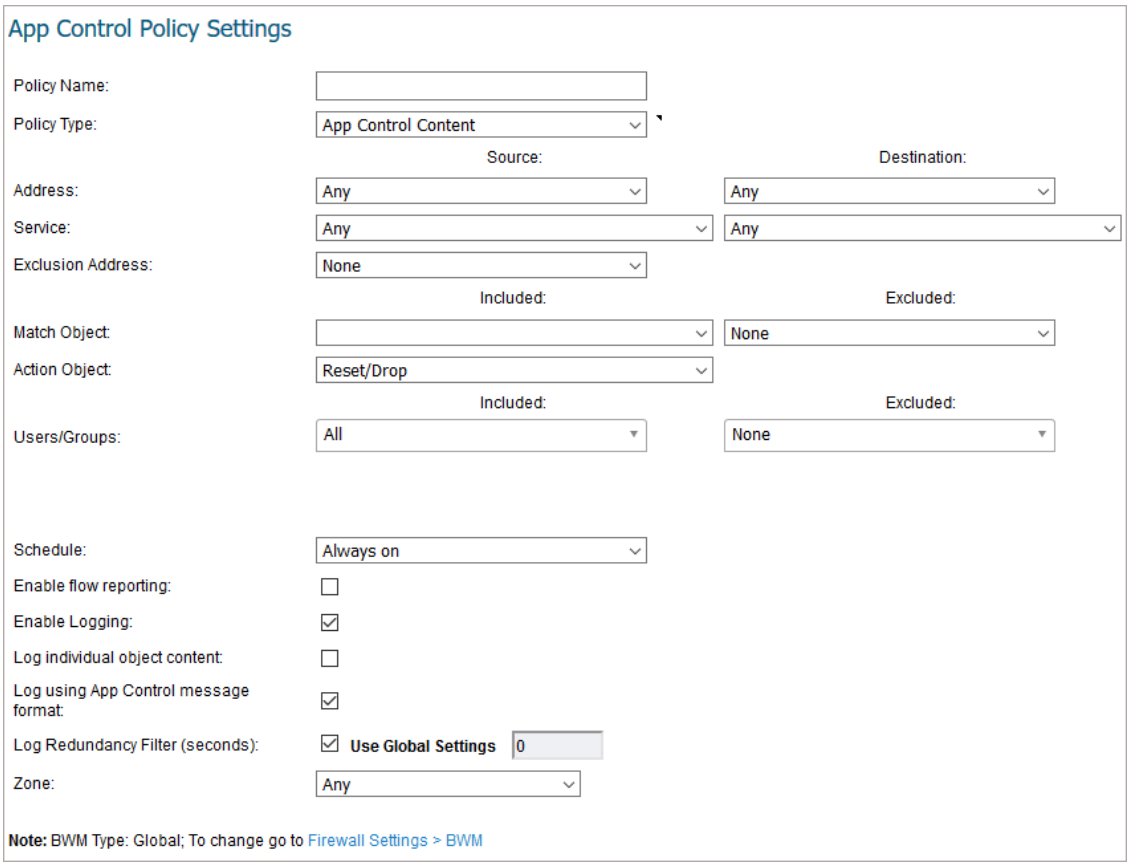

3 Under **App Control Policy Settings**, enter a meaningful name in the **Policy Name** field.

4 From **Action Object**, select the BWM action object that you want. Configure the rest of the settings as described 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.

5 Click **OK**.

Configuring App Flow Monitor

BWM can also be configured from the **INVESTIGATE | Logs > AppFlow Logs** page by selecting a service type application or a signature type application and then clicking the **Create** button. The Bandwidth Management options available there depend on the enabled priority levels in the Global Priority Queue table on the **Firewall Settings** > **Bandwidth Management** page. The priority levels enabled by default are High, Medium, and Low.

NOTE: You must have SonicWall Application Visualization enabled before proceeding.

To configure BWM using the App Flow Monitor:

1 Go to **INVESTIGATE | Logs > Appflow Logs**.

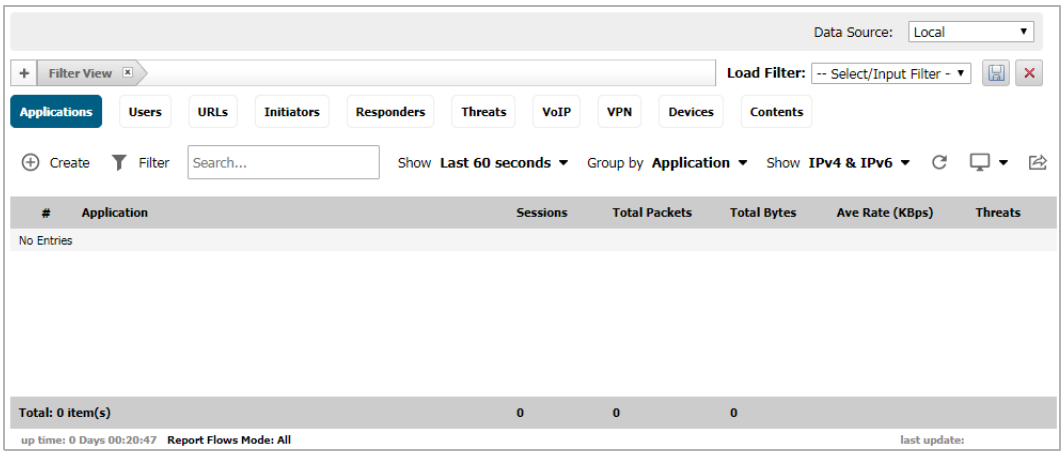

2 Check the service-based applications or signature-based applications to which you want to apply global BWM.

NOTE: General applications cannot be selected. Service-based applications and signature-based applications cannot be mixed in a single rule.

NOTE: Creating a rule for service-based applications results in creating a firewall access rule, and creating a rule for signature-based applications creates an application control policy.

3 Click **Create Rule**. The **Create Rule** dialog displays. There are slight differences between rules for service-based application options and for signature-based application options.

Service-based Application Options

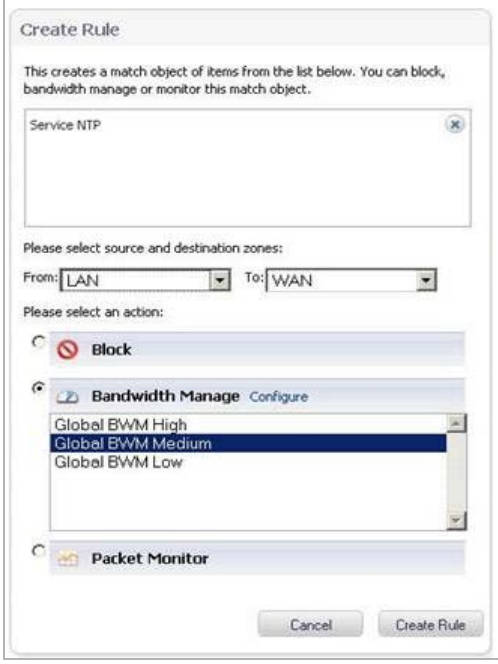

Signature-based Applications Options

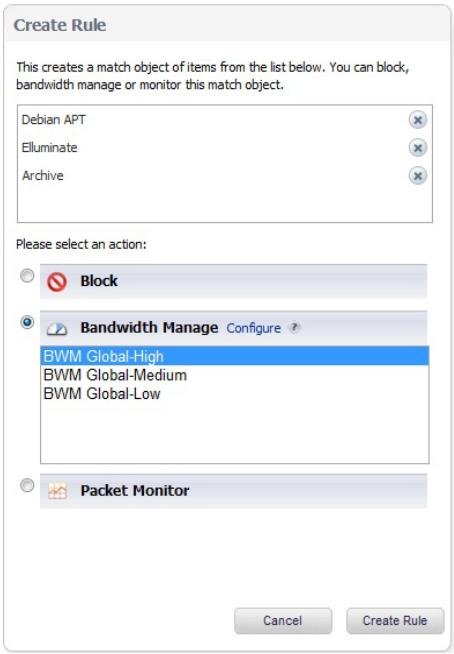

4 Select the **Bandwidth Manage** radio button.

5 Select a global BWM priority.

6 Click **Create Rule**. A confirmation dialog displays. There are slight differences between the items created for service-based application options and for signature-based application options.

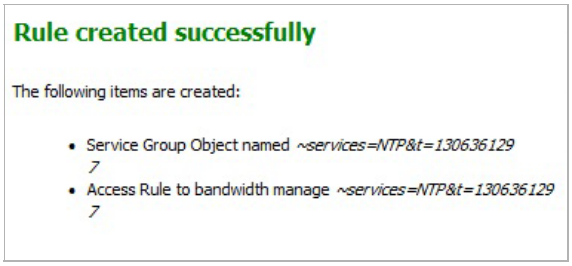

Service-based Application Successful

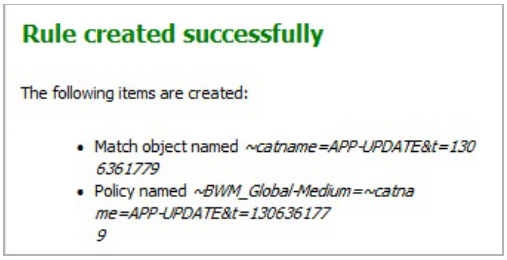

Signature-based Applications Successful

7 Click **OK**.

8 To verify that the rule was created, navigate to:

- **MANAGE | Policies > Rules > Access Rules** page for service-based applications.
- **MANAGE | Policies > Rules > App Control** for signature-based applications.

NOTE: For service-based applications, the new rule is identified with a **Tack** icon in the **Comment** column and a prefix in **Service** column of `~services=<service name>`. For example, `~services=NTP&t=1306361297`.

For signature-based applications, the new rule is identified with a prefix, `~BWM_Global-<priority>=~catname=<app_name>` in the **Name** column and a prefix in the **Object** column of `~catname=<app_name>`.

Advanced Bandwidth Management

Advanced Bandwidth Management allows you to manage specific classes of traffic based on their priority and maximum bandwidth settings. Advanced Bandwidth Management consists of three major components:

- **Classifier** – classifies packets that pass through the firewall into the appropriate traffic class.
- **Estimator** – estimates and calculates the bandwidth used by a traffic class during a time interval to determine if that traffic class has available bandwidth.
- **Scheduler** – schedules traffic for transmission based on the bandwidth status of the traffic class provided by the estimator.

Advanced Bandwidth Management: Basic concepts illustrates the basic concepts of Advanced Bandwidth Management.

Advanced Bandwidth Management: Basic concepts

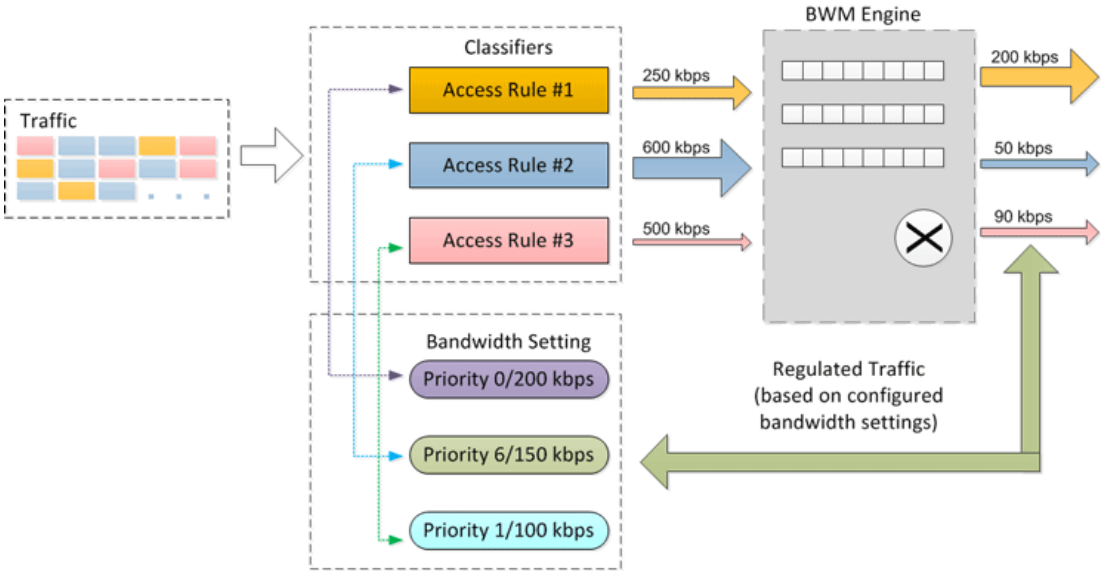

Bandwidth management configuration is based on policies that specify bandwidth limitations for traffic classes. A complete bandwidth management policy consists of two parts: a classifier and a bandwidth rule.

A **bandwidth rule** specifies the actual parameters, such as priority, guaranteed bandwidth, maximum bandwidth, and per-IP bandwidth management, and is configured in a bandwidth object. Bandwidth rules identify and organize packets into traffic classes by matching specific criteria.

A **classifier** is an access rule or application rule in which a bandwidth object is enabled. Access rules and application rules are configured for specific interfaces or interface zones.

The first step in bandwidth management is that all packets that pass through the SonicOS firewall are assigned a classifier (class tag). The classifiers identify packets as belonging to a particular traffic class. Classified packets are then passed to the BWM engine for policing and shaping. The SonicOS uses two types of classifiers:

- Access Rules
- Application Rules

A rule that has sub elements is known as a parent rule.

Configuring a bandwidth object: Parameters shows the parameters that are configured in a bandwidth object:

Configuring a bandwidth object: Parameters

Name	Description
Guaranteed Bandwidth	The bandwidth that is guaranteed to be provided for a particular traffic class.
Maximum Bandwidth	The maximum bandwidth that a traffic class can utilize.
Traffic Priority	The priority of the traffic class. • **0** – highest priority • **7** – lowest priority
Violation Action	The firewall action that occurs when traffic exceeds the maximum bandwidth. • **Delay** – packets are queued and sent when possible. • **Drop** – packets are dropped immediately.
Enable Per-IP Bandwidth Management	The elemental feature that enables the firewall to support time-critical traffic, such as voice and video, effectively. When per-IP BWM is enabled, the elemental bandwidth settings are applied to each individual IP under its parent rule.

After packets have been tagged with a specific traffic class, the BWM engine gathers them for policing and shaping based on the bandwidth settings that have been defined in a bandwidth object, enabled in an access rule, and attached to application rules.

Classifiers also identify the direction of packets in the traffic flow. Classifiers can be set for either the egress, ingress, or both directions. For Bandwidth Management, the terms ingress and egress are defined as follows:

- **Ingress** – Traffic from initiator to responder in a particular traffic flow.
- **Egress** – Traffic from responder to initiator in a particular traffic flow.

For example, a client behind Interface X0 has a connection to a server which is behind Interface X1. Direction of traffic shows:

- Direction of traffic flow in each direction for client and server
- Direction of traffic on each interface
- Direction indicated by the BWM classifier

Direction of traffic

Direction of Traffic Flow	Direction of Interface X0	Direction of Interface X1	BWM Classifier
Client to Server	Egress	Ingress	Egress
Server to Client	Ingress	Egress	Ingress

To be compatible with traditional bandwidth management settings in WAN zones, the terms inbound and outbound are still supported to define traffic direction. These terms are only applicable to active WAN zone interfaces.

- **Outbound** – Traffic from LAN\DMZ zone to WAN zone (Egress).
- **Inbound** – Traffic from WAN zone to LAN\DMZ zone (Ingress).

Topics:

- Elemental Bandwidth Settings on page 38
- Zone-Free Bandwidth Management on page 39
- Weighted Fair Queuing on page 39

Elemental Bandwidth Settings

Elemental bandwidth settings provide a method of allowing a single BWM rule to apply to the individual elements of that rule. Per-IP Bandwidth Management is an “Elemental” feature that is a sub-option of Bandwidth Object. When Per-IP BWM is enabled, the elemental bandwidth settings are applied to each individual IP under its parent rule.

The Elemental Bandwidth Settings feature enables a bandwidth object to be applied to individual elements under a parent traffic class. Elemental Bandwidth Settings is a sub-option of **MANAGE | Policies > Objects > Bandwidth Objects**, the parent rule or traffic class.

To display the Elemental Bandwidth Object dialog:

1 Click **+ Add** to see the **Bandwidth Object Settings**.

2 Click **Elemental** and click the checkbox next to **Enable Per-IP Bandwidth Management.**

3 Enter the **Maximum Bandwidth** in either **kbps** (default) or **Mbps** from the drop-down and click **OK**.



The following table shows the parameters that are configured under **Elemental Bandwidth Settings**; 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.

Elemental Bandwidth settings: Parameters

Name	Description
Enable Per-IP Bandwidth Management	When enabled, the maximum elemental bandwidth setting applies to each IP address under the parent traffic class, which allows the firewall to support time-critical traffic, such as voice and video, effectively.
Maximum Bandwidth	The maximum elemental bandwidth that can be allocated to an IP address under the parent traffic class. The maximum elemental bandwidth cannot be greater than the maximum bandwidth of its parent class.

When you enable Per-IP Bandwidth Management, each individual IP under its parent rule will be applied to the elemental bandwidth settings.

Zone-Free Bandwidth Management

The zone-free bandwidth management feature enables bandwidth management on all interfaces regardless of their zone assignments. Previously, bandwidth management only applied to these zones:

- LAN/DMZ to WAN/VPN
- WAN/VPN to LAN/DMZ

In SonicOS 6.2 and above, zone-free bandwidth management can be performed across all interfaces regardless of zone.

Zone-free bandwidth management allows you to configure the maximum bandwidth limitation independently, in either the ingress or egress direction, or both, and apply it to any interfaces using Access Rules and Application Rules.

NOTE: Interface bandwidth limitation is only available on physical interfaces. Failover and load balancing configuration does not affect interface bandwidth limitations.

Weighted Fair Queuing

Traditionally, SonicOS bandwidth management distributes traffic to 8 queues based on the priority of the traffic class of the packets. These 8 queues operate with strict priority queuing. Packets with the highest priority are always transmitted first.

Strict priority queuing can cause high priority traffic to monopolize all of the available bandwidth on an interface, and low priority traffic will consequently be stuck in its queue indefinitely. Under strict priority queuing, the scheduler always gives precedence to higher priority queues. This can result in bandwidth starvation to lower priority queues.

Weighted Fair queuing (WFQ) alleviates the problem of bandwidth starvation by servicing packets from each queue in a round robin manner, so that all queues are serviced fairly within a given time interval. High priority queues get more service and lower priority queues get less service. No queue gets all the service because of its high priority, and no queue is left unserviced because of its low priority.

For example, Traffic Class A is configured as Priority 1 with a maximum bandwidth of 400 kbps. Traffic Class B is configured as Priority 3 with a maximum bandwidth of 600 kbps. Both traffic classes are queued to an interface that has a maximum bandwidth of only 500kbps. Both queues will be serviced based on their priority in a round robin manner. So, both queues will be serviced, but Traffic Class A will be transmitted faster than Traffic Class B.

Shaped bandwidth for consecutive sampling intervals shows the shaped bandwidth for each consecutive sampling interval:

Shaped bandwidth for consecutive sampling intervals

Sampling Interval	Traffic Class A		Traffic Class B	
	Incoming kbps	Shaped kbps	Incoming kbps	Shaped kbps
1	500	380	500	120
2	500	350	500	150
3	400	300	800	200
4	600	400	400	100
5	200	180	600	320
6	200	200	250	250

Configuring Advanced Bandwidth Management

Topics:

- Enabling Advanced Bandwidth Management on page 40
- Configuring Bandwidth Policies on page 41
- Setting Interface Bandwidth Limitations with Advanced BWM on page 47
- Setting Interface Bandwidth Limitations with Advanced BWM on page 47
- Setting Interface Bandwidth Limitations with Global BWM on page 49

Enabling Advanced Bandwidth Management

To enable Advanced bandwidth management:

1 Navigate to **MANAGE | Security Configuration > Firewall Settings > Bandwidth Management**.

2 Set the **Bandwidth Management Type** option to **Advanced**.

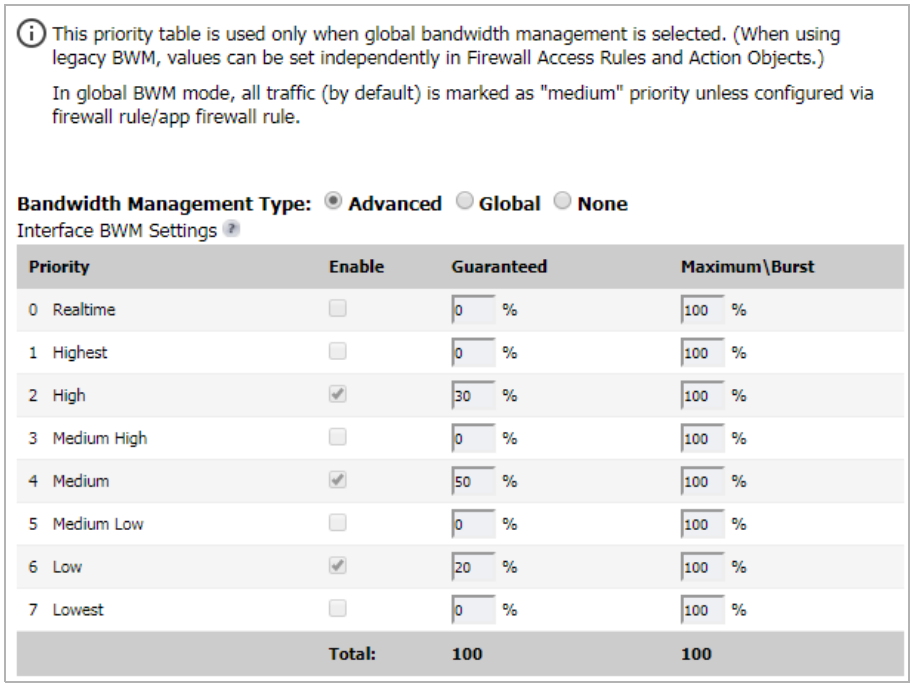

3 Click **ACCEPT**.

NOTE: When **Advanced** BWM is selected, the priorities fields are disabled and cannot be set here. Under Advanced BWM, the priorities are set in bandwidth policies. See Configuring Bandwidth Policies on page 41.

Enabling Global Bandwidth Management

To enable Global bandwidth management:

1 Navigate to **MANAGE | Security Configuration** > **Firewall Settings > Bandwidth Mangement**.

2 Set the **Bandwidth Management Type** option to **Global**.

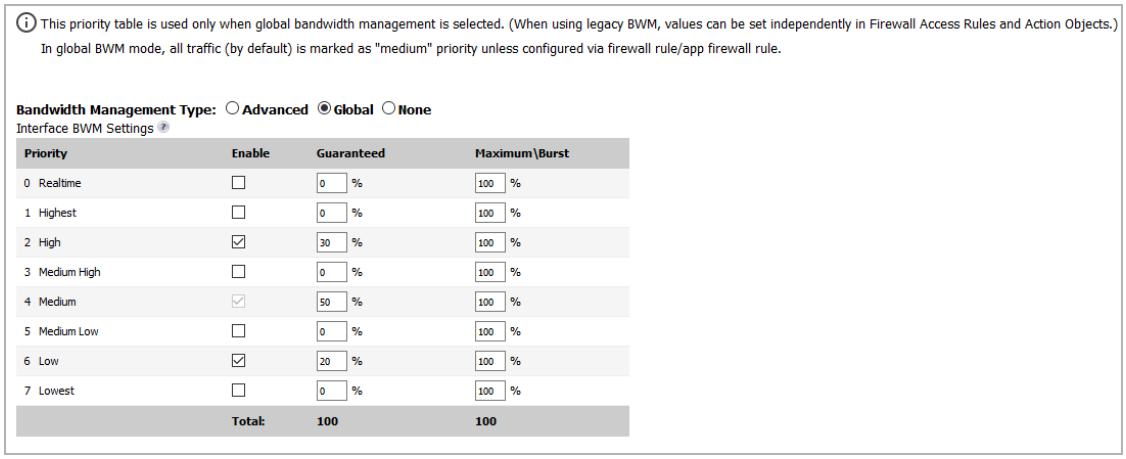

By default, only three priorities are set:

- **2 High** — 30% guaranteed
- **4 Medium** — 50% guaranteed (cannot be deselected, but the percentages can be changed)
- **6 Low** — 20% guaranteed

3 Enable or disable any priority (except 4 Medium) by selecting its Enable checkbox.

4 Specify the guaranteed bandwidth for each selected priority by entering a percentage in its Guaranteed field.

IMPORTANT: The sum of all guaranteed bandwidth must not exceed 100%. If the bandwidth exceeds 100%, the **Total** number becomes red. Also, the guaranteed bandwidth must not be greater than the maximum bandwidth per queue.

5 Specify the maximum (burst) bandwidth for each selected priority by entering a percentage in its Maximum/Burst field.

TIP: All the priorities can have the same **Maximum/Burst** bandwidth.

6 Click **ACCEPT**.

Configuring Bandwidth Policies

Topics:

- Configuring a Bandwidth Object on page 42
- Enabling Elemental Bandwidth Management on page 43
- Enabling a Bandwidth Object in an Access Rule on page 44
- Enabling a Bandwidth Priority in an Access Rule on page 44

- Enabling a Bandwidth Object in an Action Object on page 45
- Enabling a Bandwidth Priority and Bandwidth Objects in an Action Object on page 46

Configuring a Bandwidth Object

To configure a bandwidth object:

1 Navigate to **MANAGE | Policies > Objects > Bandwidth Objects**.

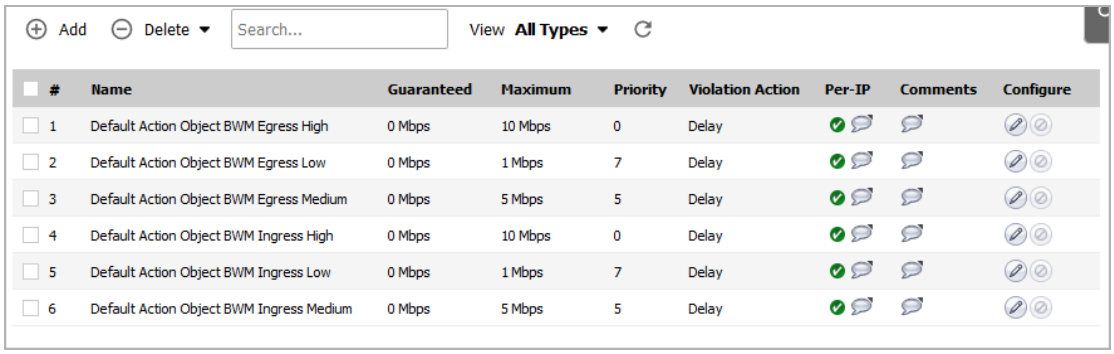

#	Name	Guaranteed	Maximum	Priority	Violation Action	Per-IP	Comments	Configure
1	Default Action Object BWM Egress High	0 Mbps	10 Mbps	0	Delay			
2	Default Action Object BWM Egress Low	0 Mbps	1 Mbps	7	Delay			
3	Default Action Object BWM Egress Medium	0 Mbps	5 Mbps	5	Delay			
4	Default Action Object BWM Ingress High	0 Mbps	10 Mbps	0	Delay			
5	Default Action Object BWM Ingress Low	0 Mbps	1 Mbps	7	Delay			
6	Default Action Object BWM Ingress Medium	0 Mbps	5 Mbps	5	Delay			

2 Do one of the following:

- Click the **Add** icon to create a new Bandwidth Object.
- Click the **Edit** icon of the Bandwidth Object you want to change.

The **Add/Edit Bandwidth Object** dialog displays.

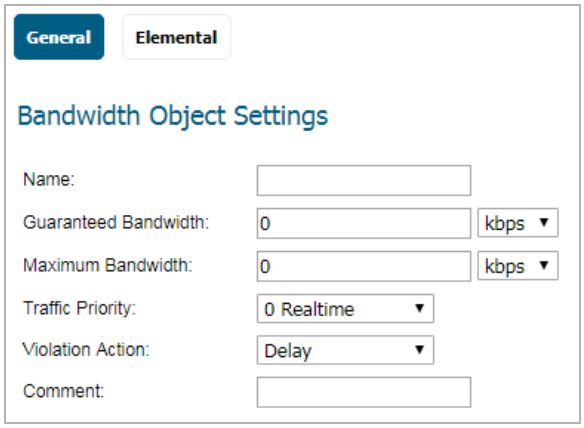

3 In the **Name** field, enter a name for this bandwidth object.

4 In the **Guaranteed Bandwidth** field, enter the amount of bandwidth that this bandwidth object will guarantee to provide for a traffic class (in kbps or Mbps).

a Specify whether the bandwidth is **kbps** (default) or **Mbps** from the drop-down menu.

5 In the **Maximum Bandwidth** field, enter the maximum amount of bandwidth that this bandwidth object will provide for a traffic class.

NOTE: The actual allocated bandwidth may be less than this value when multiple traffic classes compete for a shared bandwidth.

a Specify whether the bandwidth is **kbps** (default) or **Mbps** from the drop-down menu.

6 In the **Traffic Priority** field, enter the priority that this bandwidth object will provide for a traffic class. The highest, and default, priority is **0 Realtime**. The lowest priority is **7 Lowest**.

When multiple traffic classes compete for shared bandwidth, classes with the highest priority are given precedence.

7 In the **Violation Action** field, enter the action that this bandwidth object will provide when traffic exceeds the maximum bandwidth setting:

- **Delay** – Specifies that excess traffic packets are queued and sent when possible.
- **Drop** – Specifies that excess traffic packets are dropped immediately.

8 In the **Comment** field, enter a text comment or description for this bandwidth object.

9 Click **OK**.

Enabling Elemental Bandwidth Management

Elemental Bandwidth Management enables SonicOS to enforce bandwidth rules and policies on each individual IP that passes through the firewall.

To enable elemental bandwidth management in a bandwidth object:

1 Navigate to **MANAGE | Policies > Objects > Bandwidth Objects**.

2 Click the **Edit** icon of the Bandwidth Object you want to change. The **Bandwidth Object Settings** dialog displays.

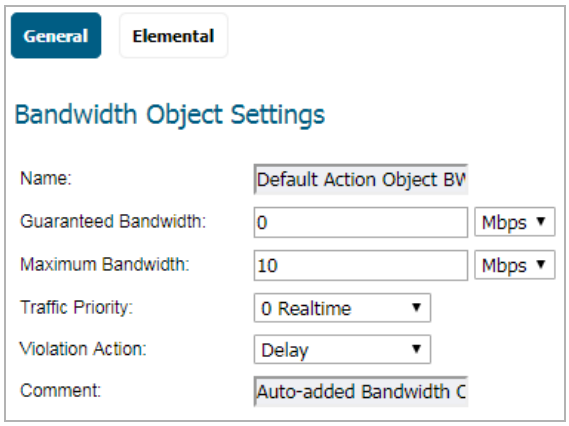

3 Click **Elemental**.



4 Select the **Enable Per-IP Bandwidth Management** option. This option is not selected by default. When enabled, the maximum elemental bandwidth setting applies to each individual IP under the parent traffic class.

5 In the **Maximum Bandwidth** field, enter the maximum elemental bandwidth that can be allocated to a protocol under the parent traffic class.

a Specify whether the bandwidth is **kbps** (default) or **Mbps** from the drop-down menu.

6 Click **OK**.

Enabling a Bandwidth Object in an Access Rule

If Advanced BWM is selected, you can enable bandwidth objects (and their configurations) in **MANAGE | Policies > Rules > Access Rules**.

To enable a bandwidth object in an Access Rule:

1. Navigate to **MANAGE | Policies > Rules** > **Access Rules**.
2. Do one of the following:
   - Click the **Add** icon to create a new Access Rule. The **Add Rule** dialog displays.
   - Click the **Edit** icon for the appropriate Access Rule. The **Edit Rule** dialog displays.
3. Click **BWM**.

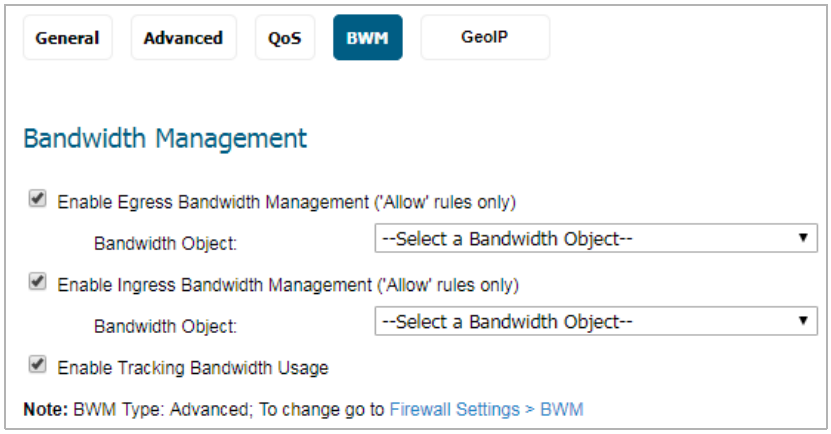

4. To enable a bandwidth object for the egress direction, under **Bandwidth Management**, select **Enable Egress Bandwidth Management ('allow' rules only)**.
5. From the **Bandwidth Object** drop-down menu, select the bandwidth object you want for the egress direction.
6. To enable a bandwidth object for the ingress direction, under **Bandwidth Management**, select **Enable Ingress Bandwidth Management ('allow' rules only)**.
7. From the **Bandwidth Object** drop-down menu, select the bandwidth object you want for the ingress direction.
8. To enable bandwidth usage tracking, select the **Enable Tracking Bandwidth Usage** option.
9. Click **OK**.

Enabling a Bandwidth Priority in an Access Rule

If **Global BWM BWM** is selected, you can enable bandwidth priority in **MANAGE | Policies > Rules > Access Rules**.

To enable bandwidth priority in an Access Rule:

1. Navigate to **MANAGE | Policies > Rules > Access Rules**.
2. Do one of the following:
   - Click the **Add** icon to create a new Access Rule. The **Add Rule** dialog displays.
   - Click the **Edit** icon for the appropriate Access Rule. The **Edit Rule** dialog displays.

3 Click **BWM**.

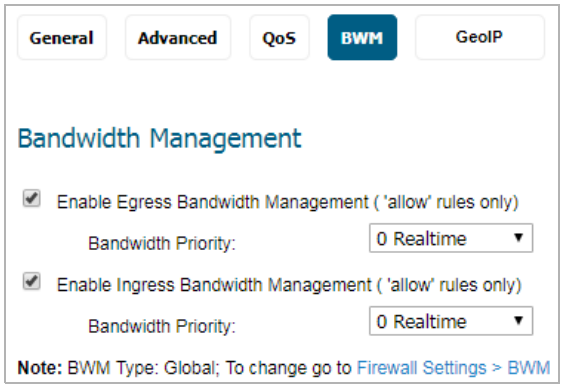

4 To enable a bandwidth object for the egress direction, under **Bandwidth Management**, select **Enable Egress Bandwidth Management ('allow' rules only)**. This option is not selected by default.

5 From the **Bandwidth Priority** drop-down menu, select the bandwidth priority you want for the egress direction. The highest, and default, priority is **0 Realtime**. The lowest priority is **7 Lowest**.

6 To enable a bandwidth object for the ingress direction, under **Bandwidth Management**, select **Enable Ingress Bandwidth Management ('allow' rules only)**. This option is not selected by default.

7 From the **Bandwidth Priority** drop-down menu, select the bandwidth priority you want for the ingress direction. The highest, and default, priority is **0 Realtime**. The lowest priority is **7 Lowest**.

8 Click **OK**.

Enabling a Bandwidth Object in an Action Object

If **Advanced BWM** is selected, you can enable bandwidth objects (and their configurations) in **Rules > Access Rules**.

To enable a bandwidth object in an action object:

1 Navigate to **MANAGE | Policies > Objects > Action Objects**.

2 Create a new action object by clicking on the **Add** icon. The **Action Object Settings** dialog displays.

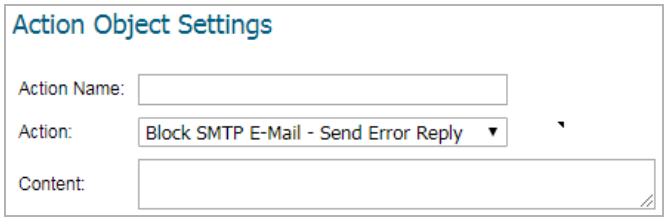

3 Enter a name for the action object in the **Action Name** field.

4 From **Action**, select **Bandwidth Management**, which allows control and monitoring of application-level bandwidth usage. The options on the **Action Object Settings** dialog change.

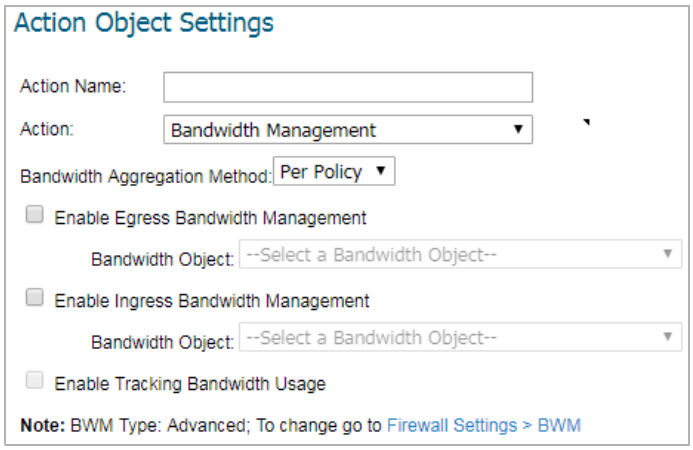

5 In the **Bandwidth Aggregation Method** drop-down menu, select the appropriate bandwidth aggregation method:

- **Per Policy** (default)
- **Per Action**

6 To enable bandwidth management in the egress direction, select the **Enable Egress Bandwidth Management** option.

a From **Bandwidth Object**, select the bandwidth object or create a new bandwidth object for the egress direction.

7 To enable bandwidth management in the ingress direction, select the **Enable Ingress Bandwidth Management** option.

a From **Bandwidth Object**, select the bandwidth object or create a new bandwidth object for the egress direction.

8 Click **OK**.

Enabling a Bandwidth Priority and Bandwidth Objects in an Action Object

If **Global BWM** is selected, you can specify BWM priority and enable bandwidth objects (and their configurations) in **MANAGE | Policies > Rules > Access Rules**.

To enable bandwidth priority and a bandwidth object in an action object:

1 Navigate to **MANAGE | Policies > Objects > Action Objects**.

2 Create a new action object by clicking on the **Add** icon. The **Action Object Settings** dialog displays.

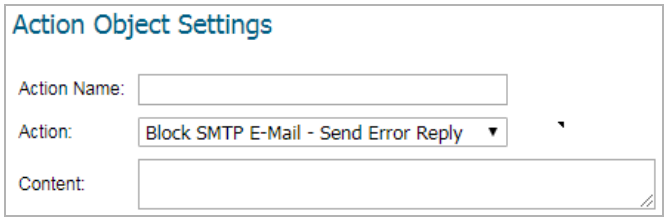

3 Enter a name for the action object in the **Action Name** field.

4 From **Action**, select **Bandwidth Management**, which allows control and monitoring of application-level bandwidth usage. The options on the **Action Object Settings** dialog change.

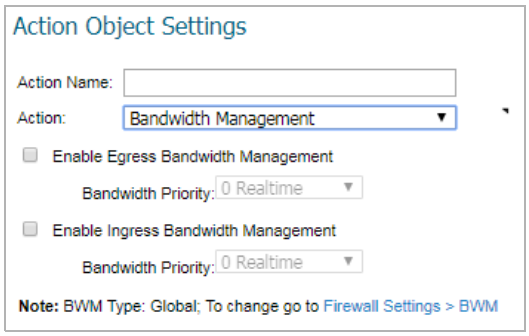

5 To enable bandwidth management in the egress direction, select the **Enable Egress Bandwidth Management** for priority option.

   a From the **Bandwidth Priority** drop-down menu, select the bandwidth object for the egress direction. The highest, and default, priority is **0 Realtime**. The lowest priority is **7 Lowest**.

6 To enable bandwidth management in the ingress direction, select the **Enable Ingress Bandwidth Management** for priority option.

   a From the **Bandwidth Priority** drop-down menu, select the bandwidth object for the ingress direction. The highest, and default, priority is **0 Realtime**. The lowest priority is **7 Lowest**.

7 Click **OK**.

Setting Interface Bandwidth Limitations with Advanced BWM

To set the bandwidth limitations for an interface:

1 Navigate to **MANAGE | System Setup > Network > Interfaces**.

2 Click the **Edit** icon for the appropriate interface. The **Edit Interface** dialog displays.

3 Click **Advanced**.

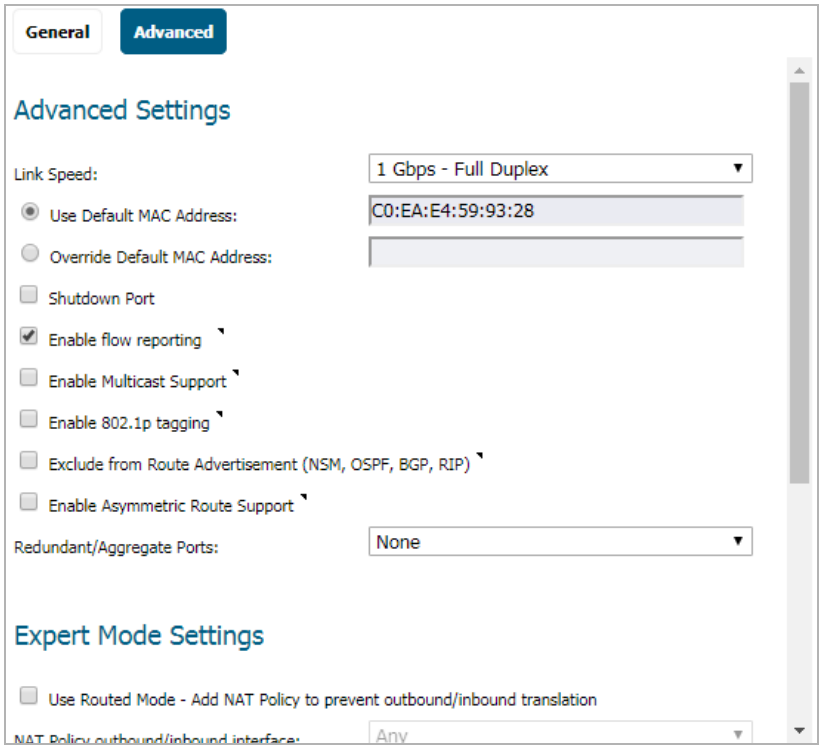

4 Scroll to the **Bandwidth Management** section.

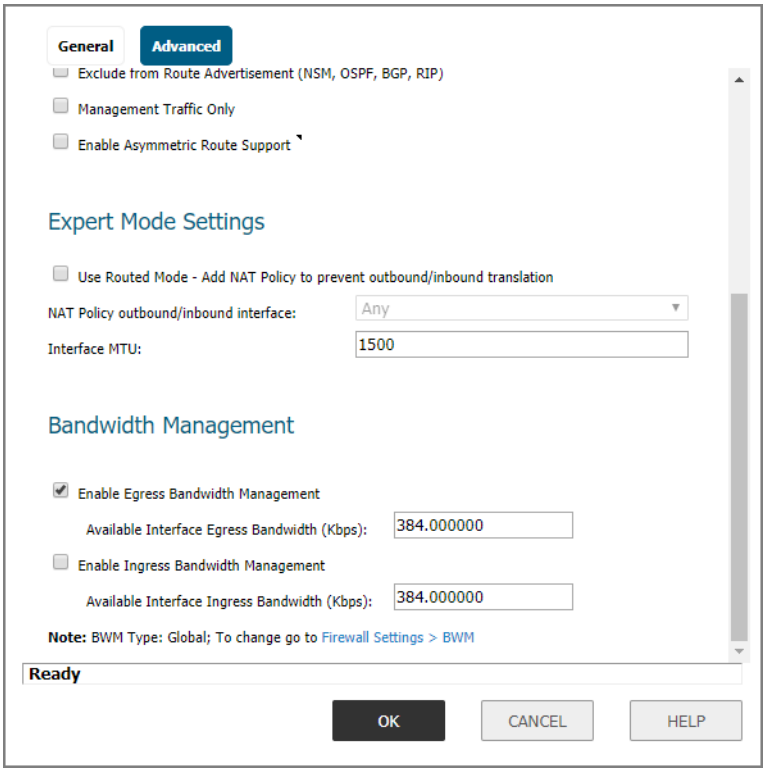

5 Select the **Enable Egress Bandwidth Management**. This option is not selected by default.

When this option is:

- Selected, the maximum available egress BWM is defined, but since 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 egress traffic can still be shaped using other options.

a In the **Available Interface Egress Bandwidth (kbps)** field, enter the maximum egress bandwidth for the interface (in kilobytes per second). The default is **384.000000** Kbps.

6 Select the **Enable Ingress Bandwidth Management**. This option is not selected by default. For information on using this option, see Step 5.

7 Click **OK**.

Setting Interface Bandwidth Limitations with Global BWM

To set the bandwidth limitations for an interface:

1 Navigate to **MANAGE | System Setup > Network > Interfaces**.

2 Click the **Edit** icon for the appropriate interface. The **Edit Interface** dialog displays.

3 Click **Advanced**.

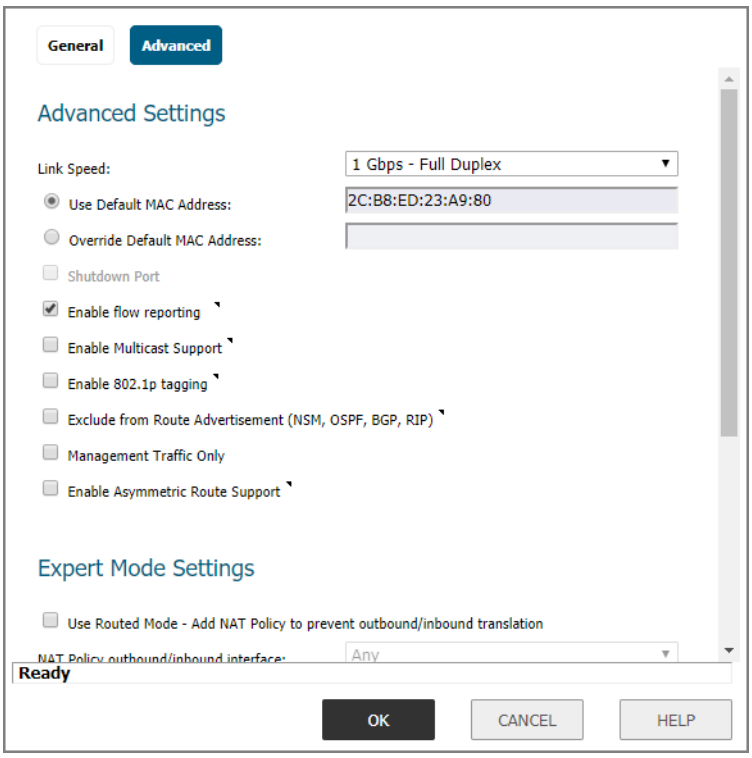

4 Scroll to the **Bandwidth Management** section.

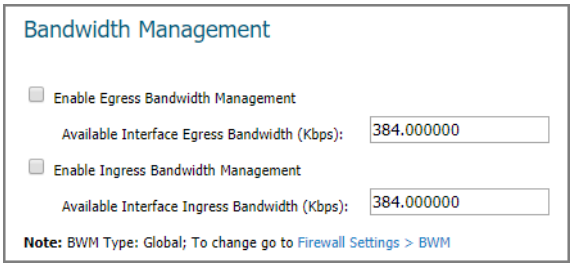

5 Select the **Enable Egress Bandwidth Management** option. This option is not selected by default.

When this option 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 egress traffic can still be shaped using other options.

a In the **Available Interface Egress Bandwidth (kbps)** field, enter the maximum egress bandwidth for the interface (in kilobytes per second). The default is **384.000000** Kbps.

6 Select the **Enable Ingress Bandwidth Management** option. This option is not selected by default. This option is not selected by default. For information on using this option, see Step 5.

a In the **Available Interface Ingress Bandwidth (kbps)** field, enter the maximum egress bandwidth for the interface (in kilobytes per second). The default is **384.000000** Kbps.

7 Click **OK**.

3

Configuring Flood Protection

NOTE: Control Plane flood protection is located on the **MANAGE | Security Configuration > Firewall Settings > Advanced Settings** page.

- Firewall Settings > Flood Protection on page 52

Firewall Settings > Flood Protection

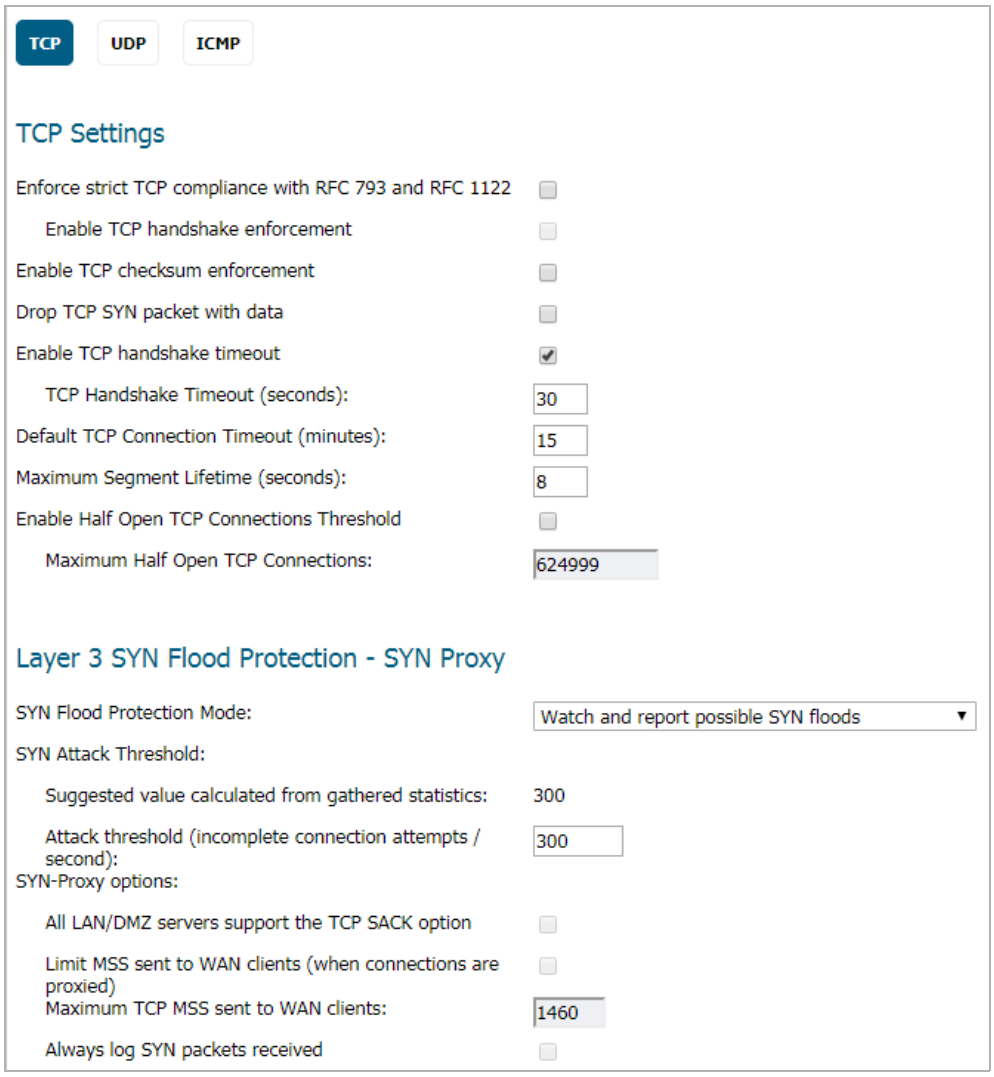

TIP: You must click **ACCEPT** to activate any settings you select.

The **Firewall Settings > Flood Protection** page lets you:

- Manage:
  - TCP (Transmission Control Protocol) traffic settings such as Layer 2/Layer3 flood protection, WAN DDOS protection
  - UDP (User Datagram Protocol) flood protection
  - ICMP (Internet Control Message Protocol) or ICMPv6 flood protection.
- View statistics on traffic through the security appliance:
  - TCP traffic
  - UDP traffic
  - ICMP or ICMPv6 traffic

SonicOS defends against UDP/ICMP flood attacks by monitoring IPv6 UDP/ICMP traffic flows to defined destinations. UDP/ICMP packets to a specified destination are dropped if one or more sources exceeds a configured threshold.

Topics:

- TCP View on page 53
- UDP View on page 63
- ICMP View on page 66

TCP View

Topics:

- TCP Settings on page 53
- Layer 3 SYN Flood Protection - SYN Proxy View on page 54
- Configuring Layer 3 SYN Flood Protection on page 56
- Configuring Layer 2 SYN/RST/FIN/TCP Flood Protection – MAC Blacklisting on page 58
- WAN DDOS Protection (Non-TCP Floods) on page 58

TCP Settings

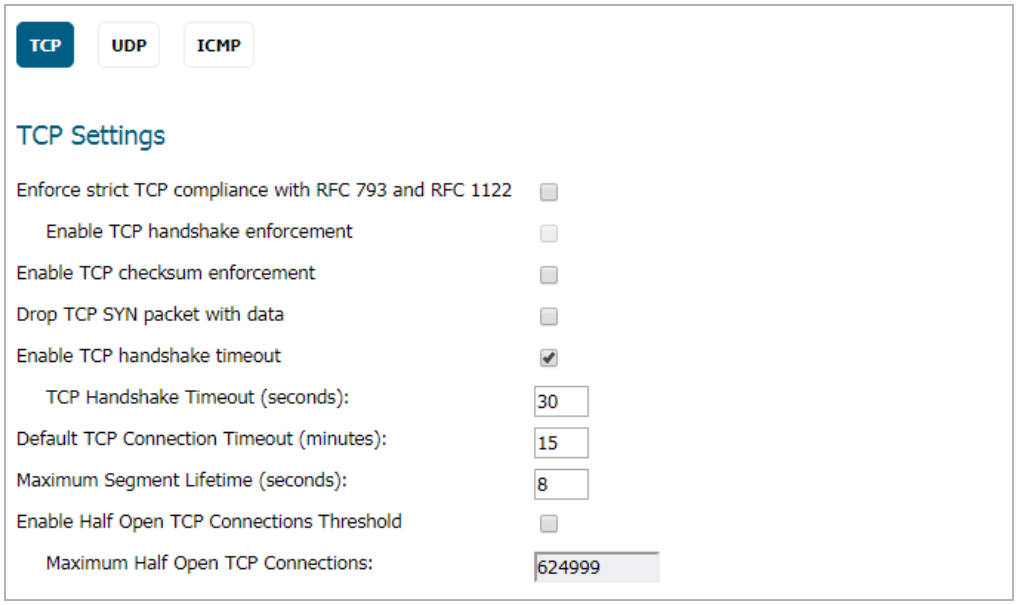

- **Enforce strict TCP compliance with RFC 793 and RFC 1122** – Ensures strict compliance with several TCP timeout rules. This setting maximizes TCP security, but it may cause problems with the Window Scaling feature for Windows Vista users. This option is not selected by default.
    - **Enable TCP handshake enforcement** – Requires a successful three-way TCP handshake for all TCP connections. This option is available only if **Enforce strict TCP compliance with RFC 793 and RFC 1122** is selected.
- **Enable TCP checksum enforcement** – If an invalid TCP checksum is calculated, the packet is dropped. This option is not selected by default.

- **Enable TCP handshake timeout** – Enforces the timeout period (in seconds) for a three-way TCP handshake to complete its connection. If the three-way TCP handshake does not complete in the timeout period, it is dropped. This option is selected by default.
    - **TCP Handshake Timeout (seconds)**: The maximum time a TCP handshake has to complete the connection. The default is **30** seconds.
- **Default TCP Connection Timeout** – The default time assigned to Access Rules for TCP traffic. If a TCP session is active for a period in excess of this setting, the TCP connection is cleared by the firewall. The default value is **15** minutes, the minimum value is 1 minute, and the maximum value is 999 minutes.

    NOTE: Setting excessively long connection time-outs slows the reclamation of stale resources, and in extreme cases, could lead to exhaustion of the connection cache.

- **Maximum Segment Lifetime (seconds)** – Determines the number of seconds that any TCP packet is valid before it expires. This setting is also used to determine the amount of time (calculated as twice the Maximum Segment Lifetime, or 2MSL) that an actively closed TCP connection remains in the TIME_WAIT state to ensure that the proper FIN / ACK exchange has occurred to cleanly close the TCP connection. The default value is **8** seconds, the minimum value is 1 second, and the maximum value is 60 seconds.
- **Enable Half Open TCP Connections Threshold** – Denies new TCP connections if the high-water mark of TCP half-open connections has been reached. By default, the half-open TCP connection is not monitored, so this option is not selected by default.
    - **Maximum Half Open TCP Connections** – Specifies the maximum number of half-open TCP connections and is available only if **Enable Half Ope TCP Connections Threshold** is selected. The default maximum is half the number of maximum connection caches.

Layer 3 SYN Flood Protection - SYN Proxy View

Topics:

- SYN Flood Protection Methods on page 54
- Configuring Layer 3 SYN Flood Protection on page 56

SYN Flood Protection Methods

SYN/RST/FIN flood protection helps to protect hosts behind the firewall from Denial of Service (DoS) or Distributed DoS attacks that attempt to consume the host's available resources by creating one of the following attack mechanisms:

- Sending TCP SYN packets, RST packets, or FIN packets with invalid or spoofed IP addresses.
- Creating excessive numbers of half-opened TCP connections.

The following sections detail some SYN flood protection methods:

- SYN Flood Protection Using Stateless Cookies on page 54
- Layer-Specific SYN Flood Protection Methods on page 55
- Understanding SYN Watchlists on page 55
- Understanding a TCP Handshake on page 55

SYN Flood Protection Using Stateless Cookies

The method of SYN flood protection employed starting with SonicOS uses stateless SYN Cookies, which increase reliability of SYN Flood detection, and also improves overall resource utilization on the firewall. With stateless

SYN Cookies, the firewall does not have to maintain state on half-opened connections. Instead, it uses a cryptographic calculation (rather than randomness) to arrive at SEQr.

Layer-Specific SYN Flood Protection Methods

SonicOS provides several protections against SYN Floods generated from two different environments: trusted (internal) or untrusted (external) networks. Attacks from *untrusted* WAN networks usually occur on one or more servers protected by the firewall. Attacks from the *trusted* LAN networks occur as a result of a virus infection inside one or more of the trusted networks, generating attacks on one or more local or remote hosts.

To provide a firewall defense to both attack scenarios, SonicOS provides two separate SYN Flood protection mechanisms on two different layers. Each gathers and displays SYN Flood statistics and generates log messages for significant SYN Flood events.

- **SYN Proxy (Layer 3)** – This mechanism shields servers inside the trusted network from WAN-based SYN flood attacks, using a SYN Proxy implementation to verify the WAN clients before forwarding their connection requests to the protected server. You can enable SYN Proxy only on WAN interfaces.
- **SYN Blacklisting (Layer 2)** – This mechanism blocks specific devices from generating or forwarding SYN flood attacks. You can enable SYN Blacklisting on any interface.

Understanding SYN Watchlists

The internal architecture of both SYN Flood protection mechanisms is based on a single list of Ethernet addresses that are the most active devices sending initial SYN packets to the firewall. This list is called a *SYN watchlist*. Because this list contains Ethernet addresses, the device tracks all SYN traffic based on the address of the device forwarding the SYN packet, without considering the IP source or destination address.

Each watchlist entry contains a value called a *hit count*. The hit count value increments when the device receives the an initial SYN packet from a corresponding device. The hit count decrements when the TCP three-way handshake completes. The hit count for any particular device generally equals the number of half-open connections pending since the last time the device reset the hit count. The device default for resetting a hit count is once a second.

The thresholds for logging, SYN Proxy, and SYN Blacklisting are all compared to the hit count values when determining if a log message or state change is necessary. When a SYN Flood attack occurs, the number of pending half-open connections from the device forwarding the attacking packets increases substantially because of the spoofed connection attempts. When you set the attack thresholds correctly, normal traffic flow produces few attack warnings, but the same thresholds detect and deflect attacks before they result in serious network degradation.

Understanding a TCP Handshake

A typical TCP handshake (simplified) begins with an initiator sending a TCP SYN packet with a 32-bit sequence (SEQi) number. The responder then sends a SYN/ACK packet acknowledging the received sequence by sending an ACK equal to SEQi+1 and a random, 32-bit sequence number (SEQr). The responder also maintains state awaiting an ACK from the initiator. The initiator's ACK packet should contain the next sequence (SEQi+1) along with an acknowledgment of the sequence it received from the responder (by sending an ACK equal to SEQr+1). The exchange looks as follows:

1 Initiator -> SYN (SEQi=0001234567, ACKi=0) -> Responder

2 Initiator <- SYN/ACK (SEQr=3987654321, ACKr=0001234568) <- Responder

3 Initiator -> ACK (SEQi=0001234568, ACKi=3987654322) -> Responder

Because the responder has to maintain state on all half-opened TCP connections, it is possible for memory depletion to occur if SYNs come in faster than they can be processed or cleared by the responder. A half-opened TCP connection did not transition to an established state through the completion of the three-way handshake. When the firewall is between the initiator and the responder, it effectively becomes the responder, brokering, or *proxying*, the TCP connection to the actual responder (private host) it is protecting.

Configuring Layer 3 SYN Flood Protection

A SYN Flood Protection mode is the level of protection that you can select to defend against half-opened TCP sessions and high-frequency SYN packet transmissions. This feature enables you to set three different levels of SYN Flood Protection.

To configure SYN Flood Protection features:

1 Go to the **Layer 3 SYN Flood Protection - SYN Proxy** section of the **MANAGE | Security Configuration > Firewall Settings > Flood Protection** page.

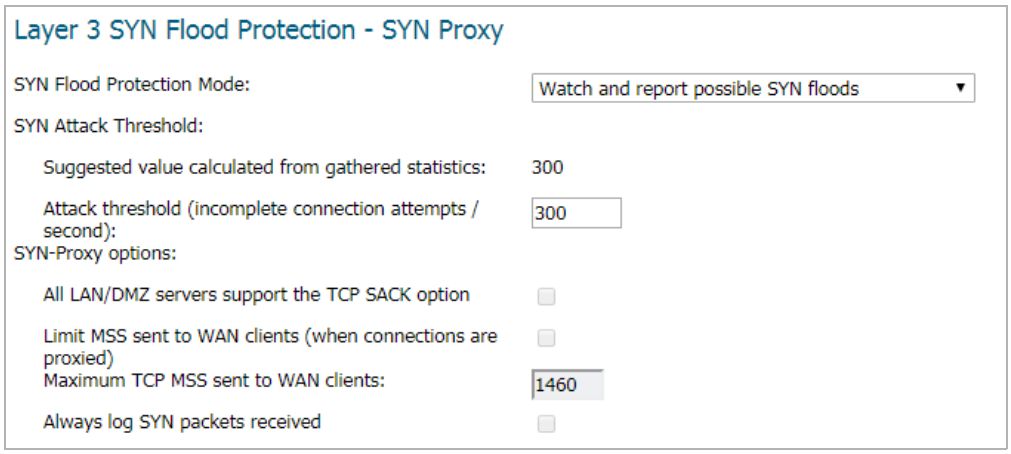

2 From the **SYN Flood Protection Mode** drop-down menu, select the type of protection mode:

- **Watch and Report Possible SYN Floods** – Enables the device to monitor SYN traffic on all interfaces on the device and to log suspected SYN flood activity that exceeds a packet count threshold. The feature does not turn on the SYN Proxy on the device, so the device forwards the TCP three-way handshake without modification.

  This is the least invasive level of SYN Flood protection. Select this option if your network is not in a high-risk environment.

  IMPORTANT: When this protection mode is selected, the **SYN-Proxy options** are not available.

- **Proxy WAN Client Connections When Attack is Suspected** – Enables the device to enable the SYN Proxy feature on WAN interfaces when the number of incomplete connection attempts per second surpasses a specified threshold. This method ensures the device continues to process valid traffic during the attack and that performance does not degrade. Proxy mode remains enabled until all WAN SYN flood attacks stop occurring or until the device blacklists all of them using the SYN Blacklisting feature.

  This is the intermediate level of SYN Flood protection. Select this option if your network experiences SYN Flood attacks from internal or external sources.

- **Always Proxy WAN Client Connections** – Sets the device to always use SYN Proxy. This method blocks all spoofed SYN packets from passing through the device.

  This is an extreme security measure that directs the device to respond to port scans on all TCP ports because the SYN Proxy feature forces the device to respond to all TCP SYN connection attempts. This can degrade performance and can generate a false positive. Select this option only if your network is in a high-risk environment.

3 Select the **SYN Attack Threshold** configuration options to provide limits for SYN Flood activity before the device drops packets. The device gathers statistics on WAN TCP connections, keeping track of the maximum and average maximum and incomplete WAN connections per second. Out of these statistics, the device suggests a value for the SYN flood threshold.

- **Suggested value calculated from gathered statistics** – The suggested attack threshold based on WAN TCP connection statistics. This value cannot be changed.
- **Attack Threshold (Incomplete Connection Attempts/Second)** – Enables you to set the threshold for the number of incomplete connection attempts per second before the device drops packets at any value between 5 and 200,000. The default is the **Suggested value calculated from gathered statistics**.

4 Select the **SYN-Proxy options** to provide more control over the options sent to WAN clients when in SYN Proxy mode.

IMPORTANT: The options in this section are not available if **Watch and report possible SYN floods** is selected for **SYN Flood Protection** Mode.

When the device applies a SYN Proxy to a TCP connection, it responds to the initial SYN packet with a manufactured SYN/ACK reply, waiting for the ACK in response before forwarding the connection request to the server. Devices attacking with SYN Flood packets do not respond to the SYN/ACK reply. The firewall identifies them by their lack of this type of response and blocks their spoofed connection attempts. SYN Proxy forces the firewall to manufacture a SYN/ACK response without knowing how the server will respond to the TCP options normally provided on SYN/ACK packets.

- **All LAN/DMZ servers support the TCP SACK option** – Enables SACK (Selective Acknowledgment) where a packet can be dropped and the receiving device indicates which packets it received. This option is not enabled by default. Enable this checkbox only when you know that all servers covered by the firewall accessed from the WAN support the SACK option.
- **Limit MSS sent to WAN clients (when connections are proxied)** – Enables you to enter the maximum MSS (Minimum Segment Size) value. This sets the threshold for the size of TCP segments, preventing a segment that is too large to be sent to the targeted server. For example, if the server is an IPsec gateway, it may need to limit the MSS it received to provide space for IPsec headers when tunneling traffic. The firewall cannot predict the MSS value sent to the server when it responds to the SYN manufactured packet during the proxy sequence. Being able to control the size of a segment, enables you to control the manufactured MSS value sent to WAN clients. This option is not selected by default.

  If you specify an override value for the default of **1460**, a segment of that size or smaller is sent to the client in the SYN/ACK cookie. Setting this value too low can decrease performance when the SYN Proxy is always enabled. Setting this value too high can break connections if the server responds with a smaller MSS value.

  - **Maximum TCP MSS sent to WAN clients.** The value of the MSS. The default is **1460**, the minimum value is 32, and the maximum is 1460.

  NOTE: When using Proxy WAN client connections, remember to set these options conservatively as they only affect connections when a SYN Flood takes place. This ensures that legitimate connections can proceed during an attack.

- **Always log SYN packets received.** Logs all SYN packets received.

Layer 2 SYN/RST/FIN Flood Protection - MAC Blacklisting

The SYN/RST/FIN Blacklisting feature lists devices that exceeded the SYN, RST, and FIN Blacklist attack threshold. The firewall device drops packets sent from blacklisted devices early in the packet evaluation process, enabling the firewall to handle greater amounts of these packets, providing a defense against attacks originating on local networks while also providing second-tier protection for WAN networks.

Devices cannot occur on the SYN/RST/FIN Blacklist and watchlist simultaneously. With blacklisting enabled, the firewall removes devices exceeding the blacklist threshold from the watchlist and places them on the blacklist. Conversely, when the firewall removes a device from the blacklist, it places it back on the watchlist. Any device

whose MAC address has been placed on the blacklist will be removed from it approximately three seconds after the flood emanating from that device has ended.

Configuring Layer 2 SYN/RST/FIN/TCP Flood Protection – MAC Blacklisting

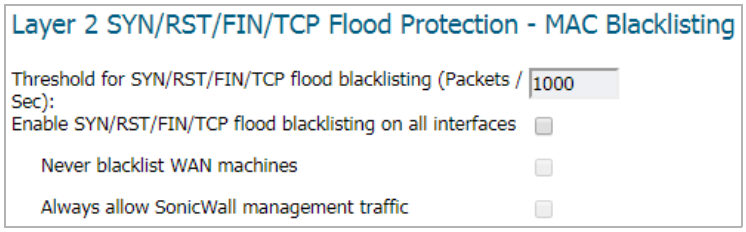

- **Threshold for SYN/RST/FIN flood blacklisting (SYNs / Sec)** – Specifies he maximum number of SYN, RST, FIN, and TCP packets allowed per second. The minimum is 10, the maximum is 800000, and default is **1,000**. This value should be larger than the SYN Proxy threshold value because blacklisting attempts to thwart more vigorous local attacks or severe attacks from a WAN network.

  NOTE: This option cannot be modified unless **Enable SYN/RST/FIN/TCP flood blacklisting** on all interfaces is enabled.

- **Enable SYN/RST/FIN/TCP flood blacklisting on all interfaces** – Enables the blacklisting feature on all interfaces on the firewall. This option is not selected by default. When it is selected, these options become available:

  - **Never blacklist WAN machines** – Ensures that systems on the WAN are never added to the SYN Blacklist. This option is recommended as leaving it cleared may interrupt traffic to and from the firewall's WAN ports. This option is not selected by default.
  - **Always allow SonicWall management traffic** – Causes IP traffic from a blacklisted device targeting the firewall's WAN IP addresses to not be filtered. This allows management traffic and routing protocols to maintain connectivity through a blacklisted device. This option is not selected by default.

WAN DDOS Protection (Non-TCP Floods)

WAN DDOS Protection provides protection against non-TCP DDOS attacks and so should be used in combination with SYN-Flood Protection if TCP SYN-flood attacks are a concern. This feature is not intended to protect a well-known server of non-TCP services on the internet (such as a central DNS server), but is intended to protect LAN and DMZ networks for which the majority of non-TCP traffic is initiated from the LAN/DMZ side, possibly in combination with limited WAN-initiated traffic.

When **WAN DDOS Protection** is enabled, it tracks the rate of non-TCP packets arriving on WAN interfaces. When the rate of non-TCP packets exceeds the specified threshold, non-TCP packets arriving on WAN interfaces will be filtered. A non-TCP packet will only be forwarded when at least one of the following conditions is true:

- Source IP address is on the Allow list
- Packet is SonicWall management traffic, and **Always allow SonicWall management traffic** is selected
- Packet is VPN Negotiation traffic (IKE) and **Always allow VPN negotiation traffic** is selected
- the packet is an ESP packet and matches the SPI of a tunnel terminating on the network security appliance
- the packet is the *n*th packet matching the value specified for **WAN DDOS Filter Bypass Rate (every n packets)**

If none of these conditions are met, the packet is dropped early in packet processing.

You can configure the **WAN DDOS Protection (Non-TCP Floods)** settings on the **MANAGE | Security Configuration > Firewall Settings > Flood Protection** page.

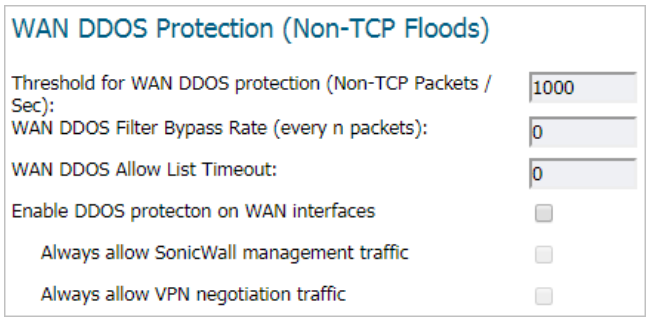

Topics:

- Threshold for WAN DDOS protection (Non-TCP Packets / Sec) on page 59
- WAN DDOS Filter Bypass Rate (every n packets) on page 59
- WAN DDOS Allow List Timeout on page 59
- Enable DDOS protection on WAN interfaces on page 59

Threshold for WAN DDOS protection (Non-TCP Packets / Sec)

Threshold for WAN DDOS protection specifies the maximum number of non-TCP packets allowed per second to be sent to a host, range, or subnet. Exceeding this threshold triggers WAN DDOS flood protection. The default number of non-TCP packets is 1000. The minimum number is 0; the maximum number is 10000000.

WAN DDOS Filter Bypass Rate (every *n* packets)

When the configured Filter Bypass Rate is non-zero, a non-TCP packet that would normally be dropped by WAN DDOS Protection will instead be passed to the LAN/DMZ network. The bypass rate allows a potential attack to be throttled, but not completely blocked. Allowing some packets to pass through even though their sources are not on the Allow List can provide a mechanism by which legitimate WAN side hosts may get a packet through to the LAN/DMZ side, and a response would populate the Allow List so the following non-TCP packets from the legitimate WAN side host would always be forwarded from that point on.

The default value of the Filter Bypass Rate is zero, so the user must modify this value before the heuristic can be attempted. When the Filter Bypass Rate is non-zero, the value determines what proportion of packets are forwarded regardless of the Allow List contents. For example, if the value was set to two, every other packet would be forwarded to the LAN/DMZ networks (assuming they passed policy, etc). If the value were 100, every 100th packet would be forwarded, and so on. The appropriate value is dependent on the capabilities of the potential LAN side target machines and the nature of the legitimate non-TCP traffic patterns in the customer's network.

WAN DDOS Allow List Timeout

If a non-zero Allow List Timeout is defined by the user, entries in the Allow List will expire in the configured time. If the Allow List Timeout is zero, they never expire. In either case, the least-recently-used entry in a particular hash-bucket may be replaced by a new entry if no unused entry is available in the list.

Enable DDOS protection on WAN interfaces

Selecting **Enable DDOS protection on WAN interfaces** (it is disabled by default) allows you to set two additional options:

- Always allow SonicWall management traffic on page 60
- Always allow VPN negotiation traffic on page 60

Always allow SonicWall management traffic

When **Always allow SonicWall management traffic** is enabled (it is disabled by default), traffic needed to manage your SonicWall appliances is allowed to pass through your WAN gateways even when the appliance is under a non-TCP DDOS attack.

Always allow VPN negotiation traffic

When **Always allow VPN Negotiation traffic** is enabled (it is disabled by default), a VPN can be negotiated even when the appliance is under a non-TCP DDOS attack.

TCP Traffic Statistics

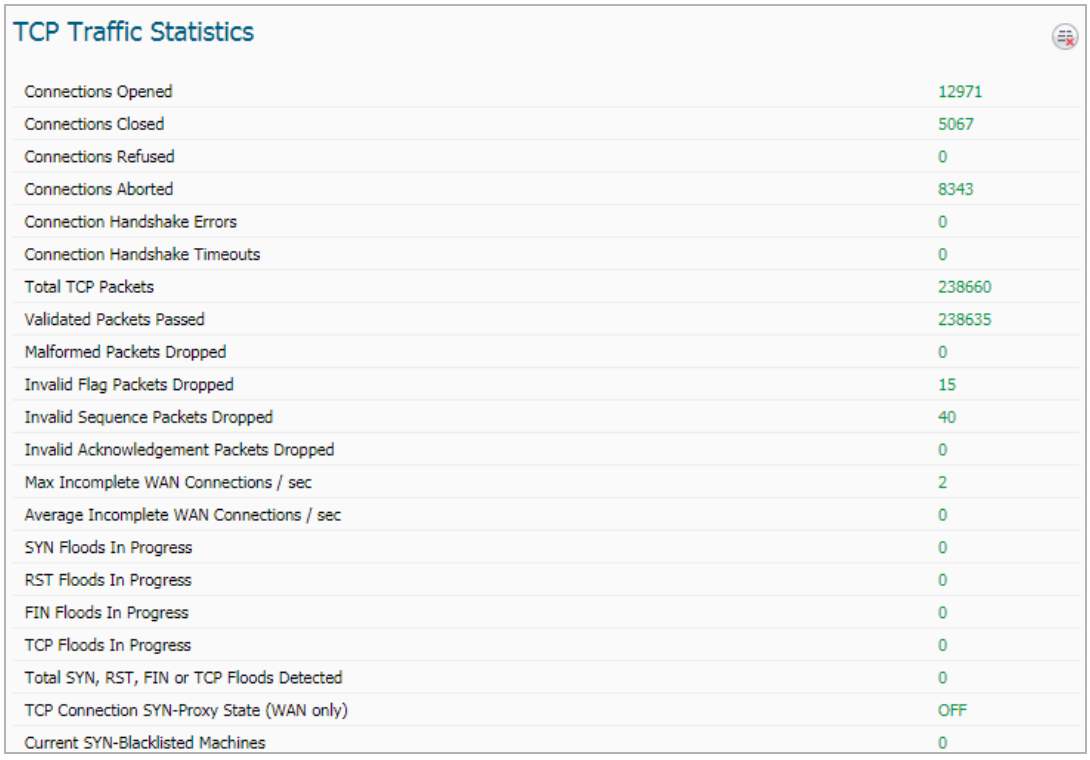

TCP Traffic Statistics describes the entries in the **TCP Traffic Statistics** table. To clear and restart the statistics displayed by a table, click the **Clear Stats** icon for the table.

TCP Traffic Statistics

This statistic	Is incremented/displays
Connections Opened	When a TCP connection initiator sends a SYN, or a TCP connection responder receives a SYN.
Connections Closed	When a TCP connection is closed when both the initiator and the responder have sent a FIN and received an ACK.
Connections Refused	When a RST is encountered, and the responder is in a SYN_RCVD state.
Connections Aborted	When a RST is encountered, and the responder is in some state other than SYN_RCVD.

TCP Traffic Statistics

This statistic	Is incremented/displays
Connection Handshake Error	When a handshake error is encountered.
Connection Handshake Timeouts	When a handshake times out.
Total TCP Packets	With every processed TCP packet.
Validated Packets Passed	When: • A TCP packet passes checksum validation (while TCP checksum validation is enabled). • A valid SYN packet is encountered (while SYN Flood protection is enabled). • A SYN Cookie is successfully validated on a packet with the ACK flag set (while SYN Flood protection is enabled).
Malformed Packets Dropped	When: • TCP checksum fails validation (while TCP checksum validation is enabled). • The TCP SACK Permitted option is encountered, but the calculated option length is incorrect. • The TCP MSS (Maximum Segment Size) option is encountered, but the calculated option length is incorrect. • The TCP SACK option data is calculated to be either less than the minimum of 6 bytes, or modulo incongruent to the block size of 4 bytes. • The TCP option length is determined to be invalid. • The TCP header length is calculated to be less than the minimum of 20 bytes. • The TCP header length is calculated to be greater than the packet's data length.

Invalid Flag Packets Dropped	When a:
	• Non-SYN packet is received that cannot be located in the connection-cache (while SYN Flood protection is disabled).
	• Packet with flags other than SYN, RST+ACK ,or SYN+ACK is received during session establishment (while SYN Flood protection is enabled).
	  • TCP XMAS Scan is logged if the packet has FIN, URG, and PSH flags set.
	  • TCP FIN Scan is logged if the packet has the FIN flag set.
	  • TCP Null Scan is logged if the packet has no flags set.
	• New TCP connection initiation is attempted with something other than just the SYN flag set.
	• Packet with the SYN flag set is received within an established TCP session.
	• Packet without the ACK flag set is received within an established TCP session.

Invalid Sequence Packets Dropped	When a: • Packet within an established connection is received where the sequence number is less than the connection's oldest unacknowledged sequence. • Packet within an established connection is received where the sequence number is greater than the connection's oldest unacknowledged sequence + the connection's last advertised window size.
Invalid Acknowledgement Packets Dropped	When an invalid acknowledgment packet is dropped.

TCP Traffic Statistics

This statistic	Is incremented/displays
Max Incomplete WAN Connections / sec	When a: • Packet is received with the ACK flag set, and with neither the RST or SYN flags set, but the SYN Cookie is determined to be invalid (while SYN Flood protection is enabled). • Packet's ACK value (adjusted by the sequence number randomization offset) is less than the connection's oldest unacknowledged sequence number. • Packet's ACK value (adjusted by the sequence number randomization offset) is greater than the connection's next expected sequence number.
Average Incomplete WAN Connections / sec	The average number of incomplete WAN connections per second.
SYN Floods In Progress	When a SYN flood is detected.
RST Floods In Progress	When a RST flood is detected.
FIN Floods In Progress	When a FIN flood is detected.
TCP Floods In Progress	When a TCP flood is detected.
Total SYN, RST, FIN or TCP Floods Detected	The total number of floods (SYN, RST, FIN, and TCP) detected.
TCP Connection SYN-Proxy State (WAN only)	For WAN only, whether the TCP connection SYN-proxy is enabled.
Current SYN-Blacklisted Machines	When a device is listed on the SYN blacklist.
Current RST-Blacklisted Machines	When a device is listed on the RST blacklist.
Current FIN-Blacklisted Machines	When a device is listed on the FIN blacklist.
Current TCP-Blacklisted Machines	When a device is listed on the TCP blacklist.
Total SYN-Blacklisting Events	When a SYN blacklisting event is detected.
Total RST-Blacklisting Events	When a RST blacklisting event is detected.
Total FIN-Blacklisting Events	When a FIN blacklisting event is detected.
Total TCP-Blacklisting Events	When a TCP blacklisting event is detected.
Total SYN Blacklist Packets Rejected	The total number of SYN packets rejected by SYN blacklisting.
Total RST Blacklist Packets Rejected	The total number of RST packets rejected by SYN blacklisting.
Total FIN Blacklist Packets Rejected	The total number of FIN packets rejected by SYN blacklisting.
Total TCP Blacklist Packets Rejected	The total number of TCP packets rejected by SYN blacklisting.
Invalid SYN Flood Cookies Received	When a SNY flood cookie is received.
WAN DDOS Filter State	Whether the DDOS filter is enabled or disabled.

TCP Traffic Statistics

This statistic	Is incremented/displays
WAN DDOS Filter – Packets Rejected	When a WAN DDOS Filter rejects a packet.
WAN DDOS Filter – Packets Leaked	When a WAN DDOS Filter rejects a leaked packet.
WAN DDOS Filter – Allow List Count	When a WAN DDOS Filter processes a packet in the Allow List.

UDP View

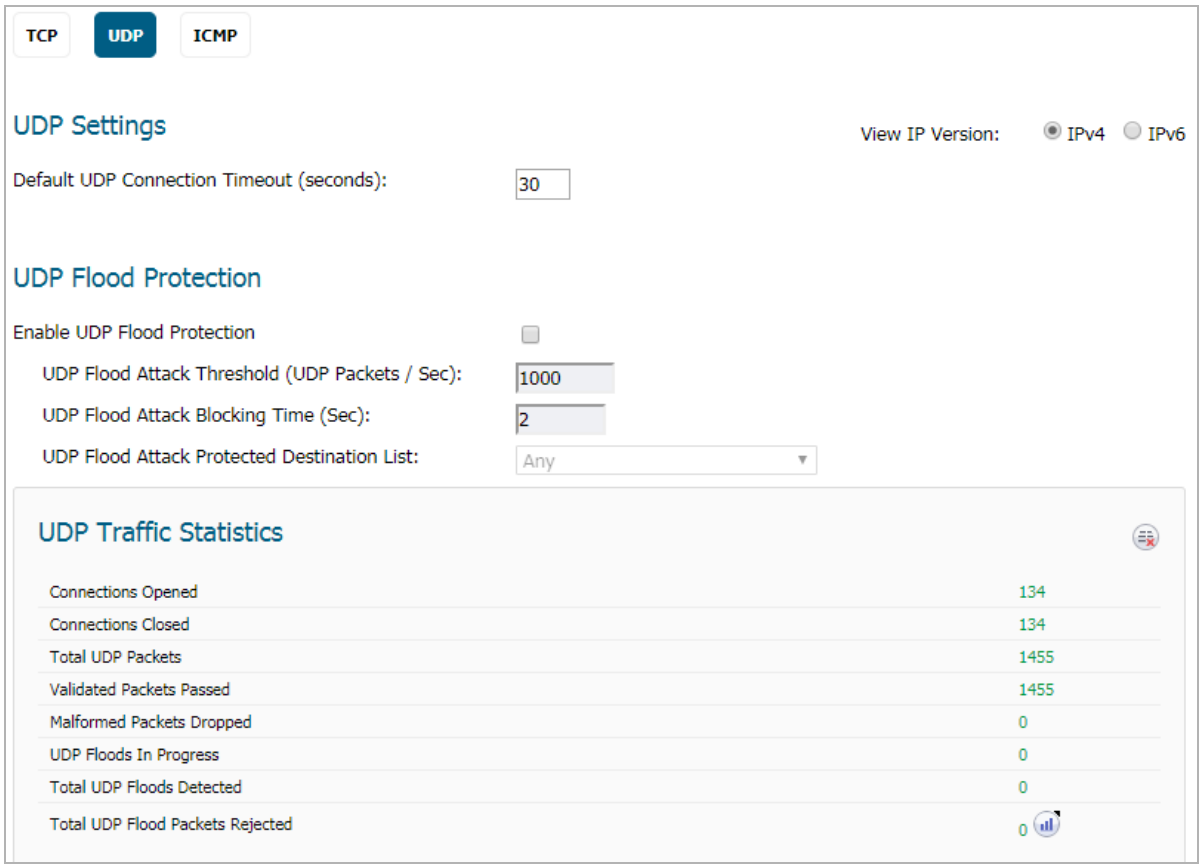

Topics:

- UDP Settings on page 63
- UDP Flood Protection on page 64
- UDP Traffic Statistics on page 65

UDP Settings

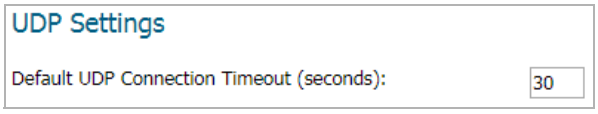

- **Default UDP Connection Timeout (seconds)** - The number of seconds of idle time you want to allow before UDP connections time out. This value is overridden by the UDP Connection timeout you set for individual rules.

UDP Flood Protection

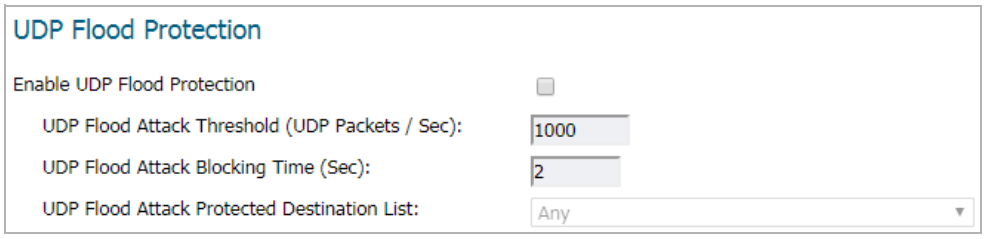

UDP Flood Attacks are a type of denial-of-service (DoS) attack. They are initiated by sending a large number of UDP packets to random ports on a remote host. As a result, the victimized system's resources are consumed with handling the attacking packets, which eventually causes the system to be unreachable by other clients.

SonicWall UDP Flood Protection defends against these attacks by using a "watch and block" method. The appliance monitors UDP traffic to a specified destination. If the rate of UDP packets per second exceeds the allowed threshold for a specified duration of time, the appliance drops subsequent UDP packets to protect against a flood attack.

UDP packets that are DNS query or responses to or from a DNS server configured by the appliance are allowed to pass, regardless of the state of UDP Flood Protection.

The following settings configure UDP Flood Protection:

- **Enable UDP Flood Protection** – Enables UDP Flood Protection. This option is not selected by default.

  NOTE: **Enable UDP Flood Protection** must be enabled to activate the other **UDP Flood Protection** options.

- **UDP Flood Attack Threshold (UDP Packets / Sec)** – The maximum number of UDP packets allowed per second to be sent to a host, range, or subnet that triggers UDP Flood Protection. Exceeding this threshold triggers ICMP Flood Protection.The minimum value is 50, the maximum value is 1000000, and the default value is **1000**.
- **UDP Flood Attack Blocking Time (Sec)** – After the appliance detects the rate of UDP packets exceeding the attack threshold for this duration of time, UDP Flood Protection is activated and the appliance begins dropping subsequent UDP packets. The minimum time is 1 second, the maximum time is 120 seconds, and the default time is **2** seconds.
- **UDP Flood Attack Protected Destination List** – The destination address object or address group that will be protected from UDP Flood Attack. The default value is **Any**.

  TIP: Select **Any** to apply the Attack Threshold to the sum of UDP packets passing through the firewall.

UDP Traffic Statistics



The **UDP Traffic Statistics** table provides statistics as shown in UDP Traffic Statistics. To clear and restart the statistics displayed by a table, click the **Clear Stats** icon for the table.

UDP Traffic Statistics

This statistic	Is incremented/displays
Connections Opened	When a connection is opened.
Connections Closed	When a connection is closed.
Total UDP Packets	With every processed UDP packet.
Validated Packets Passed	When a UDP packet passes checksum validation (while UDP checksum validation is enabled).
Malformed Packets Dropped	When: • UDP checksum fails validation (while UDP checksum validation is enabled). • The UDP header length is calculated to be greater than the packet's data length.
UDP Floods In Progress	The number of individual forwarding devices currently exceeding the UDP Flood Attack Threshold.
Total UDP Floods Detected	The total number of events in which a forwarding device has exceeded the UDP Flood Attack Threshold.
Total UDP Flood Packets Rejected	The total number of packets dropped because of UDP Flood Attack detection. Clicking on the **Statistics** icon displays a pop-up dialog showing the most recent rejected packets: 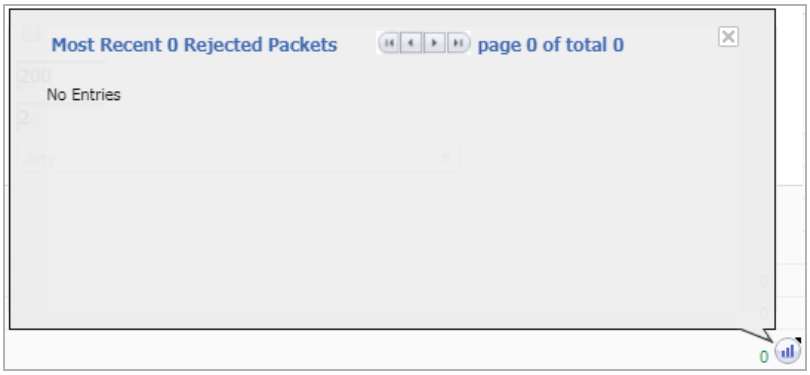

ICMP View

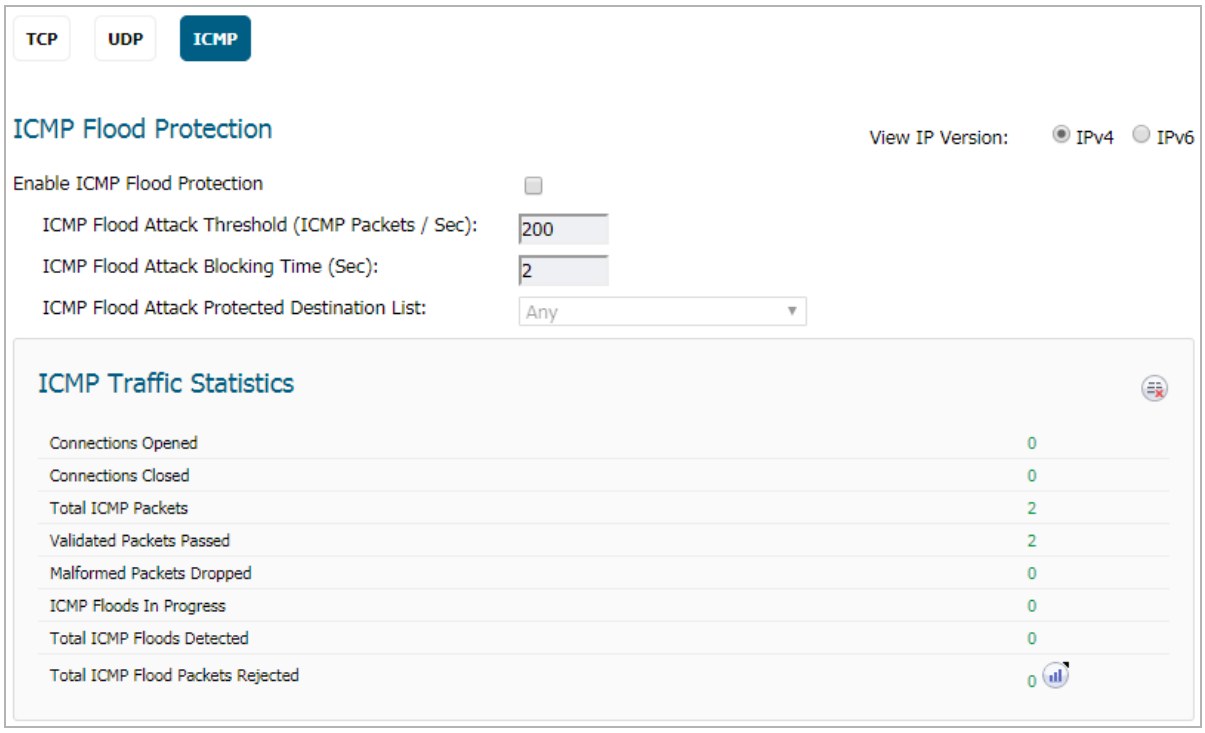

Topics:

- View IP Version on page 66
- ICMP/ICMPv6 Flood Protection on page 67
- ICMP/ICMPv6 Traffic Statistics on page 67

View IP Version

View IP Version allows you to choose the IP version: **IPv4** or **IPv6**. If you select:

- **IPv4**, the headings and options display ICMP.
- **IPv6**, the headings and options display ICMPv6.

ICMP/ICMPv6 Flood Protection

ICMP Flood Protection functions identically to UDP Flood Protection, except it monitors for ICMP/ICMPv6 Flood Attacks. The only difference is that DNS queries are not allowed to bypass ICMP Flood Protection.

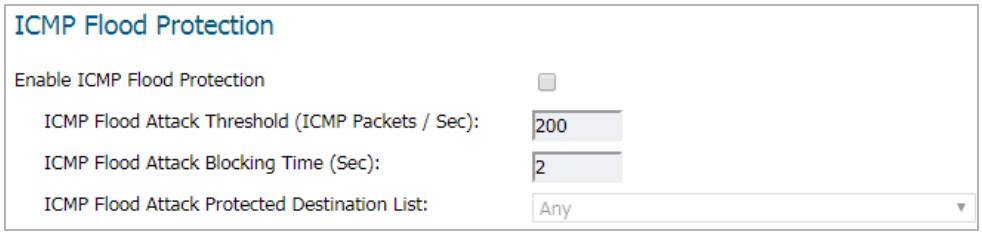

- **Enable ICMP Flood Protection** – Enables ICMP Flood Protection.

  NOTE: Enable ICMP Flood Protection must be enabled to activate the other ICMP Flood Protection options.

- **ICMP Flood Attack Threshold (ICMP Packets / Sec)** – The maximum number of ICMP packets allowed per second to be sent to a host, range, or subnet. Exceeding this threshold triggers ICMP Flood Protection. The minimum number is 10, the maximum number is 100000, and the default number is **200**.
- **ICMP Flood Attack Blocking Time (Sec)** – After the appliance detects the rate of ICMP packets exceeding the attack threshold for this duration of time, ICMP Flood Protection is activated, and the appliance will begin dropping subsequent ICMP packets. The minimum time is 1 second, the maximum time is 120 seconds, and the default time is **2** seconds.
- **ICMP Flood Attack Protected Destination List** – The destination address object or address group that will be protected from ICMP Flood Attack. The default value is **Any**.

  TIP: Select **Any** to apply the Attack Threshold to the sum of ICMP packets passing through the firewall.

ICMP/ICMPv6 Traffic Statistics



The **ICMP Traffic Statistics** table provides statistics as shown in ICMP/ICMPv6 Traffic Statistics. To clear and restart the statistics displayed by a table, click the **Clear Stats** icon for the table.

ICMP/ICMPv6 Traffic Statistics

This statistic	Is incremented/displays
Connections Opened	When a connection is opened.
Connections Closed	When a connection is closed.
Total UDP Packets	With every processed ICMP/ICMPv6 packet.

ICMP/ICMPv6 Traffic Statistics

This statistic	Is incremented/displays
Validated Packets Passed	When a ICMP/ICMPv6 packet passes checksum validation (while ICMP/ICMPv6 checksum validation is enabled).
Malformed Packets Dropped	When: • ICMP/ICMPv6 checksum fails validation (while ICMP/ICMPv6 checksum validation is enabled). • The ICMP/ICMPv6 header length is calculated to be greater than the packet's data length.
ICMP/ICMPv6 Floods In Progress	The number of individual forwarding devices currently exceeding the ICMP/ICMPv6 Flood Attack Threshold.
Total ICMP/ICMPv6 Floods Detected	The total number of events in which a forwarding device has exceeded the ICMP/ICMPv6 Flood Attack Threshold.
Total ICMP/ICMPv6 Flood Packets Rejected	The total number of packets dropped because of ICMP/ICMPv6 Flood Attack detection.Clicking on the **Statistics** icon displays a pop-up dialog showing the most recent rejected packets: 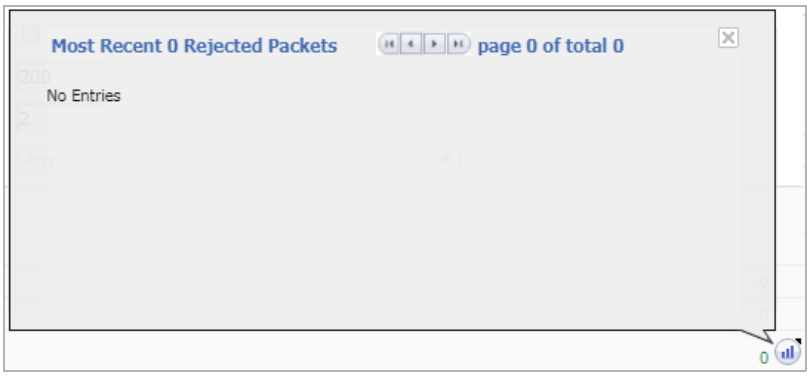

4

Configuring Firewall Multicast

- Firewall Settings > Multicast on page 69
- Multicast Snooping on page 70
- IGMP State Table on page 71
- IGMP State Table on page 71
- Enabling Multicast on page 72

Firewall Settings > Multicast

IP multicasting is a method for sending one Internet Protocol (IP) packet simultaneously to multiple hosts. Multicast is suited to the rapidly growing segment of Internet traffic - multimedia presentations and video conferencing. For example, a single host transmitting an audio or video stream and ten hosts that want to receive this stream. In multicasting, the sending host transmits a single IP packet with a specific multicast address, and the 10 hosts simply need to be configured to listen for packets targeted to that address to receive the transmission. Multicasting is a point-to-multipoint IP communication mechanism that operates in a connectionless mode - hosts receive multicast transmissions by “tuning in” to them, a process similar to tuning in to a radio.

The **MANAGE | Security Configuration > Firewall Settings > Multicast** page allows you to manage multicast traffic on the firewall.

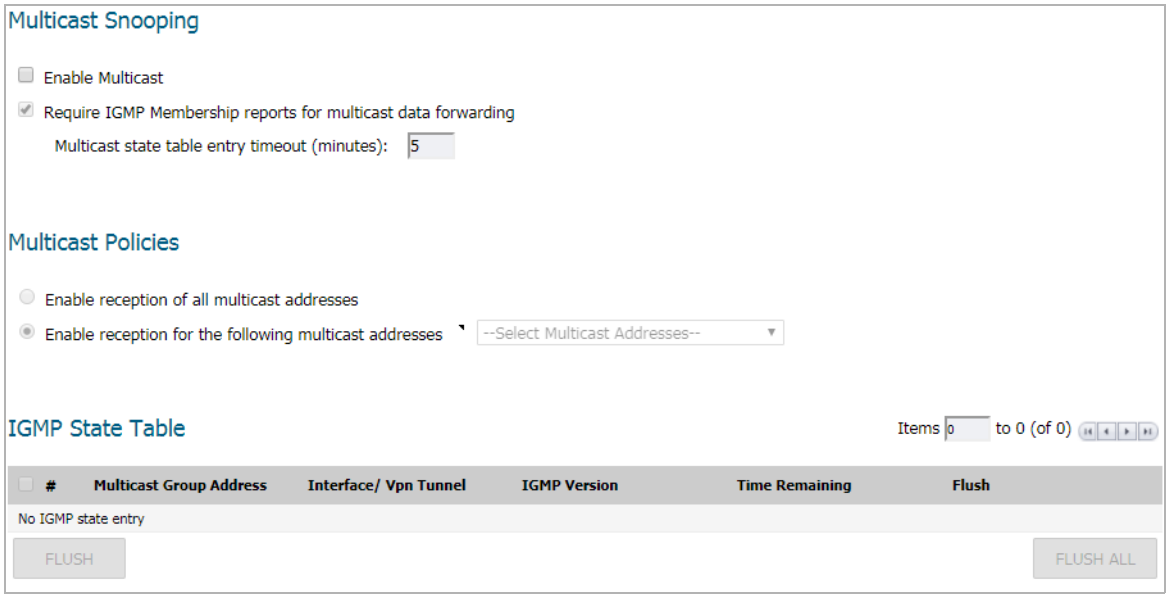

Topics:

- Multicast Snooping on page 70

- Multicast Policies on page 70
- IGMP State Table on page 71
- Enabling Multicast on LAN-Dedicated Interfaces on page 72
- Enabling Multicast Through a VPN on page 73

Multicast Snooping



- **Enable Multicast** - Select this option to support multicast traffic. This option is not selected by default.
- **Require IGMP Membership reports for multicast data forwarding** - Select this option to improve performance by regulating multicast data to be forwarded to only interfaces joined into a multicast group address using IGMP. This option is available only if Multicast is enabled. This option is selected by default.
- **Multicast state table entry timeout (minutes)** - This field has a default of **5**. The value range for this field is 5 to 60 (minutes). Update the default timer value of 5 in the following conditions:
  - You suspect membership queries or reports are being lost on the network.
  - You want to reduce the IGMP traffic on the network and currently have a large number of multicast groups or clients. This is a condition where you do not have a router to route traffic.
  - You want to synchronize the timing with an IGMP router.

Multicast Policies

TIP: Multicast must be enabled for these options to be available.

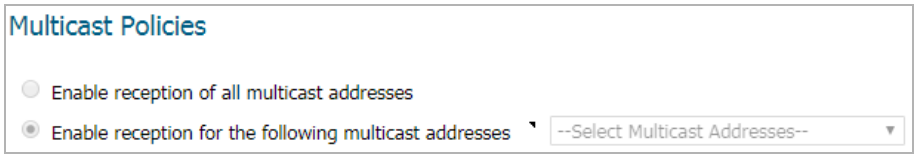

- **Enable reception of all multicast addresses** - This radio button is not enabled by default. Select this radio button to receive all (class D) multicast addresses.

  NOTE: Receiving all multicast addresses may cause your network to experience performance degradation.

- **Enable reception for the following multicast addresses** - This radio button is enabled by default. In the drop-down menu, select **Create a new multicast object** or **Create new multicast group**.

  NOTE: Only address objects and groups associated with the MULTICAST zone are available to select. Only addresses from `224.0.0.1` to `239.255.255.255` can be bound to the MULTICAST zone.

  NOTE: You can specify up to 200 total multicast addresses.

To create a multicast address object:

1 Under **Multicast Snooping**, select **Enable Multicast**.

2 Under **Multicast Policies**, in the **Enable reception for the following multicast addresses** drop-down menu, select **Create new multicast address object**. The **Add Address Object** dialog displays.

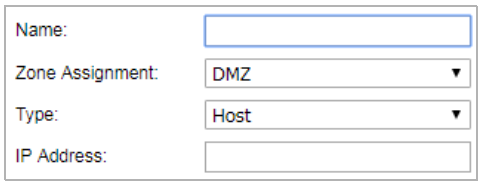

3 Configure the name of the address object in the **Name** field.

4 From the **Zone Assignment** drop-down menu, select **MULTICAST**.

5 From the **Type** drop-down menu, select **Host**, **Range**, **Network**, **MAC**, or **FQDN**.

6 Depending on your **Type** selection, the options on the dialog change. If you selected:

Type	Option(s) displayed
Host	**IP Address** – Enter the IP address of the host or network. The IP address must be in the range for multicast: `224.0.0.0` to `239.255.255.255`.
Network	• **Network** – Enter the IP address of the host or network. The IP address must be in the range for multicast: `224.0.0.0` to `239.255.255.255`. • **Netmask/Prefix Length** – Enter the netmask for the network.
Range	**Starting IP Address** and **Ending IP Address** – Enter the starting and ending IP address for the address range. The IP addresses must be in the range for multicast: `224.0.0.1` to `239.255.255.255`.
MAC	• **MAC Address** – Enter the MAC address of the host or network. • **Multi-homed Host** – Select if the MAC address is for a multihomed host. This option is selected by default.
FQDN	• **FQDN Hostname** – Enter the fully qualified domain name for the host. • **Manually set DNS entries' TTL ... (120~86400s)** – Select to enter the time-to-live (TTL or hop limit) for DNS entries. This option is not selected by default. When selected, the TTL field becomes active. The range is 120 - 86400 seconds.

7 Click **OK**.

IGMP State Table

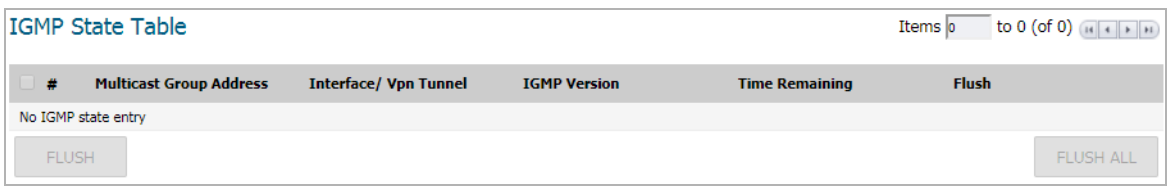

This section provides descriptions of the fields in the **IGMP State Table**.

- **Multicast Group Address**—Provides the multicast group address the interface is joined to.
- **Interface / VPN Tunnel**—Provides the interface (such as **LAN**) for the VPN policy.
- **IGMP Version**—Provides the IGMP version (such as **V2** or **V3**).

- **Time Remaining**
- **Flush** — Provides an icon to flush that particular entry.
- **FLUSH** and **FLUSH ALL** buttons—To flush a specific entry immediately, check the box to the left of the entry and click **FLUSH**. Click **FLUSH ALL** to immediately flush all entries.

Enabling Multicast

Topics:

- Enabling Multicast on LAN-Dedicated Interfaces on page 72
- Enabling Multicast Through a VPN on page 73

Enabling Multicast on LAN-Dedicated Interfaces

Topics:

- Enabling Multicast on a LAN-Dedicated Interface on page 72
- Enabling Multicast Support for Address Objects over a VPN Tunnel on page 72

Enabling Multicast on a LAN-Dedicated Interface

To enable multicast support on the LAN-dedicated interfaces of your firewall:

1. Go to the **MANAGE | Security Configuration > Firewall Settings > Multicast** page.
2. Under **Multicast Snooping**, select **Enable Multicast**.
3. Under **Multicast Policies**, choose **Enable the reception of all multicast addresses**.
4. Click **ACCEPT**.
5. Go to the **MANAGE | System Setup > Network > Interfaces** page.
6. Click the **Configure** icon for the LAN interface you want to configure. The **Edit Interface** dialog displays.
7. Click **Advanced**.
8. Select **Enable Multicast Support**.
9. Click **OK**.

Enabling Multicast Support for Address Objects over a VPN Tunnel

To enable multicast support for address objects over a VPN tunnel:

1. Go to the **MANAGE | Security Configuration > Firewall Settings > Multicast** page.
2. Under **Multicast Snooping**, select **Enable Multicast**.
3. Under **Multicast Policy**, select **Enable the reception for the following multicast addresses**.

4 From the drop-down menu, select **Create new multicast address object.** The **Add Address Object** dialog appears.

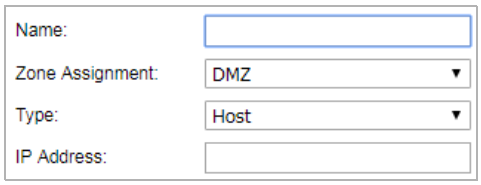

5 In the **Name** field, enter a name for your multicast address object.

6 From the **Zone Assignment** drop-down menu, select a zone: **DMZ**, **LAN**, **MULTICAST**, **SSLVPN**, **VPN**, **WAN**, or **WLAN**.

7 When you select a type from the **Type** drop-down menu**,** the other options change, depending on the selection. If you select:

- **Host**, enter an **IP address** in the **IP Address** field.
- **Range**, enter the starting and ending IP addresses in the **Starting IP Address** and the **Ending IP Address**.
- **Network**, enter the network IP address in the **Netmask** field and a netmask or prefix length in the **Netmask/Prefix Length** field.
- **MAC**, enter the MAC address in the **MAC Address** field and select the **Multi-homed host** checkbox (which is selected by default).
- **FQDN**, enter the FQDN hostname in the **FQDN Hostname** field.

8 Click **OK**.

9 Go to the **MANAGE | Connectivity > VPN > Settings** page.

10 In the **VPN Policies** table, click the **Configure** icon for the Group VPN policy you want to configure. The **VPN Policy** dialog displays.

11 Click **Advanced**.

12 In the **Advanced Settings** section, select **Enable Multicast**.

13 Click **OK**.

Enabling Multicast Through a VPN

To enable multicast across the WAN through a VPN:

1 Enable multicast globally:

   a Navigate to the **MANAGE | Security Configuration > Firewall Settings > Multicast** page.

   b Select **Enable Multicast**.

   c Click the **ACCEPT** button.

   d Repeat Step a through Step c for each interface on all participating security appliances.

2 Enable multicast support on each individual interface that will be participating in the multicast network.

   a Navigate to the **MANAGE | System Setup > Network > Interfaces** page

   b Click the **Edit** icon of the participating interface. The **Edit Interface** dialog displays.

c Click **Advanced**.

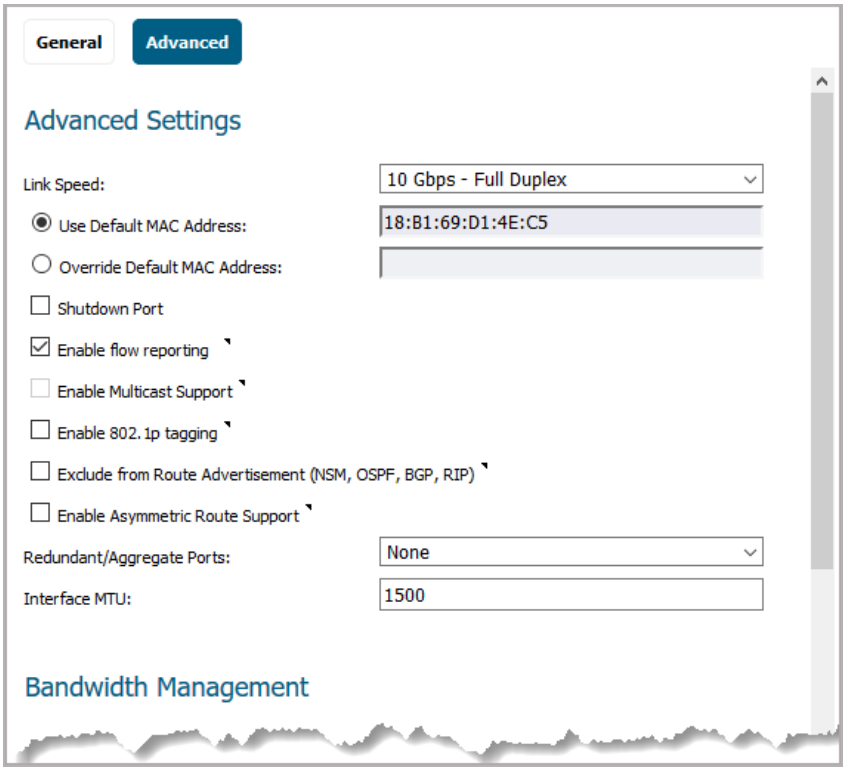

d Select the **Enable Multicast Support** checkbox.

e Click **OK**.

f Repeat Step a through Step e for each participating interface on all participating appliances.

3 Enable multicast on the VPN policies between the security appliances.

a Navigate to the **MANAGE | Connectivity > VPN > Base Settings** page.

b Click the **Edit** icon of a policy in which include multicasting. The **VPN Policy** dialog displays.

c Click **Advanced**.

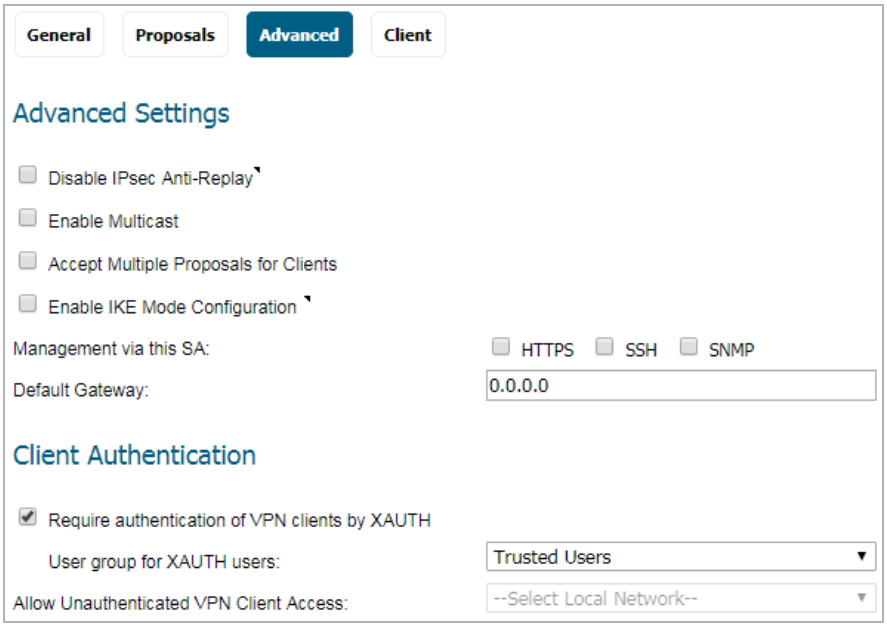

NOTE: The default WLAN'MULTICAST access rule for IGMP traffic is set to DENY. This will need to be changed to ALLOW on all participating appliances to enable multicast if they have multicast clients on their WLAN zones.

d In the **Advanced Settings** section, select **Enable Multicast**.

e Click **OK**.

4 Verify the tunnels are active between the sites.

5 Start the multicast server application and client applications. As multicast data is sent from the multicast server to the multicast group (`224.0.0.0` through `239.255.255.255`), the firewall queries its IGMP state table for that group to determine where to deliver that data. Similarly, when the appliance receives that data at the VPN zone, the appliance queries its IGMP State Table to determine where it should deliver the data.

The IGMP State Tables (upon updating) should provide information indicating that there is a multicast client on the X3 interface, and across the vpnMcastServer tunnel for the `224.15.16.17` group.

NOTE: By selecting **Enable reception of all multicast addresses**, you might see entries other than those you are expecting to see when viewing your **IGMP State Table**. These are caused by other multicast applications that might be running on your hosts.

5

Managing Quality of Service

- Firewall Settings > Quality of Service Mapping on page 76
- Classification on page 76
- Marking on page 77
- Conditioning on page 78
- 802.1p and DSCP QoS on page 79
- Bandwidth Management on page 90
- Glossary on page 90

Firewall Settings > Quality of Service Mapping

Quality of Service (QoS) refers to a diversity of methods intended to provide predictable network behavior and performance. This sort of predictability is vital to certain types of applications, such as Voice over IP (VoIP), multimedia content, or business-critical applications such as order or credit-card processing. No amount of bandwidth can provide this sort of predictability, because any amount of bandwidth will ultimately be 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.

Topics:

- Classification
- Marking
- Conditioning
- 802.1p and DSCP QoS
- Bandwidth Management
- Glossary

Classification

Classification is necessary as a first step so that traffic in need of management can be identified. SonicOS uses Access Rules as the interface to classification of traffic. This provides fine controls using combinations of Address Object, Service Object, and Schedule Object elements, allowing for classification criteria as general as **all HTTP traffic** and as specific as **SSH traffic from hostA to serverB on Wednesdays at 2:12am**.

SonicWall network security appliances have the ability to recognize, map, modify, and generate the industry-standard external CoS designators, DSCP and 802.1p (refer to the section 802.1p and DSCP QoS on page 79).

When identified, or classified, traffic can be managed. Management can be performed internally by SonicOS Bandwidth Management (BWM), which is perfectly effective as long as the network is a fully contained autonomous system. Once external or intermediate elements are introduced, such as foreign network infrastructures with unknown configurations, or other hosts contending for bandwidth (for example, the Internet) the ability to offer guarantees and predictability are diminished. In other words, as long as the endpoints of the network and everything in between are within your management, BWM will work exactly as configured. Once external entities are introduced, the precision and efficacy of BWM configurations can begin to degrade.

But all is not lost. Once SonicOS classifies the traffic, it can **tag** the traffic to communicate this classification to certain external systems that are capable of abiding by CoS tags; thus they too can participate in providing QoS.

NOTE: Many service providers do not support CoS tags such as 802.1p or DSCP. Also, most network equipment with standard configurations will not be able to recognize 802.1p tags, and could drop tagged traffic.

Although DSCP will not cause compatibility issues, many service providers will simply strip or ignore the DSCP tags, disregarding the code points.

If you wish to use 802.1p or DSCP marking on your network or your service provider's network, you must first establish that these methods are supported. Verify that your internal network equipment can support CoS priority marking, and that it is correctly configured to do so. Check with your service provider – some offer fee-based support for QoS using these CoS methods.

Marking

After the traffic has been classified, if it is to be handled by QoS capable external systems (for example, CoS aware switches or routers as might be available on a premium service provider's infrastructure, or on a private WAN), it must be tagged so that the external systems can make use of the classification, and provide the correct handling and Per Hop Behaviors (PHB).

Originally, this was attempted at the IP layer (layer 3) with RFC791's three Precedence bits and RFC1394 ToS (type of service) field, but this was used by a grand total of 17 people throughout history. Its successor, RFC2474 introduced the much more practical and widely used DSCP (Differentiated Services Code Point) which offered up to 64 classifications, as well as user-definable classes. DSCP was further enhanced by RFC2598 (Expedited Forwarding, intended to provide leased-line behaviors) and RFC2697 (Assured Forwarding levels within classes, also known as Gold, Silver, and Bronze levels).

DSCP is a safe marking method for traffic that traverses public networks because there is no risk of incompatibility. At the very worst, a hop along the path might disregard or strip the DSCP tag, but it will rarely mistreat or discard the packet.

The other prevalent method of CoS marking is IEEE 802.1p. 802.1p occurs at the MAC layer (layer 2) and is closely related to IEEE 802.1Q VLAN marking, sharing the same 16-bit field, although it is actually defined in the IEEE 802.1D standard. Unlike DSCP, 802.1p will only work with 802.1p capable equipment, and is not universally interoperable. Additionally, 802.1p, because of its different packet structure, can rarely traverse wide-area networks, even private WANs. Nonetheless, 802.1p is gaining wide support among Voice and Video over IP vendors, so a solution for supporting 802.1p across network boundaries (i.e. WAN links) was introduced in the form of **802.1p to DSCP mapping**.

802.1p to DSCP mapping allows 802.1p tags from one LAN to be mapped to DSCP values by SonicOS, allowing the packets to safely traverse WAN links. When the packets arrive on the other side of the WAN or VPN, the receiving SonicOS appliance can then map the DSCP tags back to 802.1p tags for use on that LAN. Refer to 802.1p and DSCP QoS on page 79 for more information.

Conditioning

The traffic can be conditioned (or managed) using any of the many policing, queuing, and shaping methods available. SonicOS provides internal conditioning capabilities with its Egress and Ingress Bandwidth Management (BWM), detailed in Bandwidth Management on page 90. SonicOS's BWM is a perfectly effective solution for fully autonomous private networks with sufficient bandwidth, but can become somewhat less effective as more unknown external network elements and bandwidth contention are introduced. Refer to DSCP marking: Example scenario for a description of contention issues.

Topics:

- Site to Site VPN over QoS Capable Networks on page 78
- Site to Site VPN over Public Networks on page 78

Site to Site VPN over QoS Capable Networks

If the network path between the two end points is QoS aware, SonicOS can DSCP tag the inner encapsulate packet so that it is interpreted correctly at the other side of the tunnel, and it can also DSCP tag the outer ESP encapsulated packet so that its class can be interpreted and honored by each hop along the transit network. SonicOS can map 802.1p tags created on the internal networks to DSCP tags so that they can safely traverse the transit network. Then, when the packets are received on the other side, the receiving SonicWall appliance can translate the DSCP tags back to 802.1p tags for interpretation and honoring by that internal network.

Site to Site VPN over Public Networks

SonicOS integrated BWM is very effective in managing traffic between VPN connected networks because ingress and egress traffic can be classified and controlled at both endpoints. If the network between the endpoints is non QoS aware, it regards and treats all VPN ESP equally. Because there is typically no control over these intermediate networks or their paths, it is difficult to fully guarantee QoS, but BWM can still help to provide more predictable behavior.

Site to site VPN over public networks

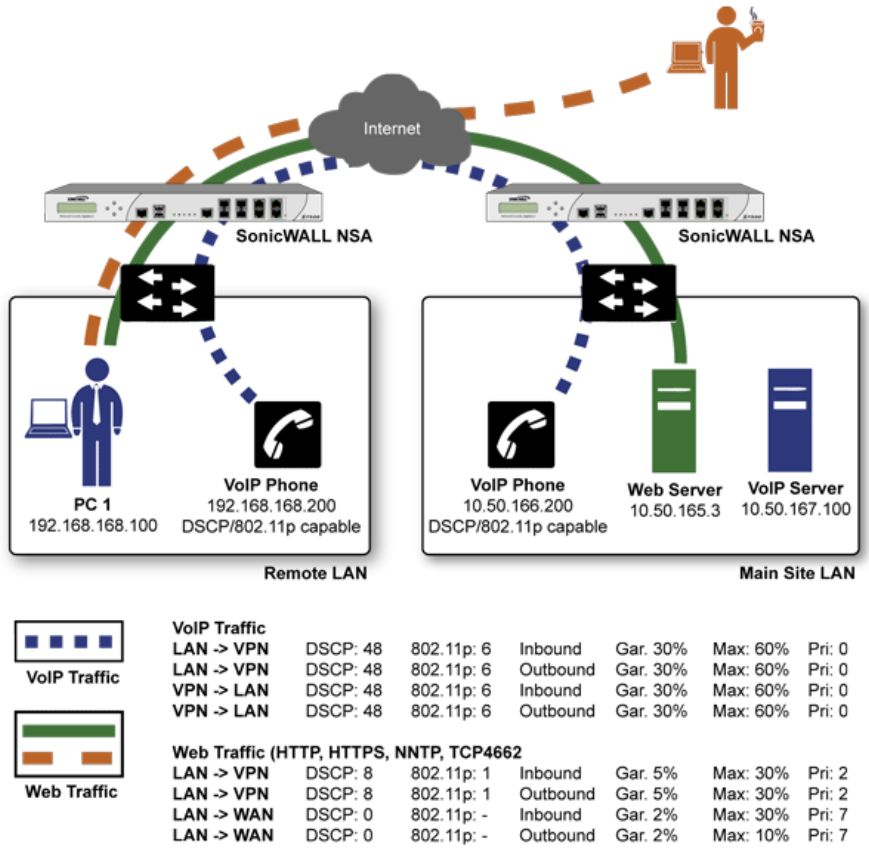

To provide end-to-end QoS, business-class service providers are increasingly offering traffic conditioning services on their IP networks. These services typically depend on the customer premise equipment to classify and tag the traffic, generally using a standard marking method such as DSCP. SonicOS has the ability to DSCP mark traffic after classification, as well as the ability to map 802.1p tags to DSCP tags for external network traversal and CoS preservation. For VPN traffic, SonicOS can DSCP mark not only the internal (payload) packets, but the external (encapsulating) packets as well so that QoS capable service providers can offer QoS even on encrypted VPN traffic.

The actual conditioning method employed by service providers varies from one to the next, but it generally involves a class-based queuing method such as Weighted Fair Queuing for prioritizing traffic, as well a congestion avoidance method, such as tail-drop or Random Early Detection.

802.1p and DSCP QoS

Topics:

- Enabling 802.1p on page 79
- DSCP Marking on page 83

Enabling 802.1p

SonicOS supports layer 2 and layer 3 CoS methods for broad interoperability with external systems participating in QoS enabled environments. The layer 2 method is the IEEE 802.1p standard wherein 3-bits of an additional

16-bits inserted into the header of the Ethernet frame can be used to designate the priority of the frame, as illustrated in the following figure:

Ethernet data frame

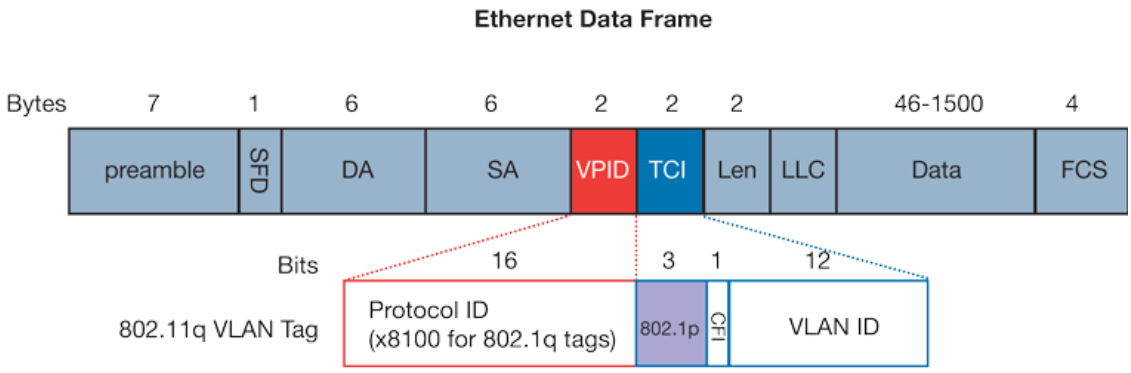

- **TPID**: Tag Protocol Identifier begins at byte 12 (after the 6 byte destination and source fields), is 2 bytes long, and has an Ether type of 0x8100 for tagged traffic.
- **802.1p**: The first three bits of the TCI (Tag Control Information – beginning at byte 14, and spanning 2 bytes) define user priority, giving eight (2^3) priority levels. IEEE 802.1p defines the operation for these 3 user priority bits.
- **CFI**: Canonical Format Indicator is a single-bit flag, always set to zero for Ethernet switches. CFI is used for compatibility reasons between Ethernet networks and Token Ring networks. If a frame received at an Ethernet port has a CFI set to 1, then that frame should not be forwarded as it is to an untagged port.
- **VLAN ID**: VLAN ID (starts at bit 5 of byte 14) is the identification of the VLAN. It has 12-bits and allows for the identification of 4,096 (2^12) unique VLAN ID's. Of the 4,096 possible IDs, an ID of 0 is used to identify priority frames, and an ID of 4,095 (FFF) is reserved, so the maximum possible VLAN configurations are 4,094.

802.1p support begins by enabling 802.1p marking on the interfaces which you wish to have process 802.1p tags. 802.1p can be enabled on any Ethernet interface on any SonicWall appliance.

The behavior of the 802.1p field within these tags can be controlled by Access Rules. The default 802.1p Access Rule action of **None** will reset existing 802.1p tags to 0, unless otherwise configured (see Managing QoS Marking on page 86 for details).

Enabling 802.1p marking will allow the target interface to recognize incoming 802.1p tags generated by 802.1p capable network devices, and will also allow the target interface to generate 802.1p tags, as controlled by Access Rules. Frames that have 802.1p tags inserted by SonicOS will bear VLAN ID 0.

802.1p tags will only be inserted according to Access Rules, so enabling 802.1p marking on an interface will not, at its default setting, disrupt communications with 802.1p-incapable devices.

802.1p requires the specific support by the networking devices with which you wish to use this method of prioritization. Many voice and video over IP devices provide support for 802.1p, but the feature must be enabled. Check your equipment's documentation for information on 802.1p support if you are unsure. Similarly, many server and host network cards (NICs) have the ability to support 802.1p, but the feature is usually disabled by default. On Win32 operating systems, you can check for and configure 802.1p settings on the **Advanced** view

of the Properties page of your network card. If your card supports 802.1p, it is listed as **802.1p QoS, 802.1p Support, QoS Packet Tagging** or something similar:



To process 802.1p tags, the feature must be present and enabled on the network interface. The network interface will then be able to generate packets with 802.1p tags, as governed by QoS capable applications. By default, general network communications will not have tags inserted so as to maintain compatibility with 802.1p-incapable devices.

NOTE: If your network interface does not support 802.1p, it will not be able to process 802.1p tagged traffic, and will ignore it. Make certain when defining Access Rules to enable 802.1p marking that the target devices are 802.1p capable.

It should also be noted that when performing a packet capture (for example, with the diagnostic tool Ethereal) on 802.1p capable devices, some 802.1p capable devices will not show the 802.1q header in the packet capture. Conversely, a packet capture performed on an 802.1p-incapable device will almost invariably show the header, but the host will be unable to process the packet.

Before moving on to For more information, see Managing QoS Marking on page 86., it is important to introduce 'DSCP Marking' because of the potential interdependency between the two marking methods, as well as to explain why the interdependency exists.

DSCP marking: Example scenario

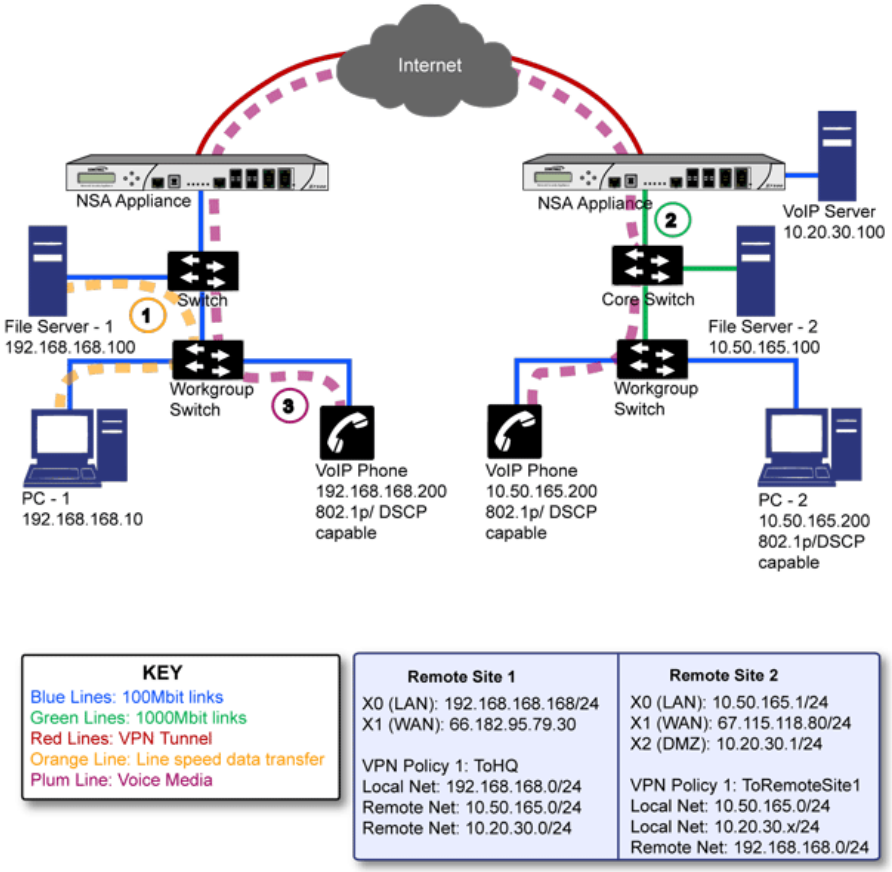

In the scenario in DSCP marking: Example scenario, we have **Remote Site 1** connected to 'Main Site' by an IPsec VPN. The company uses an internal 802.1p/DSCP capable VoIP phone system, with a private VoIP signaling server hosted at the Main Site. The Main Site has a mixed gigabit and Fast-Ethernet infrastructure, while Remote Site 1 is all Fast Ethernet. Both sites employ 802.1p capable switches for prioritization of internal traffic.

1 PC-1 at Remote Site 1 is transferring a 23 terabyte PowerPoint™ presentation to File Server 1, and the 100mbit link between the workgroup switch and the upstream switch is completely saturated.

2 At the Main Site, a caller on the 802.1p/DSCP capable VoIP Phone `10.50.165.200` initiates a call to the person at VoIP phone `192.168.168.200`. The calling VoIP phone 802.1p tags the traffic with priority tag 6 (voice), and DSCP tags the traffic with a tag of 48.

   a If the link between the Core Switch and the firewall is a VLAN, some switches will include the received 802.1p priority tag, in addition to the DSCP tag, in the packet sent to the firewall; this behavior varies from switch to switch, and is often configurable.

   b If the link between the Core Switch and the firewall is not a VLAN, there is no way for the switch to include the 802.1p priority tag. The 802.1p priority is removed, and the packet (including only the DSCP tag) is forwarded to the firewall.

   When the firewall sent the packet across the VPN/WAN link, it could include the DSCP tag in the packet, but it is not possible to include the 802.1p tag. This would have the effect of losing all prioritization information for the VoIP traffic, because when the packet arrived at the Remote Site, the switch would have no 802.1p MAC layer information with which to prioritize the traffic. The Remote Site switch would treat the VoIP traffic the same as the lower-priority file transfer because of the link saturation, introducing delay—maybe even dropped packets—to the VoIP flow, resulting in call quality degradation.

   So how can critical 802.1p priority information from the Main Site LAN persist across the VPN/WAN link to Remote Site LAN? Through the use of QoS Mapping.

QoS Mapping is a feature which converts layer 2 802.1p tags to layer 3 DSCP tags so that they can safely traverse (in mapped form) 802.1p-incapable links; when the packet arrives for delivery to the next 802.1p-capable segment, QoS Mapping converts from DSCP back to 802.1p tags so that layer 2 QoS can be honored.

In our above scenario, the firewall at the Main Site assigns a DSCP tag (for example, value **48**) to the VoIP packets, as well as to the encapsulating ESP packets, allowing layer 3 QoS to be applied across the WAN. This assignment can occur either by preserving the existing DSCP tag, or by mapping the value from an 802.1p tag, if present. When the VoIP packets arrive at the other side of the link, the mapping process is reversed by the receiving SonicWall, mapping the DSCP tag back to an 802.1p tag.

3 The receiving SonicWall at the Remote Site is configured to map the DSCP tag range 48-55 to 802.1p tag 6. When the packet exits the firewall, it will bear 802.1p tag 6. The Switch will recognize it as voice traffic, and will prioritize it over the file-transfer, guaranteeing QoS even in the event of link saturation.

DSCP Marking

DSCP (Differentiated Services Code Point) marking uses 6-bits of the 8-bit ToS field in the IP Header to provide up to 64 classes (or code points) for traffic. Since DSCP is a layer 3 marking method, there is no concern about compatibility as there is with 802.1p marking. Devices that do not support DSCP will simply ignore the tags, or at worst, they will reset the tag value to 0.

DSCP marking: IP packet

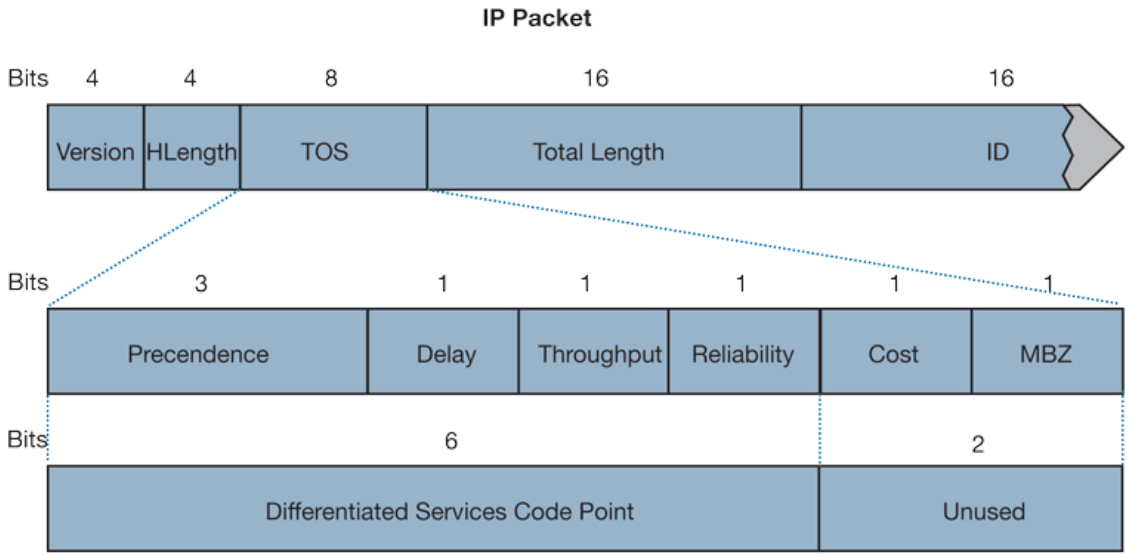

DSCP marking: IP packet depicts an IP packet, with a close-up on the ToS portion of the header. The ToS bits were originally used for Precedence and ToS (delay, throughput, reliability, and cost) settings, but were later repurposed by RFC2474 for the more versatile DSCP settings.

DSCP marking: Commonly used code points shows the commonly used code points, as well as their mapping to the legacy Precedence and ToS settings.

DSCP marking: Commonly used code points

DSCP	DSCP Description	Legacy IP Precedence	Legacy IP ToS (D, T, R)
0	Best effort	0 (Routine – 000)	-
8	Class 1	1 (Priority – 001)	-
10	Class 1, gold (AF11)	1 (Priority – 001)	T
12	Class 1, silver (AF12)	1 (Priority – 001)	D

DSCP marking: Commonly used code points

DSCP	DSCP Description	Legacy IP Precedence	Legacy IP ToS (D, T, R)
14	Class 1, bronze (AF13)	1 (Priority – 001)	D, T
16	Class 2	2 (Immediate – 010)	-
18	Class 2, gold (AF21)	2 (Immediate – 010)	T
20	Class 2, silver (AF22)	2 (Immediate – 010)	D
22	Class 2, bronze (AF23)	2 (Immediate – 010)	D, T
24	Class 3	3 (Flash – 011)	-
26	Class 3, gold (AF31)	3 (Flash – 011)	T
27	Class 3, silver (AF32)	3 (Flash – 011)	D
30	Class 3, bronze (AF33)	3 (Flash – 011)	D, T
32	Class 4	4 (Flash Override – 100)	-
34	Class 4, gold (AF41)	4 (Flash Override – 100)	T
36	Class 4, silver (AF42)	4 (Flash Override – 100)	D
38	Class 4, bronze (AF43)	4 (Flash Override – 100)	D, T
40	Express forwarding	5 (CRITIC/ECP[1] – 101)	-
46	Expedited forwarding (EF)	5 (CRITIC/ECP – 101)	D, T
48	Control	6 (Internet Control – 110)	-
56	Control	7 (Network Control – 111)	-

1. ECP: Elliptic Curve Group

DSCP marking can be performed on traffic to/from any interface and to/from any zone type, without exception. DSCP marking is controlled by Access Rules, from the **QoS** view, and can be used in conjunction with 802.1p marking, as well as with SonicOS's internal bandwidth management.

Topics:

- DSCP Marking and Mixed VPN Traffic on page 84
- Configure for 802.1p CoS 4 – Controlled load on page 85
- QoS Mapping on page 85
- Managing QoS Marking on page 86

DSCP Marking and Mixed VPN Traffic

Among their many security measures and characteristics, IPsec VPNs employ anti-replay mechanisms based upon monotonically incrementing sequence numbers added to the ESP header. Packets with duplicate sequence numbers are dropped, as are packets that do not adhere to sequence criteria. One such criterion governs the handling of out-of-order packets. SonicOS provides a replay window of 64 packets, i.e. if an ESP packet for a Security Association (SA) is delayed by more than 64 packets, the packet will be dropped.

This should be considered when using DSCP marking to provide layer 3 QoS to traffic traversing a VPN. If you have a VPN tunnel that is transporting a diversity of traffic, some that is being DSCP tagged high priority (for example, VoIP), and some that is DSCP tagged low-priority, or untagged/best-effort (for example, FTP), your service provider will prioritize the handling and delivery of the high-priority ESP packets over the best-effort ESP packets. Under certain traffic conditions, this can result in the best-effort packets being delayed for more than 64 packets, causing them to be dropped by the receiving SonicWall's anti-replay defenses.

If symptoms of such a scenario emerge (for example, excessive retransmissions of low-priority traffic), it is recommended that you create a separate VPN policy for the high-priority and low-priority classes of traffic. This

is most easily accomplished by placing the high-priority hosts (for example, the VoIP network) on their own subnet.

Configure for 802.1p CoS 4 – Controlled load

If you want to change the inbound mapping of DSCP tag 15 from its default 802.1p mapping of 1 to an 802.1p mapping of 2, it would have to be done in two steps because mapping ranges cannot overlap. Attempting to assign an overlapping mapping will give the error DSCP range already exists or overlaps with another range. First, you will have to remove 15 from its current end-range mapping to 802.1p CoS 1 (changing the end-range mapping of 802.1p CoS 1 to DSCP 14), then you can assign DSCP 15 to the start-range mapping on 802.1p CoS 2.

QoS Mapping

The primary objective of QoS Mapping is to allow 802.1p tags to persist across non-802.1p compliant links (for example, WAN links) by mapping them to corresponding DSCP tags before sending across the WAN link, and then mapping from DSCP back to 802.1p upon arriving at the other side, as shown in QoS mapping.

QoS mapping

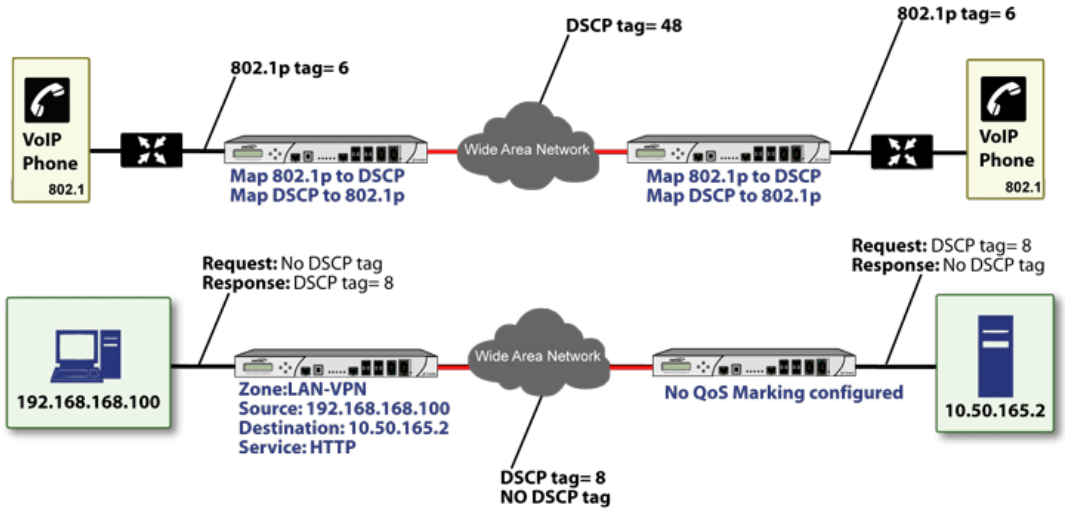

NOTE: Mapping will not occur until you assign **Map** as an action of the QoS view of an Access Rule. The mapping table only defines the correspondence that will be employed by an Access Rule's Map action.

802.1p Class Of Service	To DSCP	From DSCP Range	Configure
0 - Best effort	0 - Best effort/Default	0-7	
1 - Background	8 - Class 1	8-15	
2 - Spare	16 - Class 2	16-23	
3 - Excellent effort	24 - Class 3	24-31	
4 - Controlled load	32 - Class 4	32-39	
5 - Video (<100ms latency)	40 - Express forwarding	40-47	
6 - Voice (<10ms latency)	48 - Control	48-55	
7 - Network control	56 - Control	56-63	

RESET QOS SETTINGS

For example, according to the default table, an 802.1p tag with a value of **2** will be outbound mapped to a DSCP value of **16**, while a DSCP tag of **43** will be inbound mapped to an 802.1 value of **5**.

Each of these mappings can be reconfigured. If you wanted to change the outbound mapping of 802.1p tag **4** from its default DSCP value of **32** to a DSCP value of **43**, you can click the **Configure** icon for **4 – Controlled load** and select the new **To DSCP** value from the drop-down box:

802.1p CoS 1 end-range remap

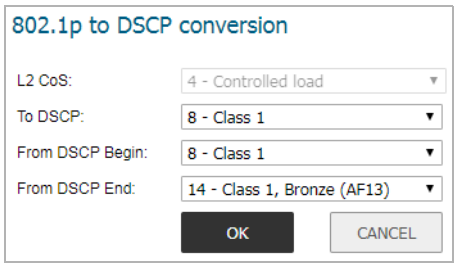

802.1p CoS 2 start-range remap

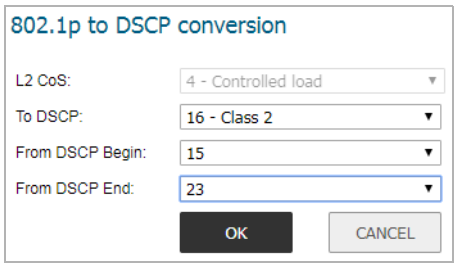

You can restore the default mappings by clicking the **Reset QoS Settings** button.

Managing QoS Marking

QoS marking is configured from the **QoS** view of the **Add/Edit Rule** dialog of the **Policies| Rules > Access Rules** page:

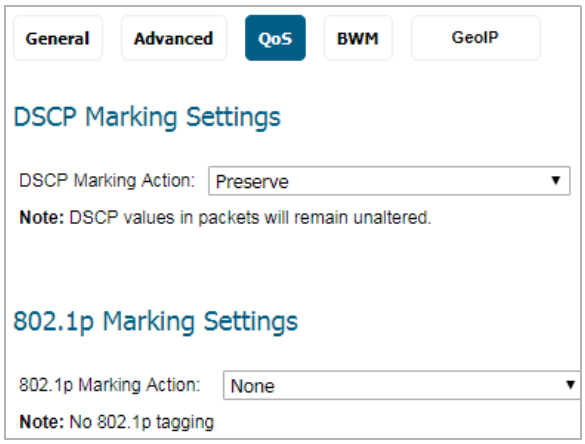

Both 802.1p and DSCP marking as managed by SonicOS Access Rules provide four actions: **None**, **Preserve**, **Explicit**, and **Map**. The default action for DSCP is **Preserve** and the default action for 802.1p is **None**.

QoS marking: Behavior describes the behavior of each action on both methods of marking:

QoS marking: Behavior

Action	802.1p (layer 2 CoS)	DSCP (layer 3)	Notes
None	When packets matching this class of traffic (as defined by the Access Rule) are sent out the egress interface, no 802.1p tag will be added.	The DSCP tag is explicitly set (or reset) to 0.	If the target interface for this class of traffic is a VLAN subinterface, the 802.1p portion of the 802.1q tag will be explicitly set to 0. If this class of traffic is destined for a VLAN and is using 802.1p for prioritization, a specific Access Rule using the **Preserve**, **Explicit**, or **Map** action should be defined for this class of traffic.
Preserve	Existing 802.1p tag will be preserved.	Existing DSCP tag value will be preserved.	
Explicit	An explicit 802.1p tag value can be assigned (0-7) from a drop-down menu that will be presented.	An explicit DSCP tag value can be assigned (0-63) from a drop-down menu that will be presented.	If either the 802.1p or the DSCP action is set to **Explicit** while the other is set to **Map**, the explicit assignment occurs first, and then the other is mapped according to that assignment.
Map	The mapping setting defined in the **Firewall Settings > QoS Mapping** page will be used to map from a DSCP tag to an 802.1p tag	The mapping setting defined in the **Firewall Settings > QoS Mapping** page will be used to map from an 802.1 tag to a DSCP tag. An additional checkbox will be presented to Allow **802.1p Marking to override DSCP values.** Selecting this checkbox will assert the mapped 802.1p value over any DSCP value that might have been set by the client. This is useful to override clients setting their own DSCP CoS values.	If **Map** is set as the action on both DSCP and 802.1p, mapping will only occur in one direction: if the packet is from a VLAN and arrives with an 802.1p tag, then DSCP will be mapped from the 802.1p tag; if the packet is destined to a VLAN, then 802.1p will be mapped from the DSCP tag.

For example, refer to Bi-directional DSCP tag action, which provides a bi-directional DSCP tag action.

Bi-directional DSCP tag action

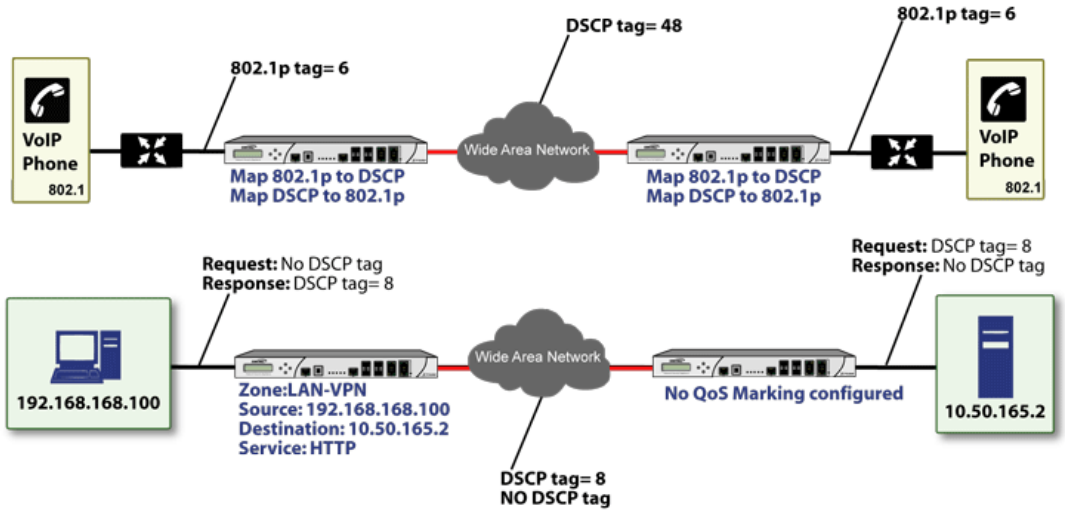

HTTP access from a Web-browser on 192.168.168.100 to the Web server on `10.50.165.2` will result in the tagging of the inner (payload) packet and the outer (encapsulating ESP) packets with a DSCP value of 8. When the packets emerge from the other end of the tunnel, and are delivered to `10.50.165.2`, they will bear a DSCP tag of 8. When `10.50.165.2` sends response packets back across the tunnel to `192.168.168.100` (beginning with the very first SYN/ACK packet) the Access Rule will tag the response packets delivered to `192.168.168.100` with a DSCP value of 8.

This behavior applies to all four QoS action settings for both DSCP and 802.1p marking.

One practical application for this behavior would be configuring an 802.1p marking rule for traffic destined for the VPN zone. Although 802.1p tags cannot be sent across the VPN, reply packets coming back across the VPN can be 802.1p tagged on egress from the tunnel. This requires that 802.1p tagging is active of the physical egress interface, and that the [Zone] > VPN Access Rule has an 802.1p marking action other than None.

After ensuring 802.1p compatibility with your relevant network devices, and enabling 802.1p marking on applicable SonicWall interfaces, you can begin configuring Access Rules to manage 802.1p tags.

The **Remote Site 1** network could have two Access Rules configured as in Remote site 1: Sample access rule configuration.

Remote site 1: Sample access rule configuration

Setting	Access Rule 1	Access Rule 2
General View		
Action	Allow	Allow
From Zone	LAN	VPN
To Zone	VPN	LAN
Service	VOIP	VOIP
Source	Lan Primary Subnet	Main Site Subnets
Destination	Main Site Subnets	Lan Primary Subnet
Users Allowed	All	All
Schedule	Always on	Always on
Enable Logging	Enabled	Enabled
Allow Fragmented Packets	Enabled	Enabled
Qos View		

Remote site 1: Sample access rule configuration

Setting	Access Rule 1	Access Rule 2
DSCP Marking Action	Map	Map
Allow 802.1p Marking to override DSCP values	Enabled	Enabled
802.1p Marking Action	Map	Map

The first Access Rule (governing **LAN>VPN**) would have the following effects:

- **VoIP** traffic (as defined by the Service Group) from **LAN Primary Subnet** destined to be sent across the VPN to **Main Site Subnets** would be evaluated for both DSCP and 802.1p tags.
    - The combination of setting both DSCP and 802.1p marking actions to **Map** is described in the table earlier in Managing QoS Marking on page 86.
    - Sent traffic containing only an 802.1p tag (for example, CoS = 6) would have the VPN-bound inner (payload) packet DSCP tagged with a value of 48. The outer (ESP) packet would also be tagged with a value of 48.
    - Assuming returned traffic has been DSCP tagged (CoS = 48) by the firewall at the Main Site, the return traffic will be 802.1p tagged with CoS = 6 on egress.
    - Sent traffic containing only a DSCP tag (for example, CoS = 48) would have the DSCP value preserved on both inner and outer packets.
    - Assuming returned traffic has been DSCP tagged (CoS = 48) by the firewall at the Main Site, the return traffic will be 802.1p tagged with CoS = 6 on egress.
    - Sent traffic containing only both an 802.1p tag (for example, CoS = 6) and a DSCP tag (for example, CoS = 63) would give precedence to the 802.1p tag and would be mapped accordingly. The VPN-bound inner (payload) packet DSCP would be tagged with a value of 48. The outer (ESP) packet would also be tagged with a value of 48.

Assuming returned traffic has been DSCP tagged (CoS = 48) by the firewall at the Main Site, the return traffic will be 802.1p tagged with CoS = 6 on egress.

To examine the effects of the second Access Rule (VPN>LAN), we'll look at the Access Rules configured at the Main Site, as shown in Main site: Sample access rule configurations.

Main site: Sample access rule configurations

Setting	Access Rule 1	Access Rule 2
General View		
Action	Allow	Allow
From Zone	LAN	VPN
To Zone	VPN	LAN
Service	VOIP	VOIP
Source	Lan Subnets	Remote Site 1 Subnets
Destination	Remote Site 1 Subnets	Lan Subnets
Users Allowed	All	All
Schedule	Always on	Always on
Enable Logging	Enabled	Enabled
Allow Fragmented Packets	Enabled	Enabled
Qos View		
DSCP Marking Action	Map	Map

Main site: Sample access rule configurations

Setting	Access Rule 1	Access Rule 2
Allow 802.1p Marking to override DSCP values	Enabled	Enabled
802.1p Marking Action	Map	Map

VoIP traffic (as defined by the Service Group) arriving from **Remote Site 1 Subnets** across the VPN destined to **LAN Subnets** on the LAN zone at the Main Site would hit the Access Rule for inbound VoIP calls. Traffic arriving at the VPN zone will not have any 802.1p tags, only DSCP tags.

- Traffic exiting the tunnel containing a DSCP tag (for example, CoS = 48) would have the DSCP value preserved. Before the packet is delivered to the destination on the LAN, it will also be 802.1p tagged according to the **QoS Mapping** settings (for example, CoS = 6) by the firewall at the Main Site.
- Assuming returned traffic has been 802.1p tagged (for example, CoS = 6) by the VoIP phone receiving the call at the Main Site, the return traffic will be DSCP tagged according to the conversion map (CoS = 48) on both the inner and outer packet sent back across the VPN.
- Assuming returned traffic has been DSCP tagged (for example, CoS = 48) by the VoIP phone receiving the call at the Main Site, the return traffic will have the DSCP tag preserved on both the inner and outer packet sent back across the VPN.
- Assuming returned traffic has been both 802.1p tagged (for example, CoS = 6) and DSCP tagged (for example, CoS = 14) by the VoIP phone receiving the call at the Main Site, the return traffic will be DSCP tagged according to the conversion map (CoS = 48) on both the inner and outer packet sent back across the VPN.

Bandwidth Management

For information on Bandwidth Management (BWM), see Firewall Settings > BWM on page 25.

Glossary

- **802.1p** – IEEE 802.1p is a Layer 2 (MAC layer) Class of Service mechanism that tags packets by using 3 priority bits (for a total of 8 priority levels) within the additional 16-bits of an 802.1q header. 802.1p processing requires compatible equipment for tag generation, recognition and processing, and should only be employed on compatible networks.
- **Bandwidth Management (BWM)** – Refers to any of a variety of algorithms or methods used to shape traffic or police traffic. Shaping often refers to the management of outbound traffic, while policing often refers to the management of inbound traffic (also known as admission control). There are many different methods of bandwidth management, including various queuing and discarding techniques, each with their own design strengths. SonicWall employs a Token Based Class Based Queuing method for inbound and outbound BWM, as well as a discard mechanism for certain types of inbound traffic.
- **Class of Service (CoS)** – A designator or identifier, such as a layer 2 or layer 3 tag, that is applied to traffic after classification. CoS information will be used by the Quality of Service (QoS) system to differentiate between the classes of traffic on the network, and to provide special handling (for example, prioritized queuing, low latency) as defined by the QoS system administrator.
- **Classification** – The act of identifying (or differentiating) certain types (or classes) of traffic. Within the context of QoS, this is performed for the sake of providing customized handling, typically prioritization or de-prioritization, based on the traffic's sensitivity to delay, latency, or packet loss. Classification within SonicOS uses Access Rules, and can occur based on any or all of the following elements: source zone, destination zone, source address object, destination address object, service object, schedule object.

- **Code Point** – A value that is marked (or tagged) into the DSCP portion of an IP packet by a host or by an intermediate network device. There are currently 64 Code Points available, from 0 to 63, used to define the ascending prioritized class of the tagged traffic.
- **Conditioning** – A broad term used to describe a plurality of methods of providing Quality of Service to network traffic, including but not limited to discarding, queuing, policing, and shaping.
- **DiffServ** (Differentiated Services) – A standard for differentiating between different types or classes of traffic on an IP network for the purpose of providing tailored handling to the traffic based on its requirements. DiffServ primarily depends upon Code Point values marked in the ToS header of an IP packet to differentiate between different classes of traffic. DiffServ service levels are executed on a Per Hop Basis at each router (or other DiffServ enabled network device) through which the marked traffic passes. DiffServ Service levels currently include at a minimum **Default**, **Assured Forwarding, Expedited Forwarding,** and **DiffServ**. Refer to DSCP Marking on page 83 for more information.
- **Discarding** – A congestion avoidance mechanism that is employed by QoS systems in an attempt to predict when congestion might occur on a network, and to prevent the congestion by dropping over-limit traffic. Discarding can also be thought of as a queue management algorithm, since it attempts to avoid situations of full queues. Advanced discard mechanisms will abide by CoS markings so as to avoid dropping sensitive traffic. Common methods are:
    - **Tail Drop** – An indiscriminate method of dealing with a full queue wherein the last packets into the queue are dropped, regardless of their CoS marking.
    - **Random Early Detection (RED)** – RED monitors the status of queues to try to anticipate when a queue is about to become full. It then randomly discards packets in a staggered fashion to help minimize the potential of Global Synchronization. Basic implementations of RED, like Tail Drop, do not consider CoS markings.
    - **Weighted Random Early Detection (WRED)** – An implementation of RED that factors DSCP markings into its discard decision process.
- **DSCP** (Differentiate Services Code Points) – The repurposing of the ToS field of an IP header as described by RFC2747. DSCP uses 64 Code Point values to enable DiffServ (Differentiated Services). By marking traffic according to its class, each packet can be treated appropriately at every hop along the network.
- **Global Synchronization** – A potential side effect of discarding, the congestion avoidance method designed to deal with full queues. Global Synchronization occurs when multiple TCP flows through a congested link are dropped at the same time (as can occur in Tail Drop). When the native TCP slow-start mechanism commences with near simultaneity for each of these flows, the flows will again flood the link. This leads to cyclical waves of congestion and under-utilization.
- **Guaranteed Bandwidth** – A declared percentage of the total available bandwidth on an interface which will always be granted to a certain class of traffic. Applicable to both inbound and outbound BWM. The total Guaranteed Bandwidth across all BWM rules cannot exceed 100% of the total available bandwidth. SonicOS enhances the Bandwidth Management feature to provide rate limiting functionality. You can now create traffic policies that specify maximum rates for Layer 2, 3, or 4 network traffic. This enables bandwidth management in cases where the primary WAN link fails over to a secondary connection that cannot handle as much traffic. The Guaranteed Bandwidth can also be set to 0%.
- **Inbound (Ingress or IBWM)** – The ability to shape the rate at which traffic enters a particular interface. For TCP traffic, actual shaping can occur where the rate of the ingress flow can be adjusted by delaying egress acknowledgments (ACKs) causing the sender to slow its rate. For UDP traffic, a discard mechanism is used since UDP has no native feedback controls.
- **IntServ** (Integrated Services) – As defined by RFC1633. An alternative CoS system to DiffServ, IntServ differs fundamentally from DiffServ in that it has each device request (or reserve) its network requirements before it sends its traffic. This requires that each hop on the network be IntServ aware, and it also requires each hop to maintain state information for every flow. IntServ is not supported by SonicOS. The most common implementation of IntServ is RSVP.

- **Maximum Bandwidth** – A declared percentage of the total available bandwidth on an interface defining the maximum bandwidth to be allowed to a certain class of traffic. Applicable to both inbound and outbound BWM. Used as a throttling mechanism to specify a bandwidth rate limit. The Bandwidth Management feature is enhanced to provide rate limiting functionality. You can now create traffic policies that specify maximum rates for Layer 2, 3, or 4 network traffic. This enables bandwidth management in cases where the primary WAN link fails over to a secondary connection that cannot handle as much traffic.The Maximum Bandwidth can be set to 0%, which will prevent all traffic.
- **Outbound (Egress or OBWM)** – Conditioning the rate at which traffic is sent out an interface. Outbound BWM uses a credit (or token) based queuing system with 8 priority rings to service different types of traffic, as classified by Access Rules.
- **Priority** – An additional dimension used in the classification of traffic. SonicOS uses 8 priority rings (0 = highest, 7 = lowest) to comprise the queue structure used for BWM. Queues are serviced in the order of their priority ring.
- **Mapping** – With regard to SonicOS's implementation of QoS, mapping is the practice of converting layer 2 CoS tags (802.1p) to layer 3 CoS tags (DSCP) and back again for preserving the 802.1p tags across network links that do not support 802.1p tagging. The map correspondence is fully user-definable, and the act of mapping is controlled by Access Rules.
- **Marking** – Also known as **tagging** or **coloring** – The act of applying layer 2 (802.1p) or layer 3 (DSCP) information to a packet for the purpose of differentiation, so that it can be properly classified (recognized) and prioritized by network devices along the path to its destination.
- **MPLS** (Multi Protocol Label Switching) – A term that comes up frequently in the area of QoS, but which is natively unsupported by most customer premise IP networking devices, including SonicWall appliances. MPLS is a carrier-class network service that attempts to enhance the IP network experience by adding the concept connection-oriented paths (Label Switch Paths – LSPs) along the network. When a packet leaves a customer premise network, it is tagged by a Label Edge Router (LER) so that the label can be used to determine the LSP. The MPLS tag itself resides between layer 2 and layer 3, imparting upon MPLS characteristics of both network layers. MPLS is becoming quite popular for VPNs, offering both layer 2 and layer 3 VPN services, but remains interoperable with existing IPsec VPN implementation. MPLS is also very well known for its QoS capabilities, and interoperates well with conventional DSCP marking.
- **Per Hop Behavior (PHB)** – The handling that will be applied to a packet by each DiffServ capable router it traverses, based upon the DSCP classification of the packet. The behavior can be among such actions as discard, re-mark (re-classify), best-effort, assured forwarding, or expedited forwarding.
- **Policing** – A facility of traffic conditioning that attempts to control the rate of traffic into or out of a network link. Policing methods range from indiscriminate packet discarding to algorithmic shaping, to various queuing disciplines.
- **Queuing** – To effectively make use of a link's available bandwidth, queues are commonly employed to sort and separately manage traffic after it has been classified. Queues are then managed using a variety of methods and algorithms to ensure that the higher priority queues always have room to receive more traffic, and that they can be serviced (de-queued or processed) before lower priority queues. Some common queue disciplines include:
    - **FIFO** (First In First Out) – A very simple, undiscriminating queue where the first packet in is the first packet to be processed.
    - **Class Based Queuing (CBQ)** – A queuing discipline that takes into account the CoS of a packet, ensuring that higher priority traffic is treated preferentially.
    - **Weighted Fair Queuing (WFQ)** – A discipline that attempts to service queues using a simple formula based upon the packets' IP precedence and the total number of flows. WFQ has a tendency to become imbalanced when there is a disproportionately large number of high-priority flows to be serviced, often having the opposite of the desired effect.

    - **Token Based CBQ** – An enhancement to CBQ that employs a token, or a credit-based system that helps to smooth or normalize link utilization, avoiding burstiness as well as under-utilization. Employed by SonicOS BWM.

- **RSVP** (Resource Reservation Protocol) – An IntServ signaling protocol employed by some applications where the anticipated need for network behavior (for example, delay and bandwidth) is requested so that it can be reserved along the network path. Setting up this Reservation Path requires that each hop along the way be RSVP capable, and that each agrees to reserve the requested resources. This system of QoS is comparatively resource intensive, since it requires each hop to maintain state on existing flows. Although IntServ's RSVP is quite different from DiffServ's DSCP, the two can interoperate. RSVP is not supported by SonicOS.
- **Shaping** – An attempt by a QoS system to modify the rate of traffic flow, usually by employing some feedback mechanism to the sender. The most common example of this is TCP rate manipulation, where acknowledgements (ACKs) sent back to a TCP sender are queued and delayed so as to increase the calculated round-trip time (RTT), leveraging the inherent behavior of TCP to force the sender to slow the rate at which it sends data.
- **Type of Service (ToS)** – A field within the IP header wherein CoS information can be specified. Historically used, albeit somewhat rarely, in conjunction with IP precedence bits to define CoS. The ToS field is now rather commonly used by DiffServ's code point values.

6

Configuring SSL Control

- About SSL Control on page 94
- Firewall Settings > SSL Control on page 102
- SSL Control Configuration on page 102
- Enabling SSL Control on Zones on page 106
- SSL Control Events on page 106

About SSL Control

SonicOS includes SSL Control, a system for providing visibility into the handshake of SSL sessions and a method for constructing policies to control the establishment of SSL connections. SSL (Secure Sockets Layer) is the dominant standard for the encryption of TCP-based network communications, with its most common and well-known application being HTTPS (HTTP over SSL); see HTTP over SSL communication. SSL provides digital certificate-based endpoint identification, and cryptographic and digest-based confidentiality to network communications.

HTTP over SSL communication

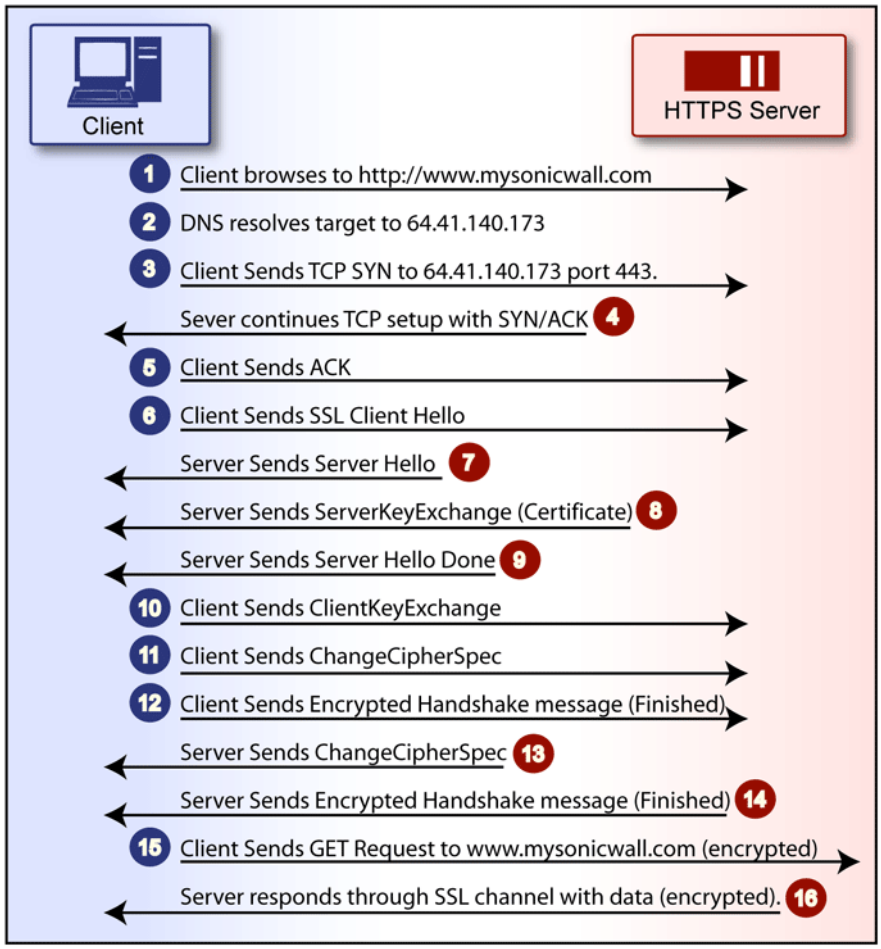

An effect of the security provided by SSL is the obscuration of all payload, including the URL (Uniform Resource Locator, for example, https://www.mysonicwall.com) being requested by a client when establishing an HTTPS session. This is due to the fact that HTTP is transported within the encrypted SSL tunnel when using HTTPS. It is not until the SSL session is established (see HTTP over SSL communication) that the actual target resource (`www.mysonicwall.com`) is requested by the client, but as the SSL session is already established, no inspection of the session data by the firewall or any other intermediate device is possible. As a result, URL-based content filtering systems cannot consider the request to determine permissibility in any way other than by IP address.

While IP address based filtering does not work well for unencrypted HTTP because of the efficiency and popularity of host-header-based virtual hosting (defined in Key Concepts to SSL Control on page 97), IP filtering can work effectively for HTTPS due to the rarity of host-header-based HTTPS sites. But this trust relies on the integrity of the HTTPS server operator, and assumes that SSL is not being used for deceptive purposes.

For the most part, SSL is employed legitimately, being used to secure sensitive communications, such as online shopping or banking, or any session where there is an exchange of personal or valuable information. The ever decreasing cost and complexity of SSL, however, has also spurred the growth of more dubious applications of SSL, designed primarily for the purposes of obfuscation or concealment rather than security.

An increasingly common camouflage is the use of SSL encrypted Web-based proxy servers for the purpose of hiding browsing details, and bypassing content filters. While it is simple to block well known HTTPS proxy services of this sort by their IP address, it is virtually impossible to block the thousands of privately-hosted proxy servers that are readily available through a simple Web-search. The challenge is not the ever-increasing number of such services, but rather their unpredictable nature. Since these services are often hosted on home networks

using dynamically addressed DSL and cable modem connections, the targets are constantly moving. Trying to block an unknown SSL target would require blocking all SSL traffic, which is practically infeasible.

SSL Control provides a number of methods to address this challenge by arming the security administrator with the ability to dissect and apply policy based controls to SSL session establishment. While the current implementation does not decode the SSL application data, it does allow for gateway-based identification and disallowance of suspicious SSL traffic.

Topics:

- Key Features of SSL Control on page 96
- Key Concepts to SSL Control on page 97
- Caveats and Advisories on page 101

Key Features of SSL Control

SSL control: Features and benefits

Feature	Benefit
Common Name-based White and Black Lists	You can define lists of explicitly allowed or denied certificate subject common names (described in Key Concepts). Entries are matched on substrings, for example, a blacklist entry for `prox` will match `www.megaproxy.com`, `www.proxify.com` and "`roxify.net`. This allows you to easily block all SSL exchanges employing certificates issued to subjects with potentially objectionable names. Inversely, you can easily authorize all certificates within an organization by whitelisting a common substring for the organization. Each list can contain up to 1,024 entries. As the evaluation is performed on the subject common name embedded in the certificate, even if the client attempts to conceal access to these sites by using an alternative hostname or even an IP address, the subject is always detected in the certificate, and policy is applied.
Self-Signed Certificate Control	It is common practice for legitimate sites secured by SSL to use certificates issued by well-known certificate authorities, as this is the foundation of trust within SSL. It is almost equally common for network appliances secured by SSL (such as SonicWall network security appliances) to use self-signed certificates for their default method of security. So while self-signed certificates in closed environments are not suspicious, the use of self-signed certificates by publicly or commercially available sites is. A public site using a self-signed certificate is often an indication that SSL is being used strictly for encryption rather than for trust and identification. While not absolutely incriminating, this sometimes suggests that concealment is the goal, as is commonly the case for SSL encrypted proxy sites. The ability to set a policy to block self-signed certificates allows you to protect against this potential exposure. To prevent discontinuity of communications to known/trusted SSL sites using self-signed certificates, the whitelist feature can be used for explicit allowance.

SSL control: Features and benefits

Feature	Benefit
Untrusted Certificate Authority Control	Like the use of self-signed certificates, encountering a certificate issued by an untrusted CA is not an absolute indication of disreputable obscuration, but it does suggest questionable trust. SSL Control can compare the issuer of the certificate in SSL exchanges against the certificates in the firewall's certificate store. The certificate store contains approximately 100 well-known CA certificates, exactly like today's Web-browsers. If SSL Control encounters a certificate that was issued by a CA not in its certificate store, it can disallow the SSL connection. For organizations running their own private certificate authorities, the private CA certificate can easily be imported into the firewall's certificate store to recognize the private CA as trusted. The store can hold up to 256 certificates.
SSL version, Cipher Strength, and Certificate Validity Control	SSL Control provides additional management of SSL sessions based on characteristics of the negotiation, including the ability to disallow the potentially exploitable SSLv2, the ability to disallow weak encryption (ciphers less than 64 bits), and the ability to disallow SSL negotiations where a certificate's date ranges are invalid. This enables the administrator to create a rigidly secure environment for network users, eliminating exposure to risk through unseen cryptographic weaknesses, or through disregard for or misunderstanding of security warnings.
Zone-Based Application	SSL Control is applied at the zone level, allowing you to enforce SSL policy on the network. When SSL Control is enabled on the zone, the firewall looks for Client Hellos sent from clients on that zone through the firewall, which triggers inspection. The firewall looks for the Server Hello and Certificate that is sent in response for evaluation against the configured policy. Enabling SSL Control on the LAN zone, for example, inspects all SSL traffic initiated by clients on the LAN to any destination zone.
Configurable Actions and Event Notifications	When SSL Control detects a policy violation, it can log the event and block the connection, or it can simply log the event while allowing the connection to proceed.

Key Concepts to SSL Control

- **SSL**- Secure Sockets Layer (SSL) is a network security mechanism introduced by Netscape in 1995. SSL was designed to provide privacy between two communicating applications (a client and a server) and also to authenticate the server, and optionally the client. SSL's most popular application is HTTPS, designated by a URL beginning with `https://` rather than simply `http://`, and it is recognized as the standard method of encrypting Web traffic on the Internet. An SSL HTTP transfer typically uses TCP port 443, whereas a regular HTTP transfer uses TCP port 80. Although HTTPS is what SSL is best known for, SSL is not limited to securing HTTP, but can also be used to secure other TCP protocols such as SMTP, POP3, IMAP, and LDAP. SSL session establishment occurs as shown in Establishing an SSL session:

Establishing an SSL session

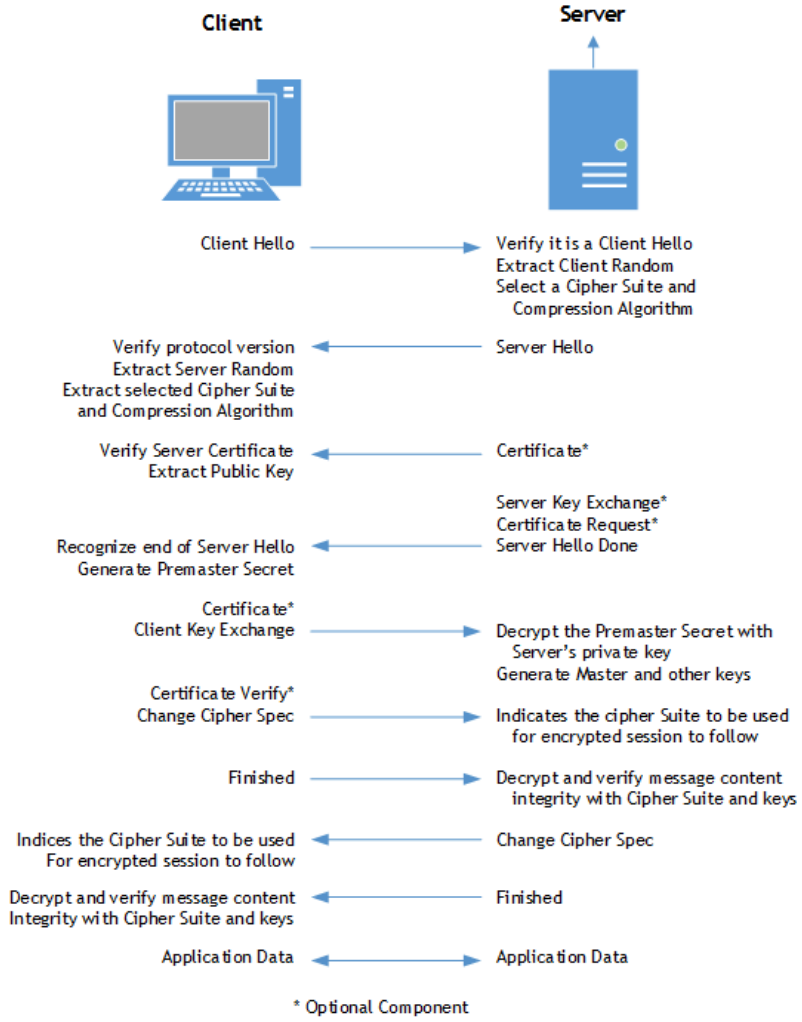

- **SSLv2** – The earliest version of SSL still in common use. SSLv2 was found to have a number of weaknesses, limitations, and theoretical deficiencies (comparatively noted in the SSLv3 entry), and is looked upon with scorn, disdain, and righteous indignation by security purists.
- **SSLv3** – SSLv3 was designed to maintain backward compatibility with SSLv2, while adding the following enhancements:
  - Alternate key exchange methods, including Diffie-Hellman.
  - Hardware token support for both key exchange and bulk encryption.
  - SHA, DSS, and Fortezza support.
  - Out-of-Band data transfer.
  - TLS – Transport Layer Security, also known as SSLv3.1, is very similar to SSLv3, but improves upon SSLv3 in the ways shown in Differences between SSL and TLS:

Differences between SSL and TLS

SSL	TLS
Uses a preliminary HMAC algorithm	Uses HMAC as described in RFC 2104
Does not apply MAC to version info	Applies MAC to version info

Differences between SSL and TLS

SSL	TLS
Does not specify a padding value	Initializes padding to a specific value
Limited set of alerts and warning	Detailed Alert and Warning messages

NOTE: SonicOS 6.2.2.1 and above support TLS 1.1 and 1.2.

- **MAC** – A MAC (Message Authentication Code) is calculated by applying an algorithm (such as MD5 or SHA1) to data. The MAC is a message digest, or a one-way hash code that is fairly easy to compute, but which is virtually irreversible. In other words, with the MAC alone, it would be theoretically impossible to determine the message upon which the digest was based. It is equally difficult to find two different messages that would result in the same MAC. If the receiver's MAC calculation matches the sender's MAC calculation on a given piece of data, the receiver is assured that the data has not been altered in transit.
- **Client Hello** – The first message sent by the client to the server following TCP session establishment. This message starts the SSL session, and consists of the following components:
  - **Version** – The version of SSL that the client wishes to use in communications. This is usually the most recent version of SSL supported by the client.
  - **Random** – A 32-bit timestamp coupled with a 28-byte random structure.
  - **Session ID** – This can either be empty if no Session ID data exists (essentially requesting a new session) or can reference a previously issued Session ID.
  - **Cipher Suites** – A list of the cryptographic algorithms, in preferential order, supported by the clients.
  - **Compression Methods** – A list of the compression methods supported by the client (typically null).
- **Server Hello** – The SSL server's response to the Client Hello. It is this portion of the SSL exchange that SSL Control inspects. The Server Hello contains the version of SSL negotiated in the session, along with cipher, session ID and certificate information. The actual X.509 server certificate itself, although a separate step of the SSL exchange, usually begins (and often ends) in the same packet as the Server Hello.
- **Certificates** - X.509 certificates are unalterable digital stamps of approval for electronic security. There are four main characteristics of certificates:
  - Identify the subject of a certificate by a common name or distinguished name (CN or DN).
  - Contain the public key that can be used to encrypt and decrypt messages between parties
  - Provide a digital signature from the trusted organization (Certificate Authority) that issued the certificate.
  - Indicate the valid date range of the certificate
- **Subject** – The guarantee of a certificate identified by a common name (CN). When a client browses to an SSL site, such as https://www.mysonicwall.com, the server sends its certificate which is then evaluated by the client. The client checks that the certificate's dates are valid, that is was issued by a trusted CA, and that the subject CN matches the requested host name (that is, they are both `www.mysonicwall.com`). Although a subject CN mismatch elicits a browser alert, it is not always a sure sign of deception. For example, if a client browses to `https://mysonicwall.com`, which resolves to the same IP address as `www.mysonicwall.com`, the server presents its certificate bearing the subject CN of `www.mysonicwall.com`. An alert will be presented to the client, despite the total legitimacy of the connection.

- **Certificate Authority (CA**) - A Certificate Authority (CA) is a trusted entity that has the ability to sign certificates intended, primarily, to validate the identity of the certificate's subject. Well-known certificate authorities include VeriSign, Thawte, Equifax, and Digital Signature Trust. In general, for a CA to be trusted within the SSL framework, its certificate must be stored within a trusted store, such as that employed by most Web-browsers, operating systems and run-time environments. The SonicOS trusted store is accessible from the **MANAGE | System Setup > Appliance > Certificates** page. The CA model is built on associative trust, where the client trusts a CA (by having the CAs certificate in its trusted store), the CA trusts a subject (by having issued the subject a certificate), and therefore the client can trust the subject.
- **Untrusted CA** – An untrusted CA is a CA that is not contained in the trusted store of the client. In the case of SSL Control, an untrusted CA is any CA whose certificate is not present in **MANAGE | System Setup > Appliance > Certificates**.
- **Self-Signed Certificates** – Any certificate where the issuer's common-name and the subject's common-name are the same, indicating that the certificate was self-signed.
- **Virtual Hosting** – A method employed by Web servers to host more than one website on a single server. A common implementation of virtual hosting is name-based (Host-header) virtual hosting, which allows for a single IP address to host multiple websites. With Host-header virtual hosting, the server determines the requested site by evaluating the "Host:" header sent by the client. For example, both `www.website1.com` and `www.website2.com` might resolve to `64.41.140.173`. If the client sends a "`GET /`" along with "`Host: www.website1.com`", the server can return content corresponding to that site.

  Host-header virtual hosting is generally not employed in HTTPS because the host header cannot be read until the SSL connection is established, but the SSL connection cannot be established until the server sends its Certificate. Since the server cannot determine which site the client will request (all that is known during the SSL handshake is the IP address) it cannot determine the appropriate certificate to send. While sending any certificate might allow the SSL handshake to commence, a certificate name (subject) mismatch will trigger a browser alert.
- **Weak Ciphers** – Relatively weak symmetric cryptography ciphers. Ciphers are classified as weak when they are less than 64 bits. For the most part, export ciphers are weak ciphers. Common weak ciphers lists common weak ciphers:

Common weak ciphers

Cipher	Encryption	Occurs in
EXP1024-DHE-DSS-DES-CBC-SHA	DES(56)	SSLv3, TLS (export)
EXP1024-DHE-CBC-SHA	DES(56)	SSLv3, TLS (export)
EXP1024-RC2-CBC-MD5	RC2(56)	SSLv3, TLS (export)
EDH-RSA-DES-CBC-SHA	DES(56)	SSLv3, TLS
EDH-DSS-DES-CBC-SHA	DES(56)	SSLv3, TLS
DES-CBC-SHA	DES(56)	SSLv2, SSLv3, TLS
EXP1024-DHE-DSS-RC4-SHA	RC4(56)	SSLv3, TLS (export)
EXP1024-RC4-SHA	RC4(56)	SSLv3, TLS (export)
EXP1024-RC4-MD5	RC4(56)	SSLv3, TLS (export)
EXP-EDH-RSA-DES-CBC-SHA	DES(40)	SSLv3, TLS (export)
EXP-EDH-DSS-DES-CBC-SHA	DES(40)	SSLv3, TLS (export)
EXP-DES-CBC-SHA	DES(40)	SSLv3, TLS (export)
EXP-RC2-CBC-MD5	RC2(40)	SSLv2, SSLv3, TLS (export)
EXP-RC4-MD5	RC4(40)	SSLv2, SSLv3, TLS (export)

Caveats and Advisories

1. **Self-signed and Untrusted CA enforcement** – If enforcing either of these two options, it is strongly advised that you add the common names of any SSL secured network appliances within your organization to the whitelist to ensure that connectivity to these devices is not interrupted. For example, the default subject name of a SonicWall network security appliances is `192.168.168.168`, and the default common name of SonicWall SSL VPN appliances is `192.168.200.1`.
2. If your organization employs its own private Certificate Authority (CA), it is strongly advised that you import your private CAs certificate into the **System > Certificates** store, particularly if you will be enforcing blocking of certificates issued by untrusted CAs. Refer to *SonicWall SonicOS 6.5 System Setup* for more information on this process.
3. SSL Control inspection is currently only performed on TCP port 443 traffic. SSL negotiations occurring on non-standard ports will not be inspected at this time.
4. **Server Hello fragmentation** – In some rare instances, an SSL server fragments the Server Hello. If this occurs, the current implementation of SSL Control does not decode the Server Hello. SSL Control policies are not applied to the SSL session, and the SSL session is allowed.
5. **Session termination handling** – When SSL Control detects a policy violation and terminates an SSL session, it simply terminates the session at the TCP layer. Because the SSL session is in an embryonic state at this point, it is not currently possible to redirect the client or to provide any kind of informational notification of termination to the client.
6. **Whitelist precedence** – The whitelist takes precedence over all other SSL Control elements. Any SSL server certificate which matches an entry in the whitelist will allow the SSL session to proceed, even if other elements of the SSL session are in violation of the configured policy. This is by design.
7. The number of pre-installed (well-known) CA certificates is 93. The resulting repository is very similar to what can be found in most Web-browsers. Other certificate related changes:
    a. The maximum number of CA certificates was raised from 6 to 256.
    b. The maximum size of an individual certificate was raised from 2,048 to 4,096.
    c. The maximum number of entries in the whitelist and blacklist is 1,024 each.

Firewall Settings > SSL Control

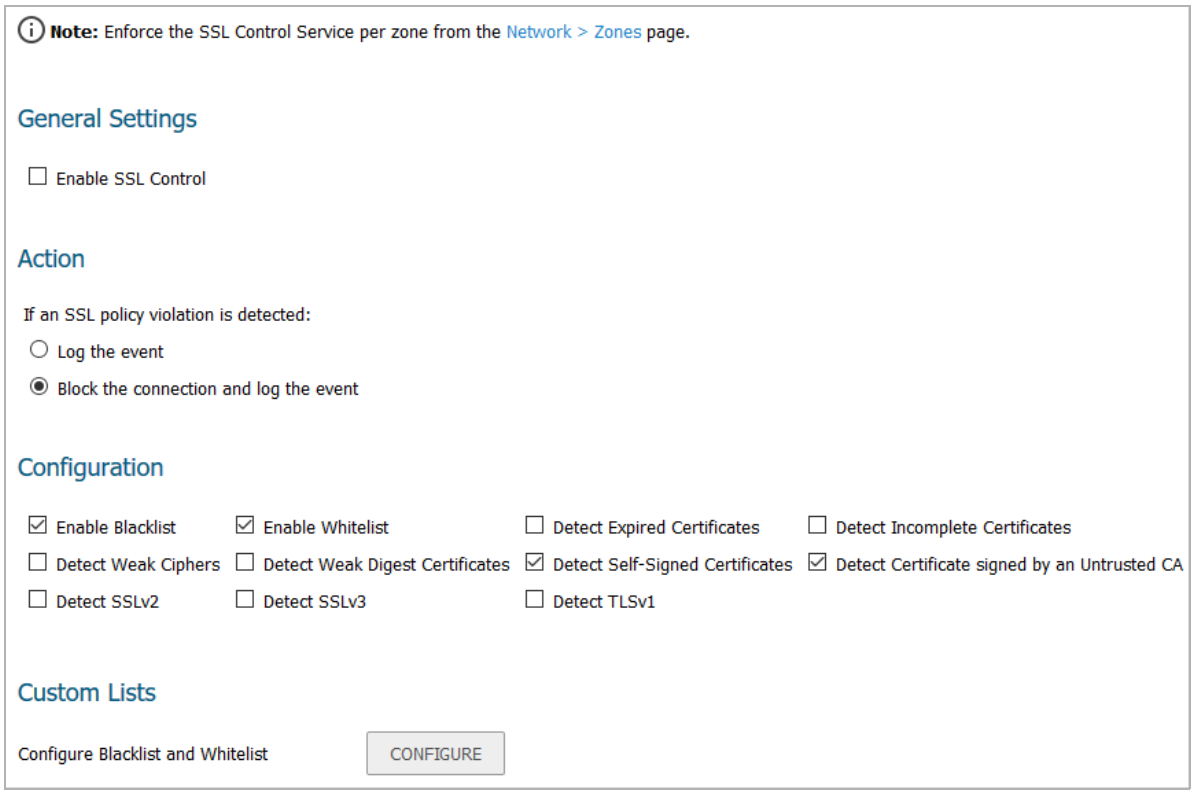

SSL Control Configuration

NOTE: Before configuring SSL Control, ensure your firewall supports IPv6. You can confirm this by using the **IPv6 Check Network Settings** tool on the **System > Diagnostics** page; see *SonicWall SonicOS 6.5 Investigate*.

SSL Control is located on the **MANAGE** view, under **Security Configuration > Firewall Settings > SSL Control**. SSL Control has a global setting, as well as a per-zone setting. By default, SSL Control is not enabled at the global or

zone level. The individual page controls are as follows (refer Key Concepts to SSL Control on page 97 for more information on terms used in this section).

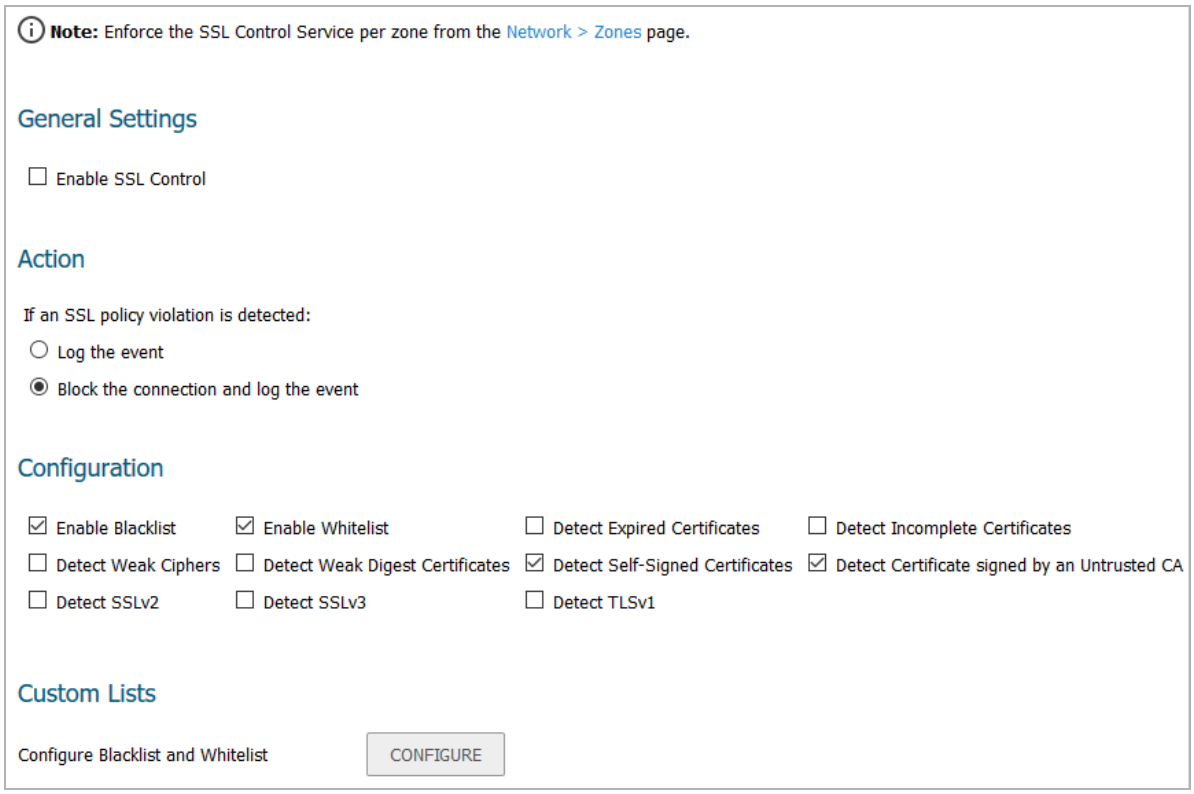

Topics:

- General Settings on page 103
- Action on page 103
- Configuration on page 104
- Custom Lists on page 105

General Settings

The **General Settings** section allows you to enable or disable SSL control:

- **Enable SSL Control** – The global setting for SSL Control. This must be enabled for SSL Control applied to zones to be effective. This option is not selected by default.

Action

The **Action** section is where you choose the action to be taken when an SSL policy violation is detected; either:

- **Log the event** – If an SSL policy violation, as defined within the **Configuration** section below, is detected, the event is logged, but the SSL connection is allowed to continue. This option is not selected by default.
- **Block the connection and log the event** – In the event of a policy violation, the connection is blocked and the event is logged. This option is selected by default.

Configuration

The **Configuration** section is where you specify the SSL policies to be enforced:

- **Enable Blacklist** – Controls detection of the entries in the blacklist, as configured in Custom Lists. This option is selected by default.
- **Enable Whitelist** – Controls detection of the entries in the whitelist, as configured in the **Configure Lists** section below. Whitelisted entries take precedence over all other SSL control settings. This option is selected by default.
- **Detect Weak Ciphers** – Controls the detection of SSL sessions negotiated with symmetric ciphers less than 64 bits, commonly indicating export cipher usage. This option is not selected by default.
- **Detect Expired Certificates** – Controls detection of certificates whose start date is before the current system time, or whose end date is beyond the current system time. Date validation depends on the firewall's System Time. Make sure your System Time is set correctly, preferably synchronized with NTP, on the **MANAGE | System Setup > System > Time** page. This option is not selected by default.
- **Detect Weak Digest Certificates** – Controls detection of certificates created using MD5 or SHA1. Both MD5 or SHA1 are not considered safe. This option is not selected by default.
- **Detect Self-Signed Certificates** – Controls the detection of certificates where both the issuer and the subject have the same common name. This option is selected by default.

  It is common practice for legitimate sites secured by SSL to use certificates issued by well-known certificate authorities, as this is the foundation of trust within SSL. It is almost equally common for network appliances secured by SSL (such as SonicWall security appliances) to use self-signed certificates for their default method of security. So while self-signed certificates in closed-environments are not suspicious, the use of self-signed certificates by publicly or commercially available sites is. A public site using a self-signed certificate is often an indication that SSL is being used strictly for encryption rather than for trust and identification. While not absolutely incriminating, this sometimes suggests that concealment is the goal, as is commonly the case for SSL encrypted proxy sites. The ability to set a policy to block self-signed certificates allows you to protect against this potential exposure. To prevent discontinuity of communications to known/trusted SSL sites using self-signed certificates, use the whitelist feature for explicit allowance.
- **Detect Certificates signed by an Untrusted CA** – Controls the detection of certificates where the issuer's certificate is not in the firewall's **MANAGE | System Setup > Appliance > Certificates** trusted store. This option is selected by default.

  Similar to the use of self-signed certificates, encountering a certificate issued by an untrusted CA is not an absolute indication of disreputable obscuration, but it does suggest questionable trust. SSL Control can compare the issuer of the certificate in SSL exchanges against the certificates stored in the SonicWall firewall where most of the well-known CA certificates are included. For organizations running their own private certificate authorities, the private CA certificate can easily be imported into the SonicWall's whitelist to recognize the private CA as trusted
- **Detect SSLv2** – Controls detection and blocking of SSLv2 exchanges. SSLv2 is known to be susceptible to cipher downgrade attacks because it does not perform integrity checking on the handshake. Best practices recommend using SSLv3 or TLS in its place. This option is selected by default. It is also dimmed and cannot be changed.
- **Detect SSLv3** – Controls detection and blocking of SSLv3 exchanges. This option is not selected by default.
- **Detect TLSv1** – Controls the detection and blocking of TLSv1 exchanges. This option is not selected by default.

Custom Lists

The **Custom Lists** section allows you to configure custom whitelists and blacklists.

- **Configure Blacklist and Whitelist** – Allows you to define strings for matching common names in SSL certificates. Entries are case-insensitive and are used in pattern-matching fashion, as shown in Blacklist and Whitelist: pattern matching:

Blacklist and Whitelist: pattern matching

Entry	Will Match	Will Not Match
sonicwall.com	https://www.sonicwall.com, https://csm.demo.sonicwall.com, https://mysonicwall.com, https://supersonicwall.computers.org, https://67.115.118.87 [1]	https://www.sonicwall.de
prox	https://proxify.org, https://www.proxify.org, https://megaproxy.com, https://1070652204 [2]	https://www.freeproxy.ru [3]

1. `67.115.118.67` is currently the IP address to which sslvpn.demo.sonicwall.com resolves, and that site uses a certificate issued to `sslvpn.demo.sonicwall.com`. This results in a match to "`sonicwall.com`" as matching occurs based on the common name in the certificate.
2. This is the decimal notation for the IP address `63.208.219.44`, whose certificate is issued to `www.megaproxy.com`.
3. `www.freeproxy.ru` will not match "`prox`" as the common name on the certificate that is currently presented by this site is a self-signed certificate issued to "`-`". This can, however, easily be blocked by enabling control of self-signed or Untrusted CA certificates.

To configure the Whitelist and Blacklist:

1 Navigate to the **MANAGE | Security Configuration > Firewall Settings > SSL Control** page.

2 Click **CONFIGURE**. The **SSL Control Custom Lists** dialog displays.

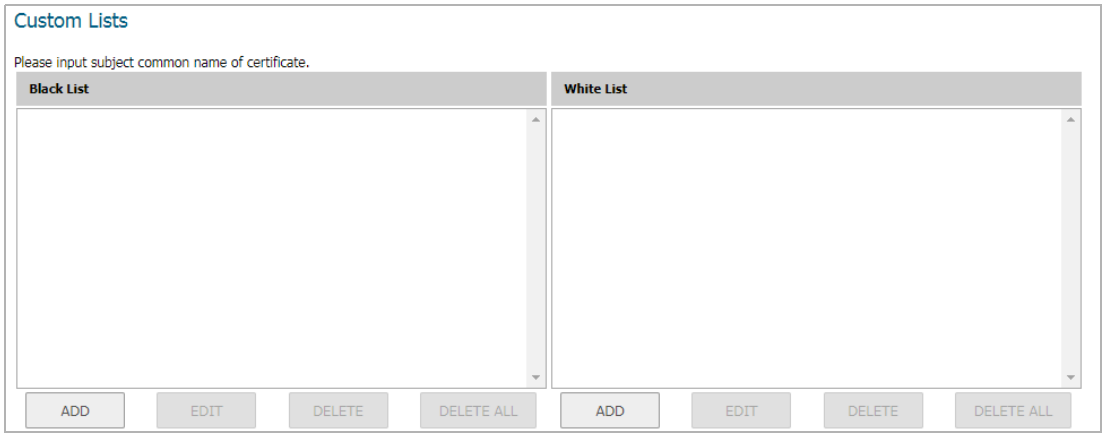

3 To add a certificate to either the Black List or White List table, click the appropriate **ADD**. The **Add Blacklist/Whitelist Domain Entry** dialog displays.

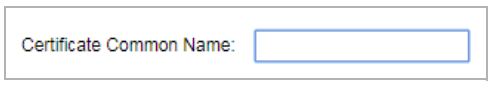

4 Enter the certificate's name in the **Certificate Common Name** field.

TIP: List matching is based on the subject common name in the certificate presented in the SSL exchange, not in the URL (resource) requested by the client.

You can edit and delete certificates with the buttons beneath each list table.

5 Click **OK**.

Changes to any of the SSL Control settings do not affect currently established connections; only new SSL exchanges that occur after the change is committed are inspected and affected.

6 Click **OK**.

7 Click **ACCEPT**.

Enabling SSL Control on Zones

After SSL Control has been globally enabled, and the desired options have been configured, SSL Control must be enabled on one or more zones. When SSL Control is enabled on the zone, the firewall looks for Client Hellos sent from clients on that zone through the firewall will trigger inspection. The firewall then looks for the Server Hello and Certificate that is sent in response for evaluation against the configured policy. Enabling SSL Control on the LAN zone, for example, will inspect all SSL traffic initiated by clients on the LAN to any destination zone.

NOTE: If you are activating SSL Control on a zone (for example, the LAN zone) where there are clients who will be accessing an SSL server on another zone connected to the firewall (for example, the DMZ zone), it is recommended that you add the subject common name of that server's certificate to the whitelist to ensure continuous trusted access.

To enable SSL Control on a zone:

1 Navigate to the **MANAGE | System Setup > Network > Zones** page.

2 Select the **Configure** icon for the desired zone. The **Edit Zone** dialog displays.

3 Select the **Enable SSL Control** option. For configuring the rest of the options on the Edit Zone dialog, see *SonicWall SonicOS 6.5 System Setup*.

4 Click **OK**. All new SSL connections initiated from that zone are now subject to inspection.

SSL Control Events

Log events include the client's username in the notes section (not shown) if the user logged in manually or was identified through CIA/Single Sign On. If the user's identity is not available, the note indicates the user is `Unidentified`.

SSL control: Event messages

#	Event Message	Conditions When it Occurs
1	SSL Control: Certificate with Invalid date	The certificate's start date is either before the SonicWall's system time or it's end date is after the system time.
2	SSL Control: Certificate chain not complete	The certificate has been issued by an intermediate CA with a trusted top-level CA, but the SSL server did not present the intermediate certificate. This log event is informational and does not affe3ct the SSL connection.

SSL control: Event messages

#	Event Message	Conditions When it Occurs
3	SSL Control: Self-signed certificate	The certificate is self-signed (the CN of the issuer and the subject match). **NOTE:** For information about enforcing self-signed certificate controls, see Caveats and Advisories on page 101.
4	SSL Control: Untrusted CA	The certificate has been issued by a CA that is not in the **System > Certificates** store of the firewall. **NOTE:** For information about enforcing self-signed certificate controls, see Caveats and Advisories on page 101.
5	SSL Control: Website found in blacklist	The common name of the subject matched a pattern entered into the blacklist.
6	SSL Control: Weak cipher being used	The symmetric cipher being negotiated was fewer than 64 bits. For a list of weak ciphers, see Common weak ciphers.
7	See #2, SSL Control: Certificate chain not complete	See #2, SSL Control: Certificate chain not complete.
8	SSL Control: Failed to decode Server Hello	The Server Hello from the SSL server was undecipherable. Also occurs when the certificate and Server Hello are in different packets, as is the case when connecting to a SSL server on a SonicWall appliance. This log event is informational, and does not affect the SSL connection.
9	SSL Control: Website found in whitelist	The common name of the subject (typically a website) matched a pattern entered into the Whitelist. Whitelist entries are always allowed, even if there are other policy violations in the negotiation, such as SSLv2 or weak ciphers.
10	SSL Control: HTTPS via SSLv2	The SSL session was being negotiated using SSLv2, which is known to be susceptible to certain man-in-the-middle attacks. Best practices recommend using SSLv3 or TLS instead.

7

Configuring Cipher Control

- About Cipher Control on page 108
- Firewall Settings > Cipher Control on page 108

About Cipher Control

Starting with SonicOS 6.5.4.1, you can allow or block any or all TLS and SSH ciphers. This functionality applies to:

- DPI-SSL (TLS traffic inspected by the firewall)
- HTTPS MGMT (TLS sessions accessing the firewall)
- SSL Control (inspect TLS traffic passing through the firewall: non DPI-SSL)

Any change to the TLS ciphers apply to all TLS traffic.

The list of ciphers displayed in the **Firewall Settings > Cipher Control** page are a list of known TLS ciphers. The list of ciphers is a super set of supported ciphers. While this list contains all known ciphers, DPI-SSL and HTTPS MGMT support a much smaller list of ciphers. For example, DPI-SSL and HTTPS MGMT do not yet support TLS 1.3 ciphers or support some weak ciphers that are listed in **Firewall Settings > Cipher Control**.

The ciphers are ordered based on the security strengths, with ciphers on top more secure than the ones below. Both DPI-SSL and HTTPS MGMT implementations use the relative ordering of their supported ciphers based on **Firewall Settings > Cipher Control**; that is, for the DPI-SSL supported ciphers, DPI-SSL orders them based on the ciphers listed in **Firewall Settings > Cipher Control**. The same is true for HTTPS MGMT ciphers.

Firewall Settings > Cipher Control

Topics:

- TLS Ciphers on page 109
- SSH Ciphers on page 115

TLS Ciphers

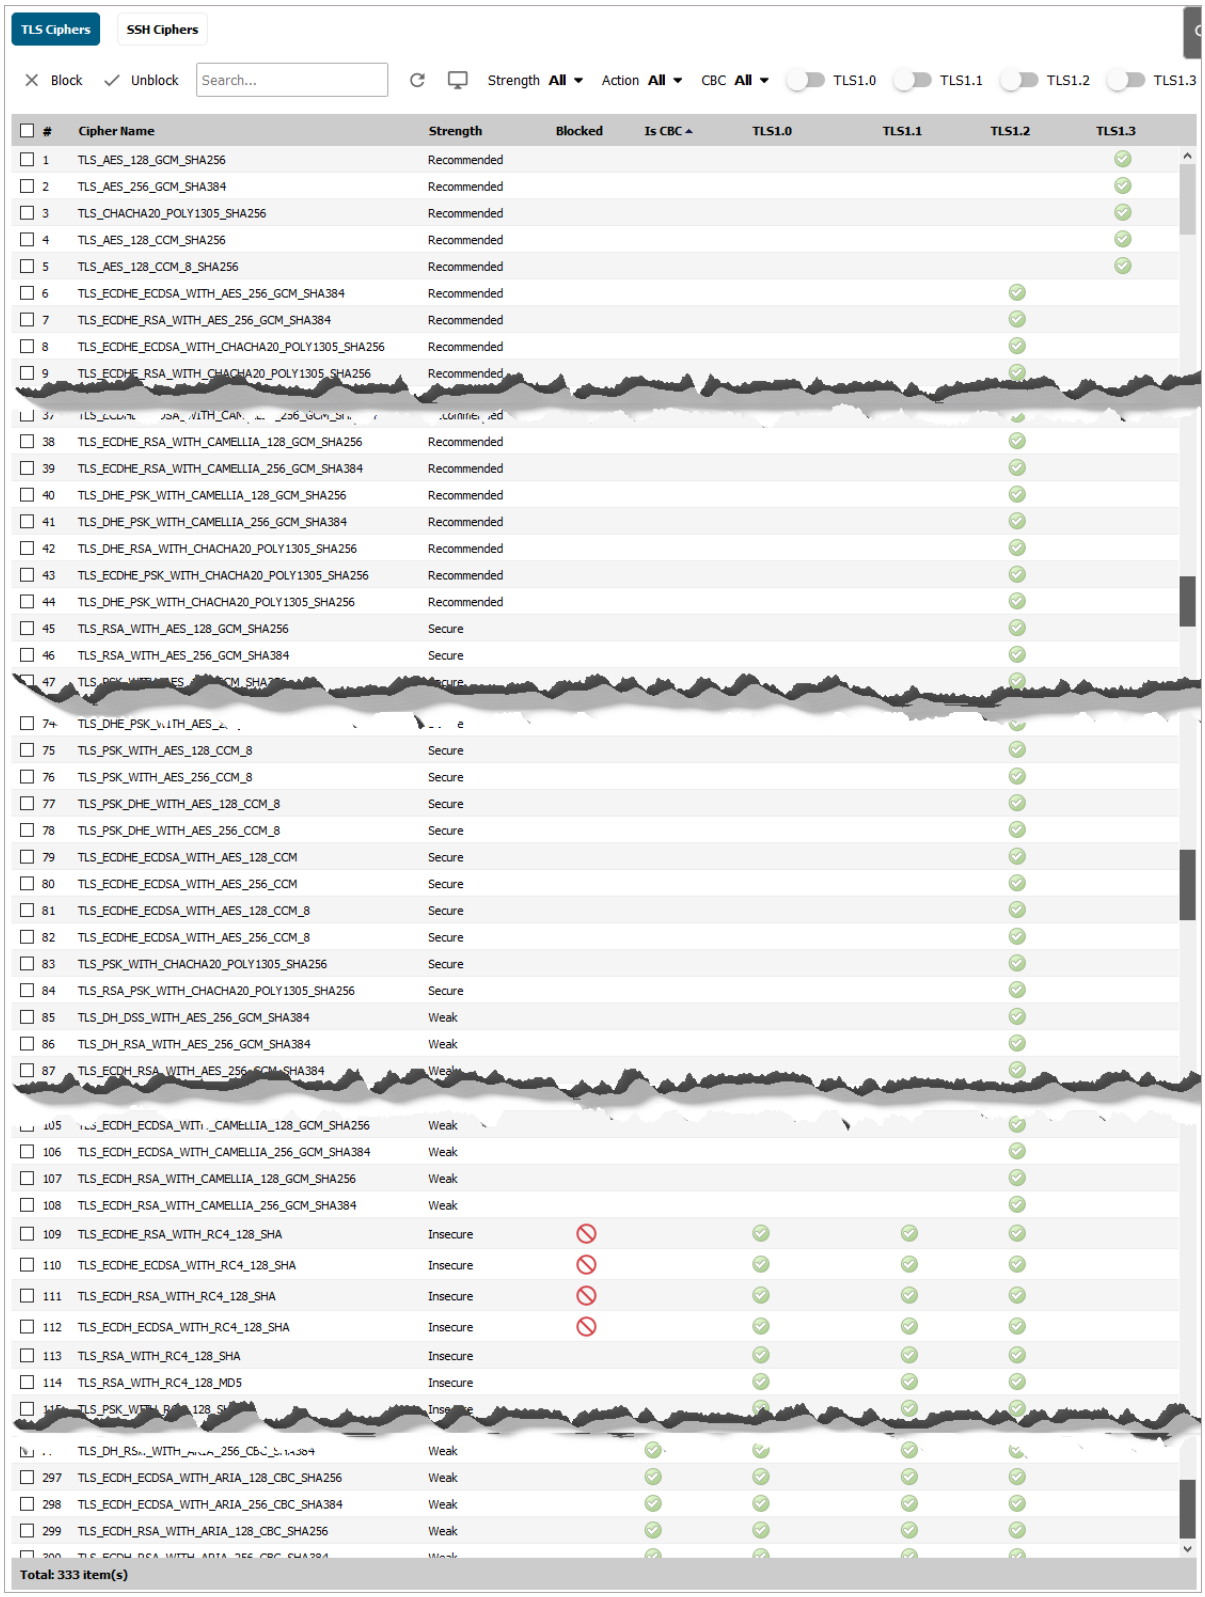
TLS Ciphers
SSH Ciphers
Block
Unblock
Search...
Strength All
Action All
CBC All
TLS1.0
TLS1.1
TLS1.2
TLS1.3
Cipher Name Strength Blocked Is CBC TLS1.0 TLS1.1 TLS1.2 TLS1.3
1 TLS_AES_128_GCM_SHA256 Recommended
2 TLS_AES_256_GCM_SHA384 Recommended
3 TLS_CHACHA20_POLY1305_SHA256 Recommended
4 TLS_AES_128_CCM_SHA256 Recommended
5 TLS_AES_128_CCM_8_SHA256 Recommended
6 TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384 Recommended
7 TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 Recommended
8 TLS_ECDHE_ECDSA_WITH_CHACHA20_POLY1305_SHA256 Recommended
9 TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256 Recommended
38 TLS_ECDHE_RSA_WITH_CAMELLIA_128_GCM_SHA256 Recommended
39 TLS_ECDHE_RSA_WITH_CAMELLIA_256_GCM_SHA384 Recommended
40 TLS_DHE_PSK_WITH_CAMELLIA_128_GCM_SHA256 Recommended
41 TLS_DHE_PSK_WITH_CAMELLIA_256_GCM_SHA384 Recommended
42 TLS_DHE_RSA_WITH_CHACHA20_POLY1305_SHA256 Recommended
43 TLS_ECDHE_PSK_WITH_CHACHA20_POLY1305_SHA256 Recommended
44 TLS_DHE_PSK_WITH_CHACHA20_POLY1305_SHA256 Recommended
45 TLS_RSA_WITH_AES_128_GCM_SHA256 Secure
46 TLS_RSA_WITH_AES_256_GCM_SHA384 Secure
75 TLS_PSK_WITH_AES_128_CCM_8 Secure
76 TLS_PSK_WITH_AES_256_CCM_8 Secure
77 TLS_PSK_DHE_WITH_AES_128_CCM_8 Secure
78 TLS_PSK_DHE_WITH_AES_256_CCM_8 Secure
79 TLS_ECDHE_ECDSA_WITH_AES_128_CCM Secure
80 TLS_ECDHE_ECDSA_WITH_AES_256_CCM Secure
81 TLS_ECDHE_ECDSA_WITH_AES_128_CCM_8 Secure
82 TLS_ECDHE_ECDSA_WITH_AES_256_CCM_8 Secure
83 TLS_PSK_WITH_CHACHA20_POLY1305_SHA256 Secure
84 TLS_RSA_PSK_WITH_CHACHA20_POLY1305_SHA256 Secure
85 TLS_DH_DSS_WITH_AES_256_GCM_SHA384 Weak
86 TLS_DH_RSA_WITH_AES_256_GCM_SHA384 Weak
106 TLS_ECDH_ECDSA_WITH_CAMELLIA_256_GCM_SHA384 Weak
107 TLS_ECDH_RSA_WITH_CAMELLIA_128_GCM_SHA256 Weak
108 TLS_ECDH_RSA_WITH_CAMELLIA_256_GCM_SHA384 Weak
109 TLS_ECDHE_RSA_WITH_RC4_128_SHA Insecure
110 TLS_ECDHE_ECDSA_WITH_RC4_128_SHA Insecure
111 TLS_ECDH_RSA_WITH_RC4_128_SHA Insecure
112 TLS_ECDH_ECDSA_WITH_RC4_128_SHA Insecure
113 TLS_RSA_WITH_RC4_128_SHA Insecure
114 TLS_RSA_WITH_RC4_128_MD5 Insecure
297 TLS_ECDH_ECDSA_WITH_ARIA_128_CBC_SHA256 Weak
298 TLS_ECDH_ECDSA_WITH_ARIA_256_CBC_SHA384 Weak
299 TLS_ECDH_RSA_WITH_ARIA_128_CBC_SHA256 Weak
Total: 333 item(s)

Cipher Name	Name of the cipher.
Strength	Strength of the cipher: • **Recommended** • **Secure** • **Weak** • **Insecure**
Blocked	Indicates, with a **Blocked** icon, whether the cipher has been blocked from being used
Is CBC	Indicates, with an **Enabled** icon, whether the cipher uses CBC (Cipher-Block Chaining) mode
TLS1.0 **TLS1.1** **TLS1.2** **TLS1.3**	Indicates, with an **Enabled** icon, whether the cipher is used in the TLS (Transport Layer Security) protocol version
Total	Indicates the total number of cipher entries in the table

Topics:

- Blocking/Unblocking Ciphers on page 110
- Filtering Ciphers on page 110

Blocking/Unblocking Ciphers

To block ciphers:

1. Navigate to **MANAGE | Security Configuration > Firewall Settings > Cipher Control**.
2. Click **TLS Ciphers**.
3. Either:
   - Select the cipher(s) to block.
   - Click the checkbox in the table header.
4. Click **X Block**. A **Blocked** icon displays in the **Blocked** column for each blocked cipher.

To unblock ciphers:

1. Navigate to **MANAGE | Security Configuration > Firewall Settings > Cipher Control**.
2. Click **TLS Ciphers**.
3. Either:
   - Select the cipher(s) to unblock.
   - Click the checkbox in the table header.
4. Click **Unblock**. The **Blocked** icon no longer displays in the **Blocked** column for the blocked cipher(s).

Filtering Ciphers

You can filter ciphers to easily configure which ciphers should be allowed or blocked.

Topics:

- Selecting Display Options on page 111

- Displaying Ciphers by Strength on page 112
- Displaying Ciphers by Block/Unblock on page 112
- Displaying Ciphers by CBC Mode on page 113
- Displaying Ciphers by TLS Protocol Version on page 114

Selecting Display Options

The **TLS Ciphers** table displays which TSL protocols support which ciphers. You can also display other protocols that support the ciphers:

- DPI-SSL
- HTTPS management
- SSL control

The **Display** icon helps you filter ciphers based on functional use cases (DPI-SSL, HTTPS MGMT, and pass-through traffic). For example, if cipher X is blocked, the expected behavior is:

- **DPI-SSL** – Cipher X is no longer a part of the TLS context and is not a part of the client advertised ciphers sent by the firewall handshaking with origin server.
- **HTTPS MGMT** – Cipher X is not a part of the HTTPS MGMT server application running on the firewall. Thus, if a TLS client negotiates just cipher X, the TLS handshake between client and firewall fails.
- **SSL Control** – As this refers to traffic (other than DPI-SSL decrypted sessions) passing through the firewall, the firewall blocks any TLS connection between origin client and origin server that uses/negotiates Cipher X.

To display other protocols:

1 Navigate to **MANAGE | Security Configuration > Firewall Settings > Cipher Control**.

2 Click **TLS Ciphers**.

3 Click the **Display Options** icon. The **Select Columns to Display** popup displays.

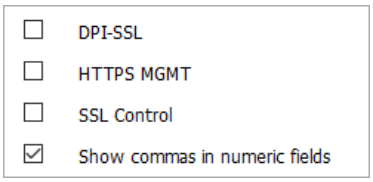

4 Select the protocol(s) to display:

- **DPI-SSL** – This option is not selected by default.
- **HTTPS MGMT** – This option is not selected by default.
- **SSL Control** – This option is not selected by default.
- **Show commons in numeric fields** – This option is selected by default.

5 Click **SAVE**. The column(s) are added to the **TLS Ciphers** table.

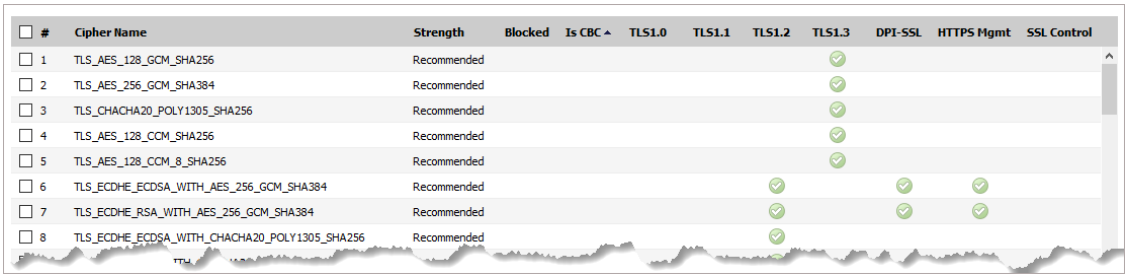

Displaying Ciphers by Strength

Ciphers are rated according to their strength:

- Recommended
- Secure
- Insecure
- Weak

The **TLS Ciphers** table displays all ciphers of all strengths. You can restrict the **TLS Cipher** table to display only those ciphers of a particular strength.

To display ciphers by strength:

1 Navigate to **MANAGE | Security Configuration > Firewall Settings > Cipher Control**.

2 Click **TLS Ciphers**.

3 Select the strength from **Strength**. The default is **All**.

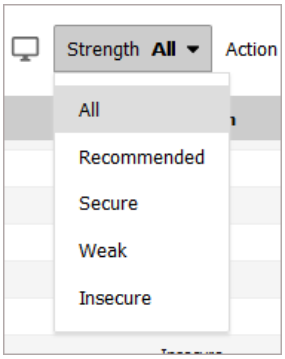

TLS Cipher table redisplays, showing only those ciphers with the corresponding strength and the **Strength** drop-down menu reflects the displayed strength.

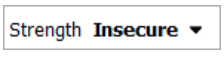

Displaying Ciphers by Block/Unblock

The **TLS Ciphers** table displays all blocked and unblocked ciphers. You can restrict the **TLS Cipher** table to display only those ciphers that are blocked or unblocked.

To display ciphers by strength:

1 Navigate to **MANAGE | Security Configuration > Firewall Settings > Cipher Control**.

2 Click **TLS Ciphers**.

3 Select the allow/block action from **Action**.

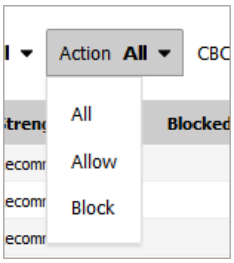

- **All** (default)
- **Allow** (unblock)
- **Block**

The **TLS Cipher** table redisplays, showing only those ciphers with the corresponding action and **Action** reflects the displayed action.



Displaying Ciphers by CBC Mode

The **TLS Ciphers** table displays all ciphers for all ciphers regardless of whether they use CBC mode. You can restrict the display to whether a cipher uses CBS mode.

To display whether ciphers use CBC mode:

1 Navigate to **MANAGE | Security Configuration > Firewall Settings > Cipher Control**.

2 Click **TLS Ciphers**.

3 Select whether the cipher uses CBC mode from **CBC**.

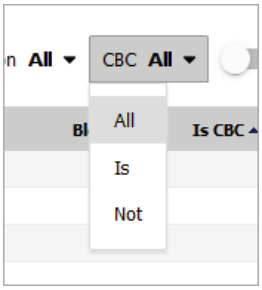

- **All** (default)
- **Is** (uses CBC mode)
- **Not** (does not use CBC mode)

The **TLS Cipher** table redisplays according to the selection, showing an **Enabled** icon in the Is **CBC** column for those ciphers using CBC mode and nothing in the **CBC** column for those that don't.

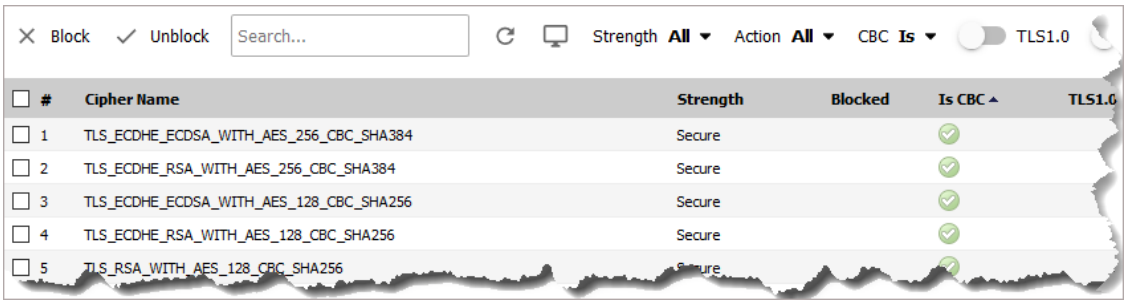

Displaying Ciphers by TLS Protocol Version

The **TLS Ciphers** table displays all ciphers for all TLS protocol versions. You can restrict the display by version of TLS protocol the cipher supports.

To display ciphers by TLS protocol:

1 Navigate to **MANAGE | Security Configuration > Firewall Settings > Cipher Control**.

2 Click **TLS Ciphers**.

3 Click the **TLS** version(s) for displaying ciphers:

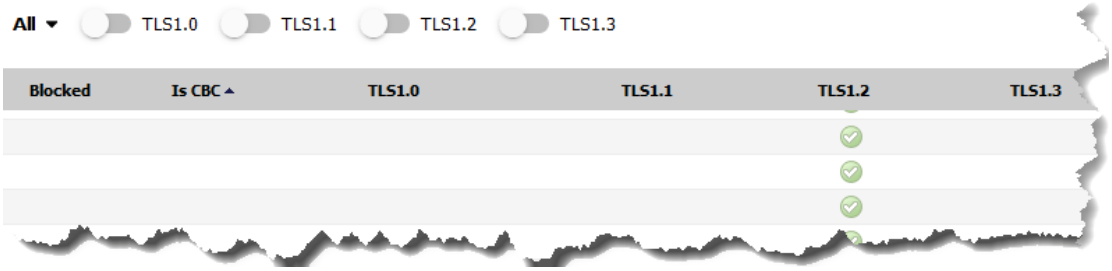

- **TLS1.0**
- **TLS1.1**
- **TLS1.2**
- **TLS1.3**

The display is restricted to only those ciphers supporting that TLS version:

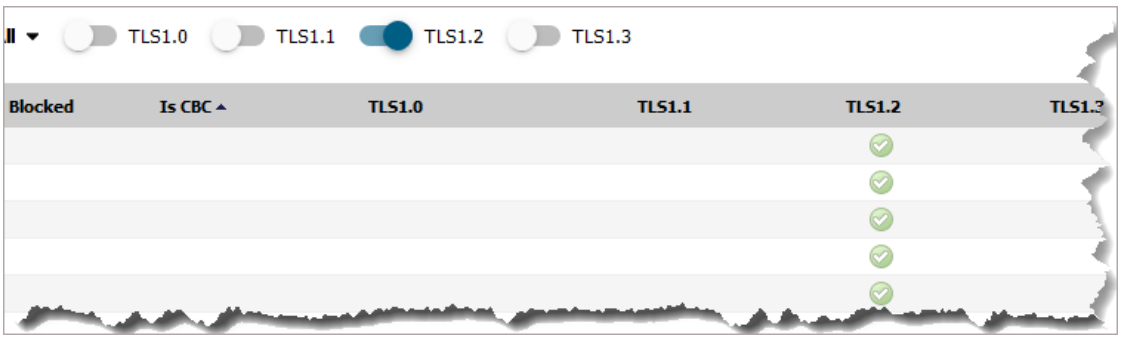

NOTE: If a cipher supports more than the selected version, the Enabled icon displays for the other supported versions as well.

SSH Ciphers

The SSH Ciphers page of **MANAGE | Security Configuration > Firewall Settings > Cipher Control** allows you to specify which cryptographic SSH ciphers SonicOS uses.

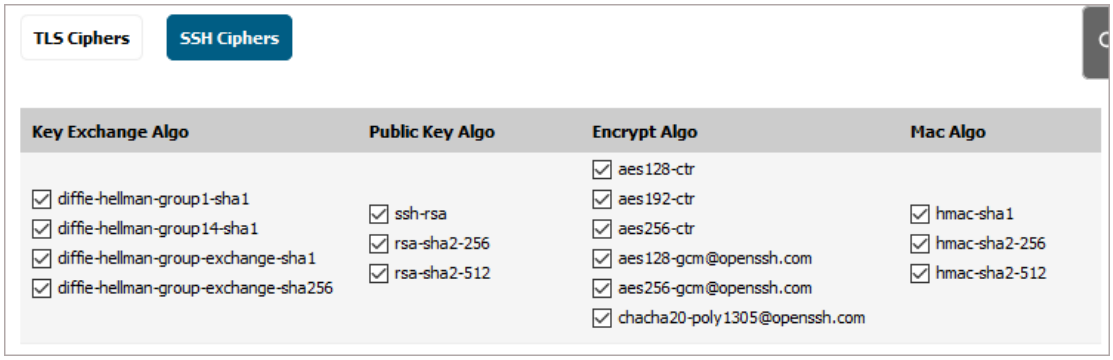

Key Exchange Algo	Lists the cryptographic algorithms used to exchange cryptographic keys between two parties
Public Key Algo	Lists the asymmetric cryptographic algorithms using pairs of public keys
Encrypt Algo	Lists the encryption algorithms used in secure transfers of files, such as FTP transfers
Mac Algo	Lists the algorithms using a MAC (message authentication code) value to authenticate messages

To select or deselect SSH ciphers:

1 Navigate to **MANAGE | Security Configuration > Firewall Settings > Cipher Control**.

2 Click **SSH Ciphers**.

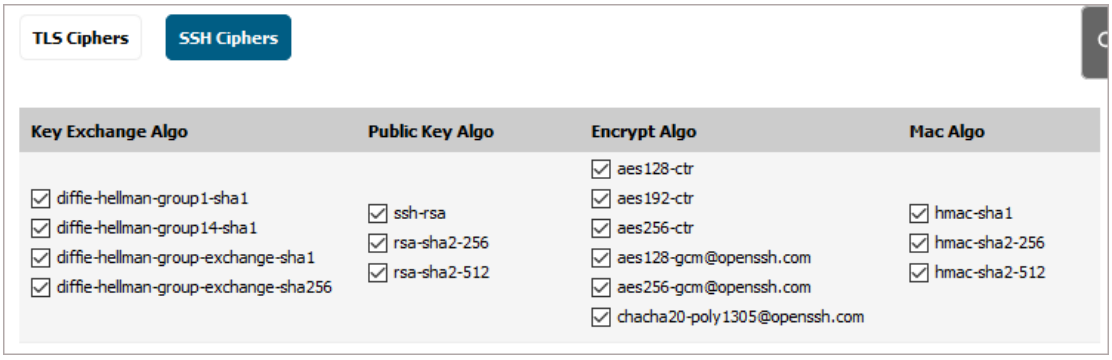

IMPORTANT: All SSH ciphers are selected by default.

3 Select the SSH algorithm to use or ignore. A status message displays at the bottom of the screen.



TIP: You may see a processing message that displays briefly.

Part 2

Security Config | Security Services

- Managing SonicWall Security Services
- Configuring Content Filtering Service
- DPI-SSL Enforcement
- Configuring Client AV Enforcement
- Configuring Client CF Enforcement
- Managing SonicWall Gateway Anti-Virus Service
- Activating Intrusion Prevention Service
- Configuring Capture ATP
- Activating Anti-Spyware Service
- Configuring SonicWall Real-Time Blacklist
- Configuring Geo-IP Filters
- Configuring Botnet Filters

8

Managing SonicWall Security Services

- About SonicWall Security Services on page 117
- Configuring Security Services on page 118

About SonicWall Security Services

SonicWall offers a variety of subscription-based security services to provide layered security for your network. SonicWall security services are designed to integrate seamlessly into your network to provide complete protection.

The following subscription-based security services are listed in **Security Services** on the firewall's management interface:

- SonicWall Content Filtering Service
- SonicWall Client Anti-Virus
- SonicWall Gateway Anti-Virus
- SonicWall Intrusion Prevention Service
- SonicWall Anti-Spyware
- SonicWall RBL Filter
- SonicWall Geo-IP Filter
- SonicWall Botnet Filter

TIP: After you register your firewall, you can try FREE TRIAL versions of SonicWall Content Filtering Service, SonicWall Client Anti-Virus, SonicWall Gateway Anti-Virus, SonicWall Intrusion Prevention Service, and SonicWall Anti-Spyware.

You can activate and manage SonicWall security services directly from the SonicWall management interface or from https://www.mysonicwall.com.

Configuring Security Services

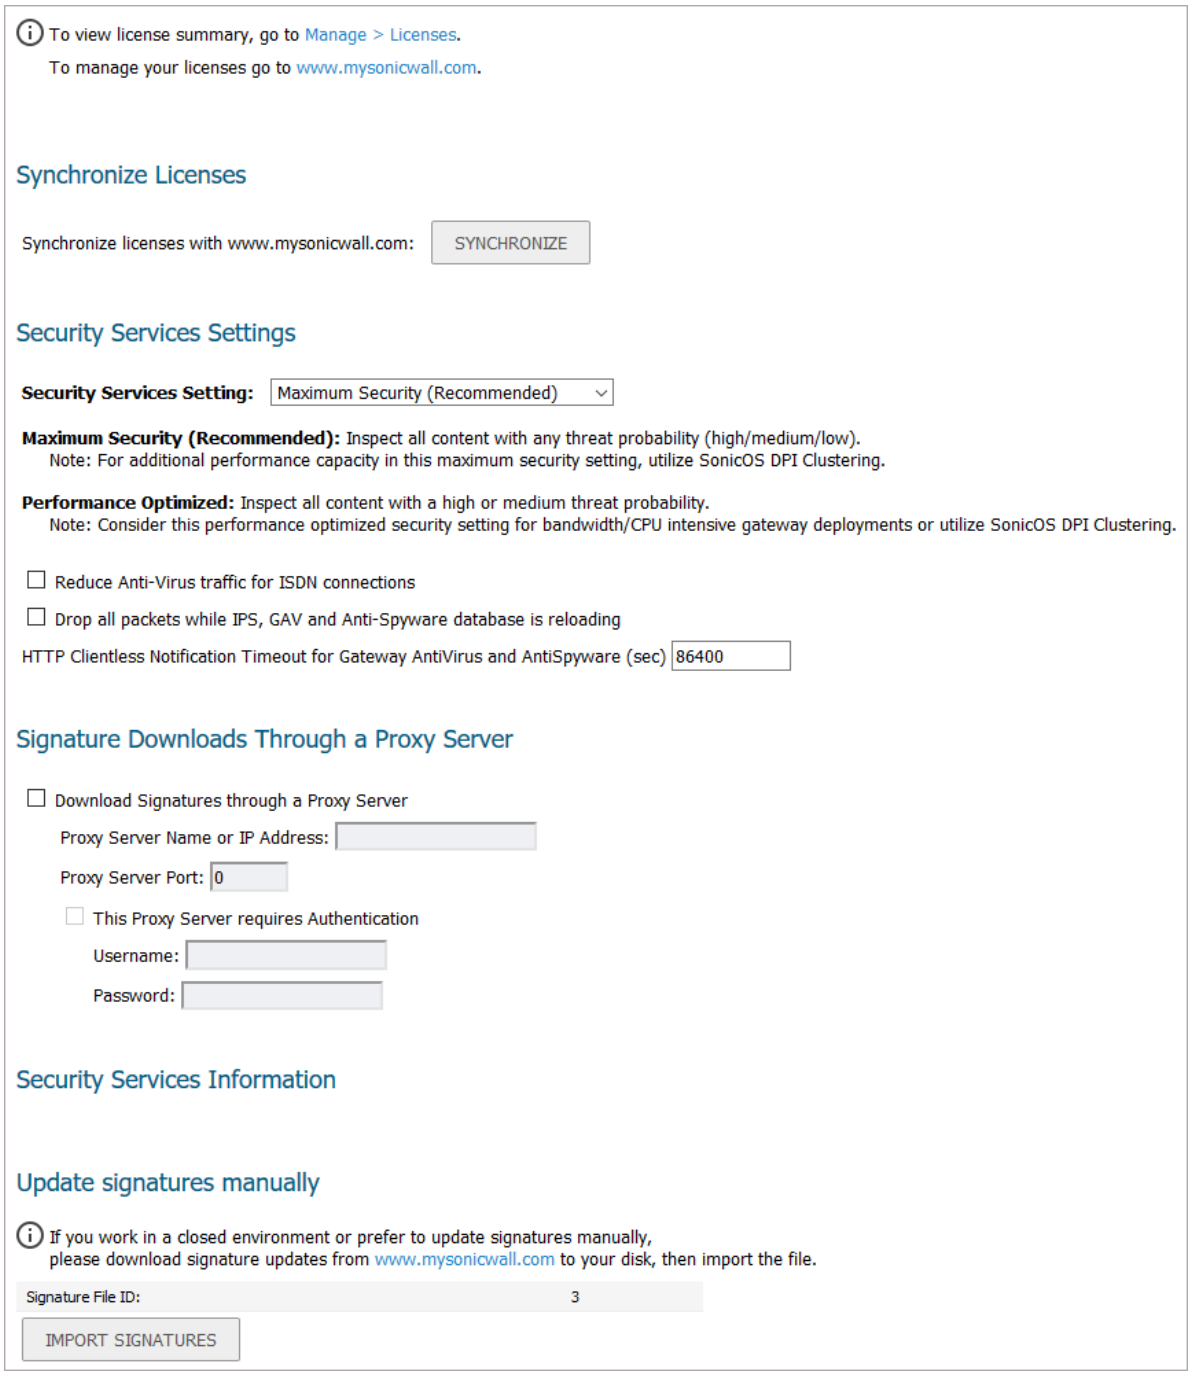

The following sections describe global configurations that are done on the panels of the**MANAGE | Security Configuration > Security Services > Base Setup** page:

- Viewing and Managing Licenses on page 119
- Synchronize Licenses on page 119
- Security Services Settings on page 119
- Signature Downloads Through a Proxy Server on page 120

- Security Services Information on page 121
- Update Signature Manually on page 121

Viewing and Managing Licenses

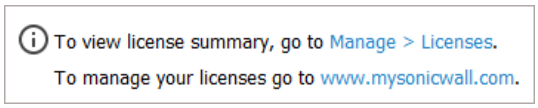

The top of the page displays two links:

- **To view license summary, go to Manage > Licenses.** – Click the link to view your licenses and their status on the **MANAGE | Updates > Licenses** page.
- **To manage your licenses go to www.mysonicwll.com.** Click the link to try, upgrade, or purchase/renew licenses at MySonicWall.

Synchronize Licenses

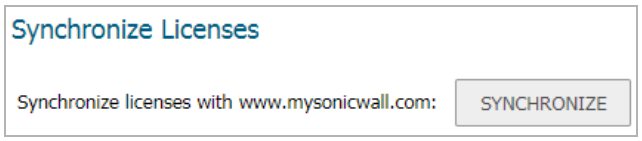

To synchronize your licenses with your `mysonicwall.com` account, click the **SYNCHRONIZE** button after **Synchronize licenses with www.mysonicwall.com**.

Security Services Settings

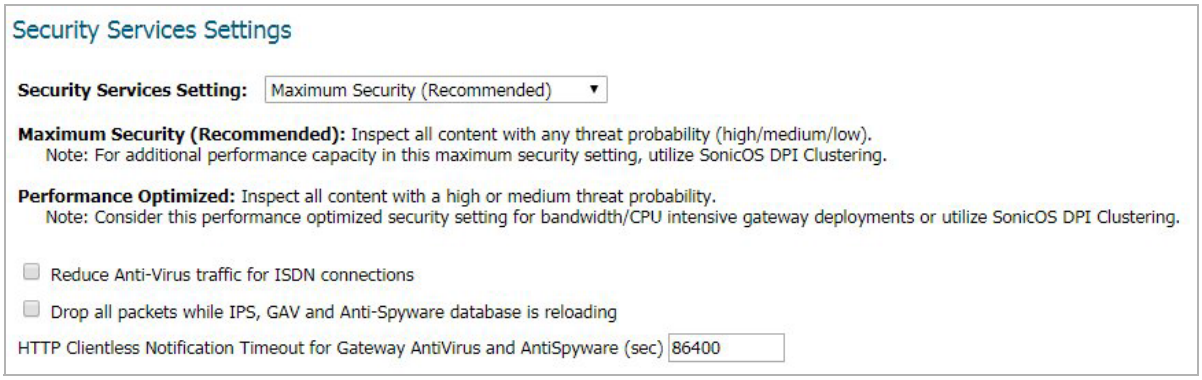

The **Security Services Settings** section provides the following options for fine-tuning SonicWall security services:

- **Security Services Settings** - This drop-down menu specifies whether SonicWall security services are applied to maximize security or to maximize performance:
    - **Maximum Security (Recommended)** - Inspect all content with any threat probability (high/medium/low). For additional performance capacity in this maximum security setting, utilize SonicOS HA Clustering.

- **Performance Optimized** - Inspect all content with a high or medium threat probability. Consider this performance optimized security setting for bandwidth or CPU intensive gateway deployments or utilize SonicOS HA Clustering.

The **Maximum Security** setting provides maximum protection. The **Performance Optimized** setting utilizes knowledge of the currently known threats to provide high protection against active threats in the threat landscape.

- **Reduce Anti-Virus traffic for ISDN connections** - Select this feature to enable the SonicWall Anti-Virus to check only once a day (every 24 hours) for updates and reduce the frequency of outbound traffic for users who do not have an “always on” Internet connection.
- **Drop all packets while IPS, GAV and Anti-Spyware database is reloading** - Select this option to instruct the firewall to drop all packets whenever the IPS, GAV, and Anti-Spyware database is updating.
- **HTTP Clientless Notification Timeout for Gateway AntiVirus and AntiSpyware** - Set the timeout duration after which the firewall notifies users when GAV or Anti-Spyware detects an incoming threat from an HTTP server. The default timeout is one day (86400 seconds).

Signature Downloads Through a Proxy Server

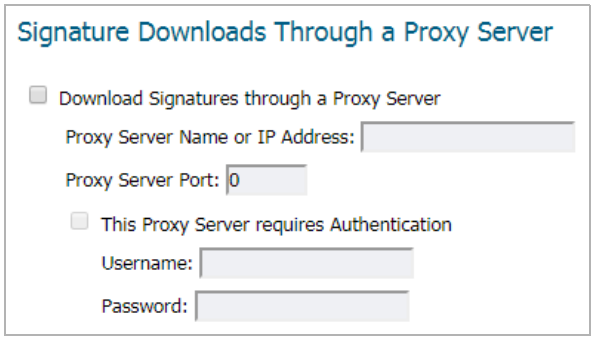

This section provides the ability for SonicWall network security appliances that operate in networks where they must access the Internet through a proxy server to download signatures. This feature also allows for registration of SonicWall network security appliances through a proxy server without compromising privacy.

To enable signature download or appliance registration through a proxy server:

1. Select the **Download Signatures through a Proxy Server** checkbox.
2. In the **Proxy Server Name or IP Address** field, enter the host name or IP address of the proxy server.
3. In the **Proxy Server Port** field, enter the port number used to connect to the proxy server.
4. Select the **This Proxy Server requires Authentication** checkbox if the proxy server requires a **username** and **password**.
5. If the appliance has not been registered with `MySonicWall.com`, two additional fields are displayed:
   - **MySonicWall Username** - Enter the username for the `mysonicwall.com` account that the appliance is to be registered to.
   - **MySonicWall Password** - Enter the `mysonicwall.com` account password.
6. Click **ACCEPT**.

Security Services Information

This panel is not currently used.

Update Signature Manually

The Manual Signature Update feature is intended for networks where reliable, broadband Internet connectivity is either not possible or not desirable (for security reasons). The Manual Signature Update feature provides a method to update the latest signatures at your discretion:

1 Download the signatures from http://www.mysonicwall.com to a separate computer, a USB drive, or other media.

2 Upload the signatures to the firewall.

The same signature update file can be used on all SonicWall network security appliances that meet these requirements:

- Devices that are registered to the same `mysonicwall.com` account
- Devices that belong to the same class of SonicWall network security appliances.

To manually update signature files:

1 On the **Security Services > Summary** page, scroll to the **Update Signatures Manually** heading at the bottom of the page. Record the **Signature File ID** for the device.



2 Log on to http://www.mysonicwall.com using the `mysonicwall.com` account that was used to register the SonicWall network security appliance.

NOTE: The signature file can only be used on firewalls that are registered to the `mysonicwall.com` account that downloaded the signature file.

3 Click on **Download Signatures** under the **Downloads** heading.

4 In the pull down window next to **Signature ID:**, select the appropriate SFID for your firewall.

5 Download the signature update file by clicking on **Click here to download the Signature file**.

NOTE: The remaining steps can be performed while disconnected from the Internet.

6 Return to the **Security Services > Summary** page on the firewall management interface.

7 Click the **Import Signatures** button.

8 In pop-up dialog that appears, click the **browse** button and navigate to the location of the signature update file.

9 Click **Import**. The signatures are uploaded for the security services that are enabled on the firewall.

9

Configuring Content Filtering Service

IMPORTANT: The **MANAGE | Security Configuration > Security Services > Content Filter** has two variants: one for SonicWall CFS and one for Websense Enterprise.

- Security Services > Content Filter: SonicWall CFS on page 123
- Security Services > Content Filter: Websense Enterprise on page 135

Security Services > Content Filter: SonicWall CFS

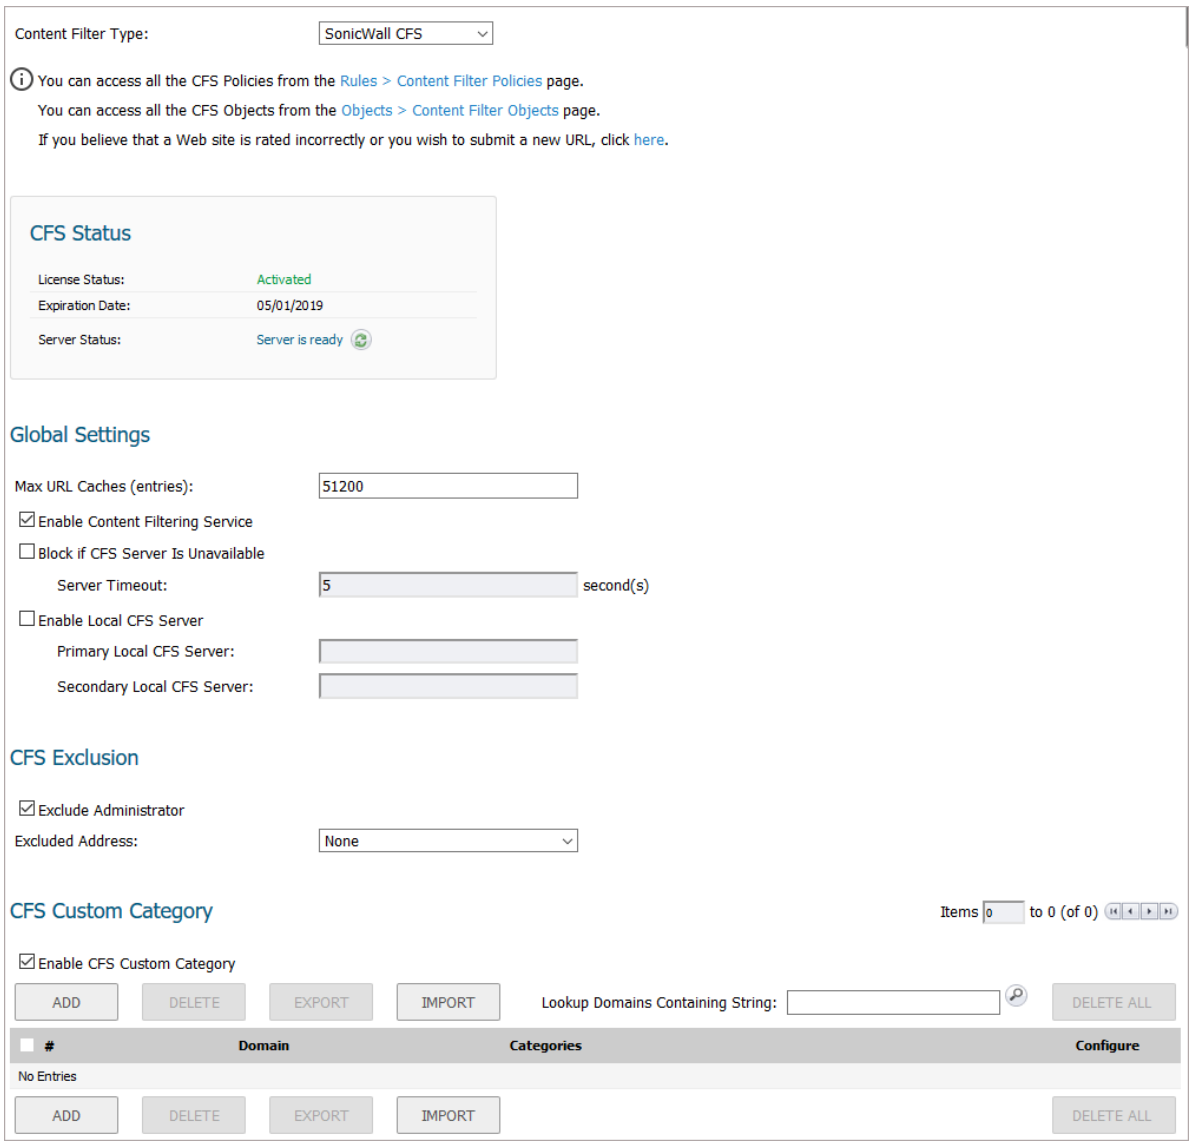

NOTE: Content Filtering Service (CFS) content is not supported in Wire Mode.

You can activate Content Filter Objects and configure SonicWall Content Filtering Service (SonicWall CFS) as well as Websense Enterprise, a third-party Content Filtering product, from the **MANAGE | Security Configuration > Security Services > Content Filter** page.

Topics:

- About CFS on page 124
- Enabling CFS on page 126
- Configuring CFS Policies on page 128
- Configuring CFS Custom Categories on page 128

About CFS

The SonicWall Content Filtering Service (CFS) delivers content filtering enforcement for educational institutions, businesses, libraries, and government agencies. With content filter objects, you can control the websites students and employees can access using their IT-issued computers while behind the organization's firewall.

NOTE: For more a detailed description of CFS, as well as how to license and install it, see the *SonicWall SonicOS 6.5 Release Notes* and the *SonicWall Content Filtering Service Upgrade Guide*. Also, for how to create Content Filter Objects for CFS policies, see *SonicWall SonicOS 6.5 Policies*.

CFS compares requested websites against a massive cloud database that contains millions of rated URIs, IP addresses, and websites. It also provide you with the tools to create and apply policies that allow or deny access to sites based on individual or group identity and/or by time of day.

Topics:

- About Threat API on page 124
- About CFS Policies on page 125
- About Content Filter Objects on page 125
- How CFS Works on page 125
- CFS Blocking of Individual Videos on page 126
- About CFS Logs on page 126

About Threat API

IMPORTANT: Before configuring Threat API, you must enable it.

NOTE: SonicOS Threat API requires that the firewall has a Content Filtering System (CFS) license.

The SonicOS Threat API provides API access to SonicWall firewall services. Compared with current firewall GUI/CLI user interfaces, Threat API is simple and makes good use of the standard HTTP protocol. With the trend toward cloud deployment, Threat API can more easily be used than traditional SonicOS GUI/CLI.

Malicious threats can originate from URLs or IP addresses. Lists of these threats can be large and change frequently. SonicOS can already block custom lists of URLs and IP addresses, but it's inconvenient because you have to log in and update the lists by hand. Using an API interface makes it much easier.

The Threat list is sent to SonicOS using the Threat API feature. Threats can be added in either of the following formats:

- URLs (`https://malicious123.example.com/malware`)
- IP addresses (`10.10.1.25`)

Third parties can generate the threat list and pass it to the firewall using Threat API.

For IP addresses in the threat list, SonicOS initially creates a default Threat API Address Group and then creates an Address Object (AO) for each IP address in the threat list. The you configure Firewall Access Rules that reference that Address Group and block the IP addresses.

SonicOS adds the URLs to its CFS Threat URI list. You enable Threat API Enforcement in the associated CFS Profile and configure a Content Filtering System (CFS) policy to block the URLs in the threat list. When a threat is blocked by CFS, the user sees a block message in their browser.

About CFS Policies

A CFS policy determines whether a packet is filtered (by applying the configured CFS Action) or simply allowed through to the user. A CFS policy defines the filtering conditions to which a packet is compared:

- Name
- Source Zone
- Destination Zone
- Source Address
- User/Group
- Schedule

If a packet matches all the defined conditions, the packet is filtered according to the corresponding CFS Profile, and the CFS Action is applied.

NOTE: If authentication data for User/Group is not available during matching, no match is made for this condition. This strategy prevents performance issues, especially when Single Sign-On is in use.

Each CFS policy has a priority level, and policies with higher priorities are checked first.

CFS uses a policy table internally to manage all the configured policies. For each policy element, the table is constructed by the configuration data and runtime data. The configuration data includes parameters that define the policy from the user interface, such as policy name, properties and others. The runtime data includes the parameters used for packet handling.

CFS also uses a policy lookup table to accelerate runtime policy lookup for matching conditions:

- Source zone
- Destination zone
- IPv4 AO
- IPv6 AO

About Content Filter Objects

CFS uses Content Filter Objects in CFS Policies to identify URIs and domains for filtering and to specify the type of action to be taken when filtering. For more information about Content Filter Objects, see *SonicWall SonicOS 6.5 Policies*.

Under the CFS rating design, a domain may be resolved to one of four ratings; from highest to lowest priority, the ratings are:

1. Block
2. Passphrase
3. Confirm
4. BWM (bandwidth management)

If the URL is not categorized into any of these ratings, then the operation will be allowed.

How CFS Works

1. A packet arrives and is examined by CFS.
2. CFS checks it against the configured exclusion addresses and allows it through if a match is found.

3 CFS checks its policies to find the first policy that matches these conditions in the packet:

- Source zone
- Destination zone
- Address object
- Users/group
- Schedule
- Enabled state

4 CFS uses the CFS Profile defined in the matching policy to do the filtering and returns the corresponding action for this packet.

NOTE: If no policy is matched, the packet is passed through without any action by CFS.

5 CFS performs the action defined in the CFS Action Object for the matching policy.

CFS Blocking of Individual Videos

SonicWall Content Filtering Service (CFS) can selectively filter and block individual YouTube videos.

NOTE: SonicWall CFS can only block *specific* YouTube videos. It cannot block categories of videos. This feature only works if the SonicWall CFS server already has a rating for the specific video identified in the "v=" parameter of the URI. Each video URI to be blocked must be added individually to SonicWall CFS.

This feature is not supported when a local CFS server; only when using the SonicWall public CFS server. This is due to a conflict with the blacklist/whitelist feature in the local CFS server.

No SonicOS configuration is required to use this feature.

About CFS Logs

In **MANAGE | Logs & Reporting > Log Settings > Base Setup**, a new subcategory, **Content Filter**, has been added to the **Security Services** category. This new subcategory lists these logs:

- CFS Alert
- Website Accessed
- Website Blocked

For information about configuring these logs, see *SonicWall SonicOS 6.5 Logs and Reporting*.

Enabling CFS

IMPORTANT: Before enabling CFS and configuring your CFS policies, configure your Content Filter Objects as described in *SonicWall SonicOS 6.5 Policies*.

To enable CFS:

1 Navigate to the **MANAGE | Security Configuration > Security Services > Content Filter** page.

2 Choose the content filtering service from the **Content Filter Type** drop-down menu:

- **SonicWall CFS** (default)
- **Websense Enterprise** (for how to configure Websense Enterprise, see Security Services > Content Filter: Websense Enterprise on page 135)

3 In the **Global Settings** section, specify the maximum URL entries that can be cached in the **Max URL Caches (entries)** field. The default is **51200**.

   The URL rating is saved with a cached URL entry, which speeds processing of known URLs.

4 To enable content filter for all packets, select the **Enable Content Filtering Service** checkbox. This option is selected by default. To bypass content filtering for all packets, deselect this option.

5 To enable content filtering for HTTPS sites, select the **Enable HTTPS content filtering** checkbox. This option is not selected by default.

   When this option is enabled, CFS performs URL rating look up in this order:

   a Searches the client `hello` for the Server Name, which CFS uses to obtain the URL rating.

   b If the Server Name is not available, searches the SSL certificate for the Common Name, which CFS uses to obtain the URL rating.

   c If neither Server Name nor Common Name is available, CFS uses the IP address to obtain the URL rating.

6 To limit the time for obtaining a rating request when filtering, select the **Block if CFS Server Is Unavailable** checkbox. This option is not selected by default.

   a When this option is selected, the **Server Timeout** field becomes available. Enter the maximum time, in seconds, the CFS service has to respond to rating requests. The minimum is 2 seconds, the maximum is 10 seconds, and the default is **5** seconds.

7 To bypass content filtering for all requests from an account with administrator privileges, select the **Exclude Administrator** checkbox in the **CFS Exclusion** section. This option is selected by default.

8 To bypass content filtering for all requests from a category of address objects, choose the address object from the **Excluded Address** drop-down menu. The default is **None**. You can also create a new address object by choosing **Create new address object**; for information about creating an address object, see *SonicWall SonicOS 6.5 Policies*.

9 Click **ACCEPT**.

Enabling the Local CFS Server

The Local CFS Responder (Local CFS) allows the Content Filtering Service to receive URL ratings directly from a local responder, rather than from a remote public responder. For information on configuring and using Local CFS, see the *Local CFS Administration Guide*.

To enable the Local CFS Responder:

1 Navigate to the **MANAGE | Security Configuration > Security Services > Content Filter** page.

2 Select **SonicWall CFS** (default) as the content filtering service from **Content Filter Type**.

3 Scroll to the **Global Settings** section.

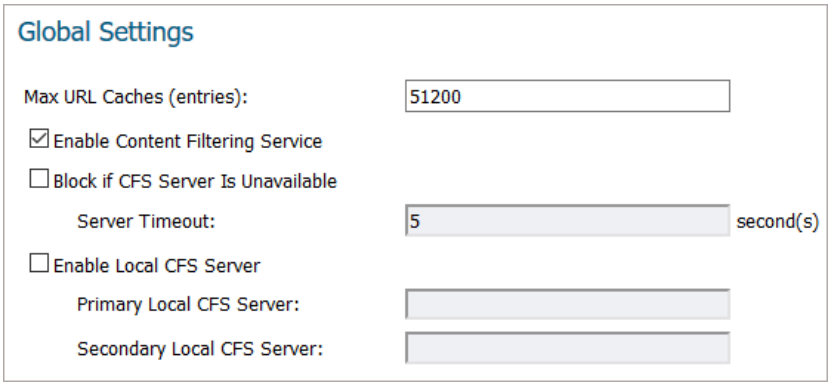

4 Select **Enable Local CFS Server**.

5 Enter the IP addresses for the primary and secondary local CFS servers in the **Primary Local CFS Server** and **Secondary Local CFS Server** fields.

6 Mousing over the **Statistics** icon to the right of the **Primary Local CFS Server** field will display information about the server entered.

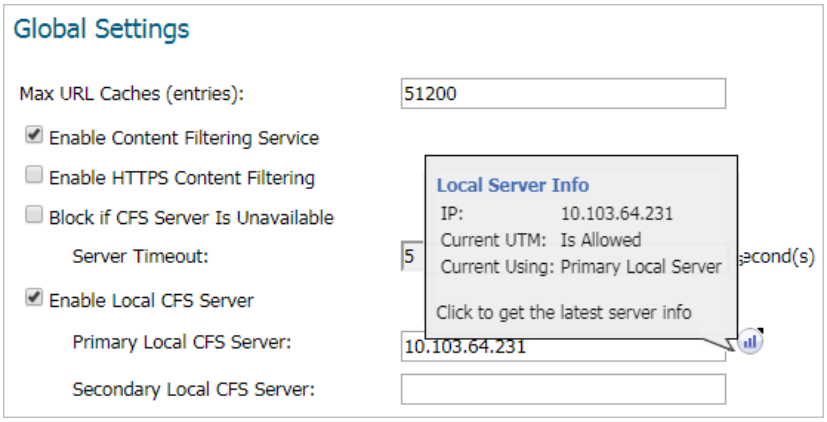

7 Click **ACCEPT**.

Configuring CFS Policies

To add, edit, or delete CFS policies, go to the **MANAGE | Policies > Objects > Content Filter Objects** page. For more information, see *SonicWall SonicOS 6.5 Policies*.

Configuring CFS Custom Categories

This section describes the CFS Custom Category table and provides instructions for configuring, editing, and deleting CFS custom categories. Importing and exporting the custom category table are also described.

Topics:

- About the CFS Custom Category Table on page 129
- Searching the CFS Custom Category Table on page 129
- Configuring a CFS Custom Category on page 130

- Exporting the CFS Custom Category Table on page 132
- Importing a CFS Custom Category Table on page 132
- Editing a CFS Custom Category on page 134
- Deleting CFS Custom Categories on page 134

About the CFS Custom Category Table

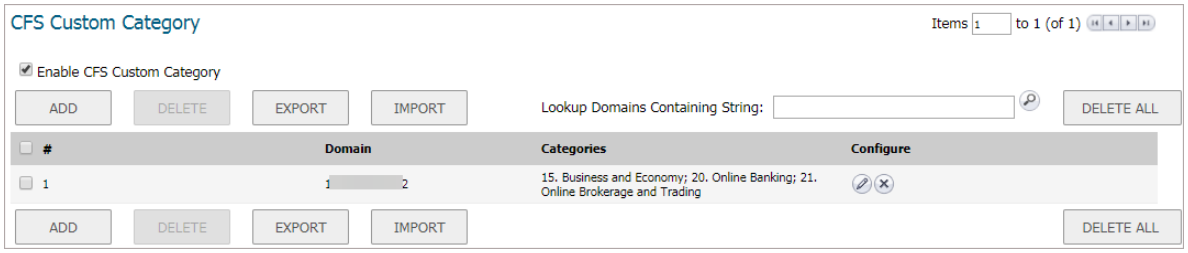

Domain	IP address of the domain to which the custom category applies.
Categories	Categories selected for the custom category.
Configure	Displays the **Edit** and **Delete** icons for each domain.

Searching the CFS Custom Category Table

You can search a long table for a specific IP address by:

1 Entering an IP address in the Lookup Policies by Address field. The IP address can be in either format:
   - `192.168.168.168`
   - `fe80::c2ea:e4ff:fe59:a634`
2 Clicking the **Search** (magnifying glass) icon.

Requesting a Rating Review

If you believe that a web site is rated incorrectly or you wish to submit a new URL, you submit a request to the SonicWall Content Filtering Service by:

- Clicking on the link at the top of the **Security Services > Content Filter** page, `If you believe that a Web site is rated incorrectly or you wish to submit a new URL, click here.`
- Going to http://cfssupport.sonicwall.com/Support/web/eng/newui/viewRating.jsp.

The **CFS URl Rating Review Request** form displays.

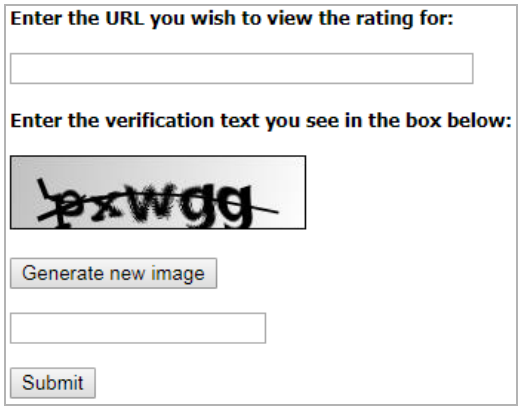

Configuring a CFS Custom Category

You can customize ratings for certain URLs. Up to 5,000 valid entries are supported. Custom categories are processes like those categories provided by the backend server. When CFS checks the ratings for one URL, it checks the user rating first and then the rating from the backend server. CFS categories are managed and built dynamically using configuration strings passed from the backend server.

Topics: :

- Enabling Custom Categories on page 130
- Configuring a Custom Category on page 130

Enabling Custom Categories

Before you can use custom categories, you must enable the service.

To enable custom categories:

1. Navigate to **MANAGE | Security Configuration > Security Services > Content Filter.**
2. Scroll to **CFS Custom Category**.



3. Select the **Enable CFS Custom Category** checkbox. This option is not selected by default.
4. Click **ACCEPT**.

Configuring a Custom Category

To define a custom category:

1. Navigate to **MANAGE | Security Configuration > Security Services > Content Filter**.

2 Scroll to **CFS Custom Categories**.

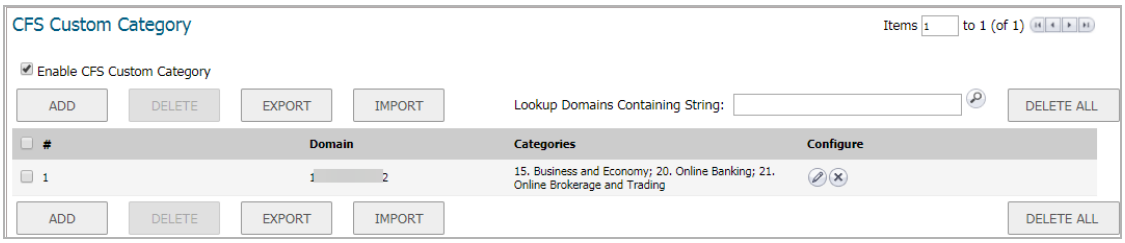

3 Click **ADD**. The **CFS Custom Category** dialog displays.

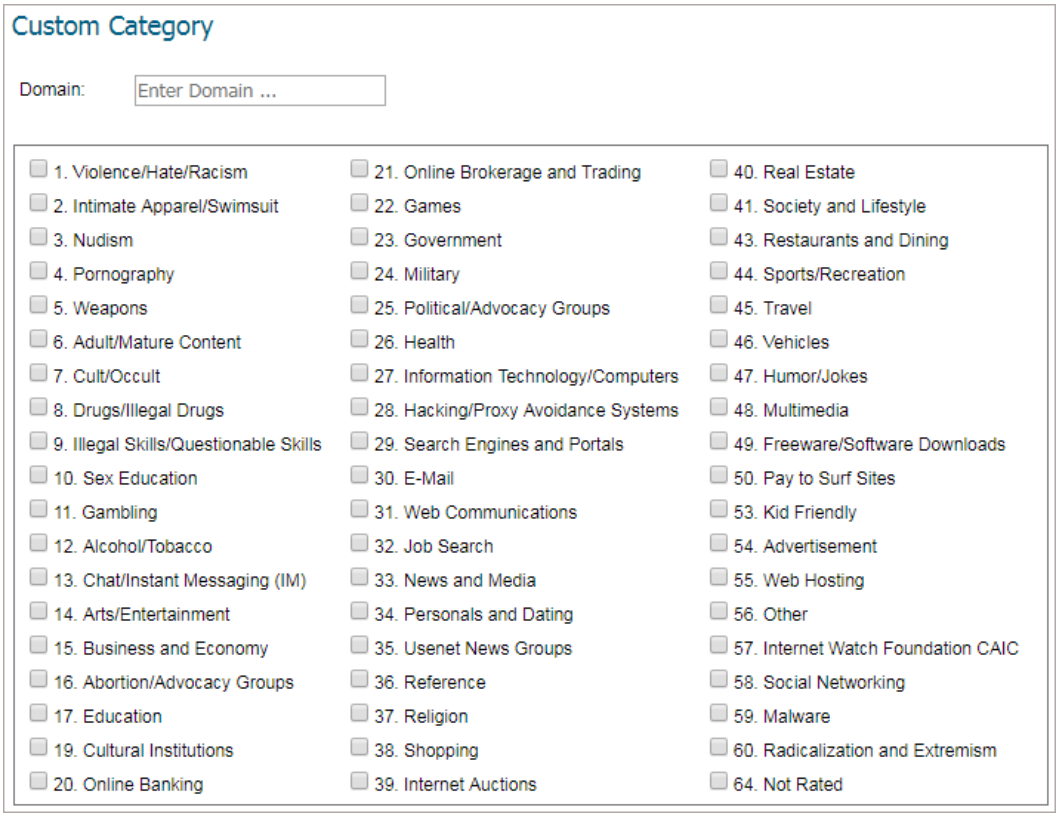

4 In the **Domain** field, enter the IP address or domain name of the domain for which the custom category applies:

- The IP address can be either of these formats:
  - `192.168.168.168`
  - `fe80::c2ea:e4ff:fe59:a634`
- Omit the `www.` prefix for a domain name. If you include it, a confirmation message displays; when you click **OK**, the prefix is removed from the domain name in the **Domain** field:



5 Select up to four categories from the list.

6 Click **ADD**.

7 To create more CFS custom categories, repeat Step 4 through Step 6 for each policy.

NOTE: Each custom category you create is a separate entry in the **CFS Custom Category** table; they are not concatenated.

8 Click **CLOSE**. The **CFS Custom Category** table is updated.

Exporting the CFS Custom Category Table

You can export the **CFS Custom Category** table to a `.wri` file you can edit and save for importing.

To export the CFS Custom Category table:

1 Navigate to **MANAGE | Security Configuration > Security Services > Content Filter**.

2 Scroll to **CFS Custom Categories**.

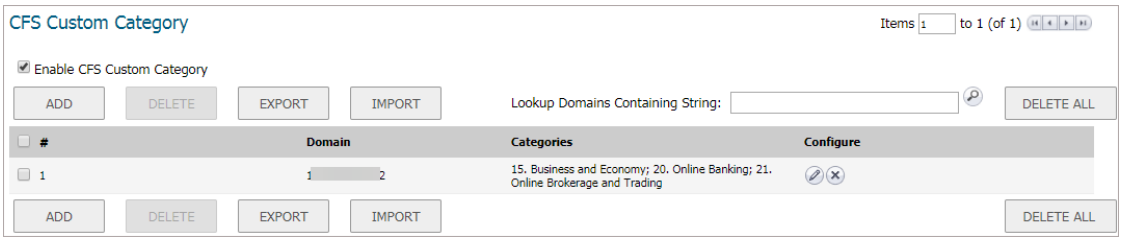

3 Click **EXPORT**. The **Opening cfsCustomCategoryData.wri** dialog displays.

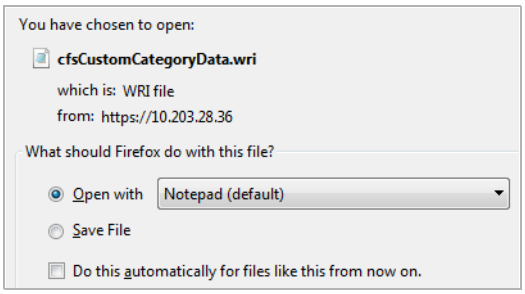

4 You can either open the file (default program is Notepad) or save it. If you:

- Open the file.
- Save the file, it is downloaded to your Downloads folder with the file name, `cfsCustomCaegoryData.wri`; new line characters are added after each entry.

NOTE: The file consists of all the **CFS Custom Category** table entries, all on one line.

5 Click **OK**.

Importing a CFS Custom Category Table

You can import a file of CFS Custom Category table entries. The entries in this file will overwrite the existing entries in the table.

The file should contain entries in this format:

`DomainName/IPAddress: Rating1[, Rating2[, Rating3[, Rating4]]] Separator`

Token	Definition
DomainName	A domain name, such as SonicWall. If you include the `www.` prefix, it is ignored.
IPAddress	A standard or IPv6 IP address, such as: • `192.168.168.168` • `fe80::c2ea:e4ff:fe59:a634`
Rating	A category rating from 1-255, as shown in the **Add CFS Custom Category** dialog. You can specify up to 4 ratings for each category.
Separator	A carriage return or new line separator:

Separator	Style
`\r\n`	Windows style, new line separator
`\n`	UNIX style, new line separator
`\r`	MAC OS style, new line separator

To import a custom category table:

1 Navigate to **MANAGE | Security Configuration > Security Services > Content Filter.**

2 Scroll to **CFS Custom Category**.

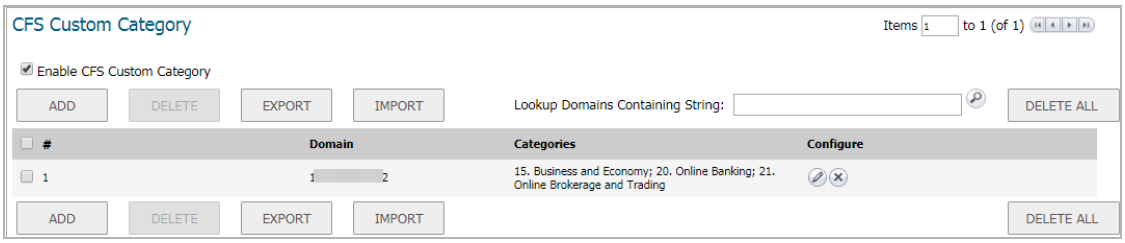

3 Click **IMPORT**. A confirmation dialog displays.

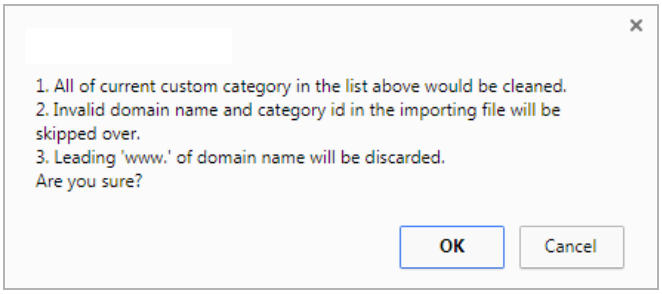

All current entries in the CFS Custom Category table are replaced with the entries in the file. Any entries you want to keep should be in the file.

TIP: Export the CFS Custom Category table and make any changes to the exported file before importing table entries.

4 Click **OK**.

Editing a CFS Custom Category

To edit a CFS custom category:

1 Click the **Edit** icon for the CFS custom category to be edited. The **CFS Custom Category** dialog displays. This dialog is the same as the **Add CFS Custom Category** dialog.

2 To make your changes, follow the appropriate procedures in Configuring a CFS Custom Category on page 130.

Deleting CFS Custom Categories

To delete CFS custom categories:

1 Do one of these:

- Click the **Delete** icon for the CFS custom categories to be deleted.
- Click the checkbox for one or more CFS custom categories to be deleted. The **DELETE** button becomes active; click it.

A confirmation message displays.

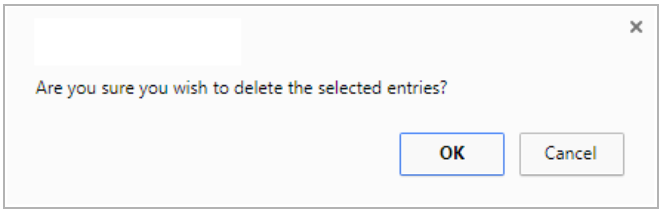

2 Click **OK**.

To delete all CFS custom categories:

1 Click the **DELETE ALL** button.

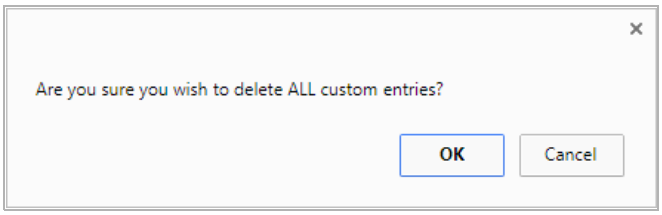

2 Click **OK**. All CFS custom categories are deleted.

Security Services > Content Filter: Websense Enterprise

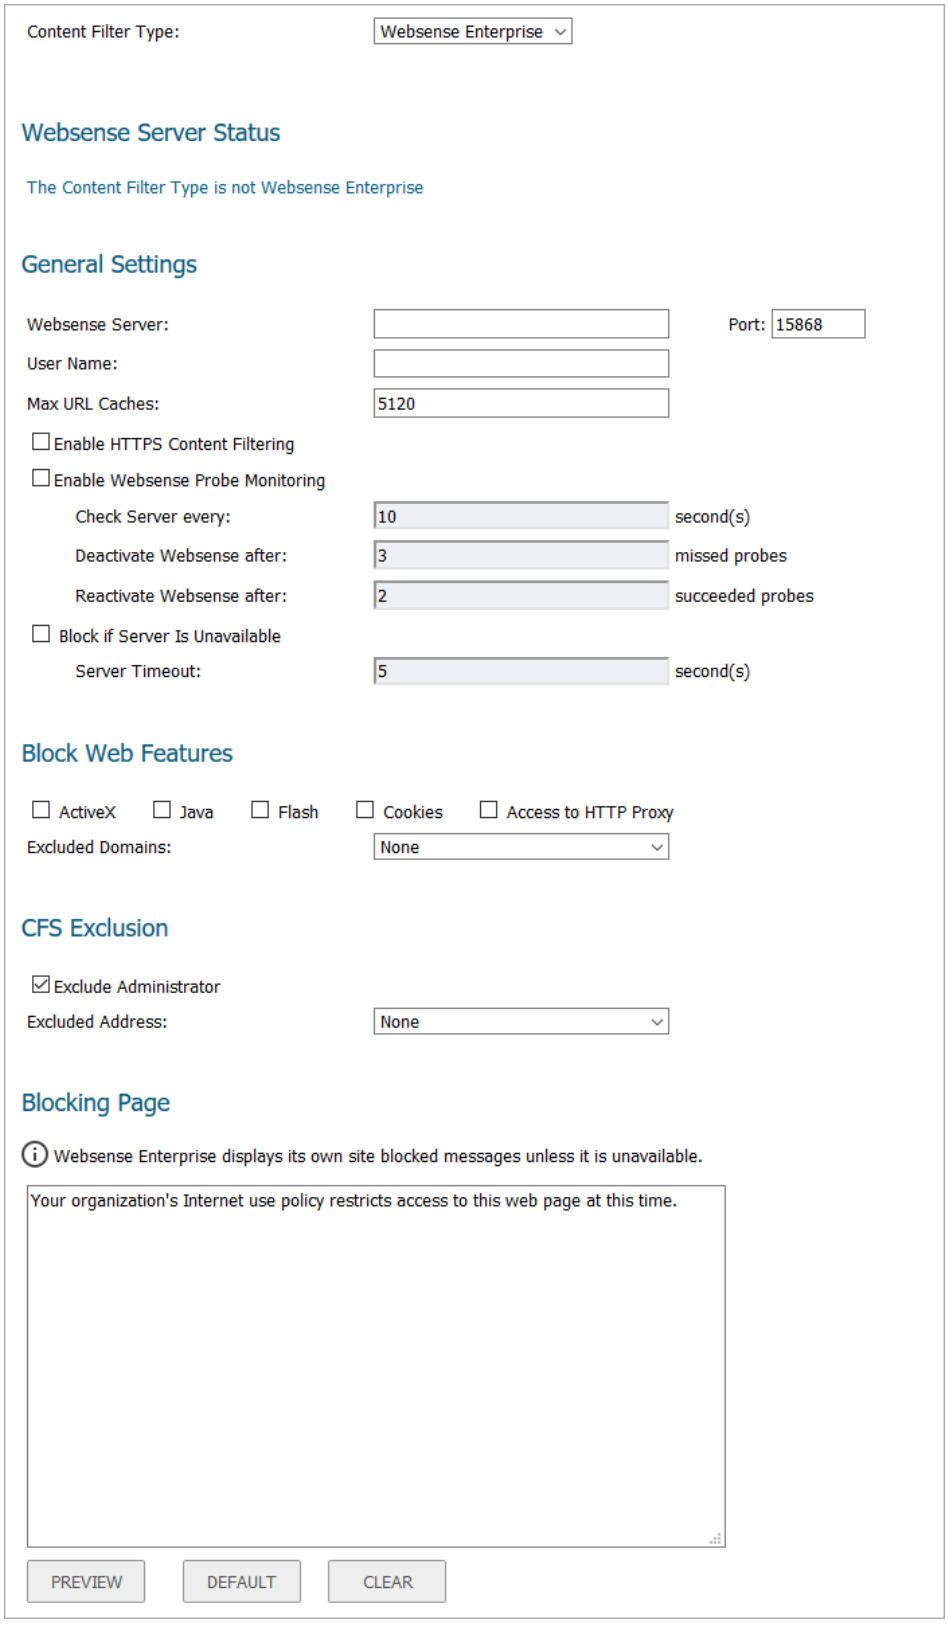

Topics:

- Selecting Websense Enterprise Content Filter Type on page 136
- Viewing Websense Server Status on page 136
- Configuring General Settings on page 137
- Configuring Web Features to Block on page 138
- Configuring CFS Exclusions on page 138
- Creating a Use Policy Blocking Page on page 139

Selecting Websense Enterprise Content Filter Type

To select Websense Filter as the content filter type:

1 Navigate to **MANAGE | Security Configuration > Security Services > Content Filter**.

2 Scroll to **Content Filter Type**.



3 Select **Websense Enterprise**. The options change.

4 Configure the Websense options.

5 Click **ACCEPT**.

TIP: Until you click **ACCEPT**, the Websense server status indicates the content filter type is not Websense Enterprise:

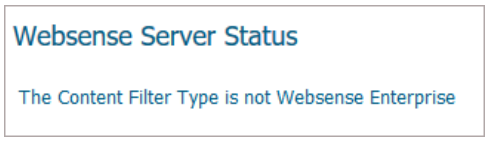

Viewing Websense Server Status

To view Websense server status:

1 Navigate to **MANAGE | Security Configuration > Security Services > Content Filter**.

2 Ensure **Websense Enterprise** is selected for **Content Filter Type**.

3 Scroll to **Content Filter Type**.

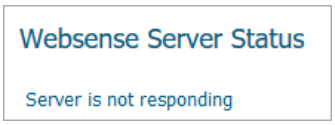

Configuring General Settings

To configure general Websense Enterprise settings:

1 Navigate to **MANAGE | Security Configuration > Security Services > Content Filter**.

2 Ensure **Websense Enterprise** is selected for **Content Filter Type**.

3 Scroll to **General Settings**.

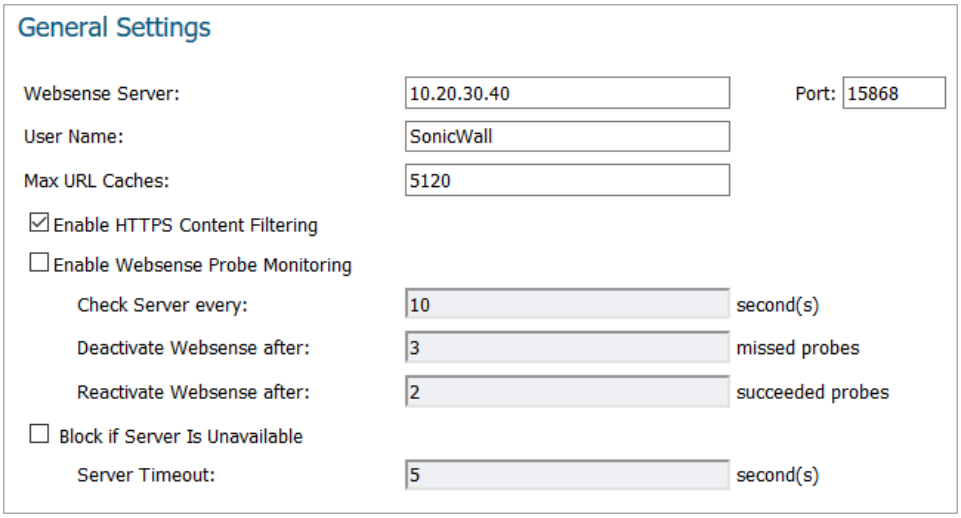

4 In the **Websense Server** field, enter the IP address of the Websense server.

5 In the **Port** field, enter the port for the Websense server. The default is **15868**.

6 In the **User Name** field, enter the username of the Websense server.

7 In the **Max URL Caches** field, enter the maximum number of URL caches. The minimum is 5120, the maximum is 51200, and the default is **5120**.

8 To enable HTTPS content filtering, select **Enable HTTPS Content Filtering**. This option is selected by default.

9 To monitor Websense probes, select **Enable Websense Probe Monitoring**. The following options become available. This option is not selected by default.

   a To specify the frequency of the probes, enter the probe interval, in seconds, in the **Check Server every … seconds** field. The minimum is 5 seconds, the maximum is 100 seconds, and the default is **10** seconds.

   b To deactivate Websense after a period of inactivity, enter the number of missed probes in the **Deactivate Websense after … missed probes** field. The minimum number is 1, the maximum number is 255, and the default is **3**.

   c To reactivate Websense after a period of inactivity, enter the number of successful probes in the **Reactivate Websense after … succeeded probes**. The minimum is 1, the maximum is 255, and the default is **2**.

10 To block web access is the server is unavailable, select **Block if Server is unavailable**. The following option becomes available. This option is not selected by default.

   a To specify the time the server is unavailable before access is blocked, enter the time in the **Server Timeout: … seconds** field. The minimum time is 1 second, the maximum is 10 seconds, and the default is **5** seconds.

11 Click **ACCEPT**.

Configuring Web Features to Block

To specify the web features to block:

1 Navigate to **MANAGE | Security Configuration > Security Services > Content Filter**.

2 Ensure **Websense Enterprise** is selected for **Content Filter Type**.

3 Scroll to **Block Web Features**.



4 Select one or more features to block (none are selected by default):

- **ActiveX**
- **Java**
- **Flash**
- **Cookies**
- **Access to HTTP Proxy**

5 Specify the domains to exclude from blocking from **Excluded Domains**. The default is **None**.

6 Click **ACCEPT**.

Configuring CFS Exclusions

To configure CFS exclusions:

1 Navigate to **MANAGE | Security Configuration > Security Services > Content Filter**.

2 Ensure **Websense Enterprise** is selected for **Content Filter Type**.

3 Scroll to **CFS Exclusion**.



4 To exclude the administrator from CFS, select **Exclude Administrator**. This option is selected by default.

5 Select the address object or group to exclude from **Excluded Address**. The default is **None**.

6 Click **ACCEPT**.

Creating a Use Policy Blocking Page

Websense Enterprise displays a default message when a web page is blocked. You can create a custom message to explain your company's use policy.

To create a custom blocking message:

1 Navigate to **MANAGE | Security Configuration > Security Services > Content Filter**.

2 Ensure **Websense Enterprise** is selected for **Content Filter Type**.

3 Scroll to **Blocking Page**.



The default message is displayed.

4 Replace the default message with your custom one.

5 To see the message as a user would see it, click **PREVIEW**.

A confirmation message displays.

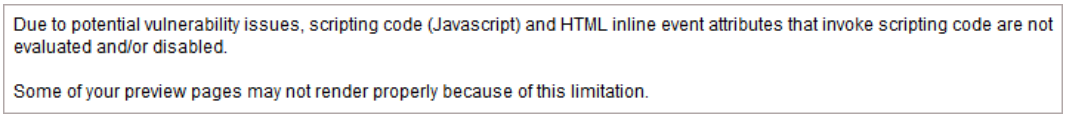

a Click **OK**. A popup page displays.



b Close the popup page.

6 To:

- Clear the contents of the **Blocking Page** field, click **CLEAR**.
- Restore the default message, click **DEFAULT**.

7 Click **ACCEPT**.

10

Configuring Client AV Enforcement

- Security Services > Client AV Enforcement
- Configuring Client Anti-Virus Enforcement

Security Services > Client AV Enforcement

SonicWall network security appliances running SonicOS 6.5.1 and higher support SonicWall Capture Client for client AV enforcement, in addition to McAfee and SentinelOne Anti-Virus. SonicWall Capture Client is powered by SentinelOne, and SentinelOne Client AV Enforcement is renamed as Capture Client AV Enforcement in SonicOS 6.5.1.3 and higher.

These client anti-virus services are all licensed separately, allowing you to purchase the desired number of each license for your deployment.

About SonicWall Capture Client

SonicWall Capture Client delivers multiple client protection capabilities. With a next-generation malware protection engine powered by SentinelOne, the SonicWall Capture Client delivers advanced threat protection with these key features:

- Continuous behavioral monitoring of the client that helps create a complete profile of file activity, application & process activity, and network activity.
- Multiple layered signatureless techniques that help protect against and remediate well known, little known, and even unknown malware, without regular scans or periodic updates.
- Roll-back capabilities that support policies that remove the threat completely and restore a targeted client to its original state, before the malware activity started. This removes the effort of manual restoration in the case of ransomware and similar attacks.
- Cloud-based management console that reduces overhead of management.

About Client AV Enforcement

By their nature, anti-virus products typically require regular, active maintenance on every PC. When a new virus is discovered, all anti-virus software deployed within an organization must be updated with the latest virus definition files. Failure to do so severely limits the effectiveness of anti-virus software and disrupts productive work time. With more than 50,000 known viruses and new virus outbreaks occurring regularly, the task of maintaining and updating virus protection can become unwieldy. Unfortunately, many small to medium businesses do not have adequate IT staff to maintain their anti-virus software. The resulting gaps in virus defenses may lead to data loss and decreased employee productivity.

The widespread outbreaks of viruses, such as NIMDA and Code Red, illustrate the problematic nature of virus defense for small and medium businesses. Users without the most current virus definition files allow these viruses to multiply and infect many other users and networks. SonicWall Client Anti-Virus Enforcement prevents

occurrences like these and offers a new approach to virus protection. SonicOS constantly monitors the version of the virus definition file and automatically triggers download and installation of new virus definition files to each user's computer. In addition, the firewall restricts each user's access to the Internet until they are protected, therefore acting as an enforcer of the company's virus protection policy. This new approach ensures the most current version of the virus definition file is installed and active on each PC on the network, preventing a rogue user from disabling the virus protection and potentially exposing the entire organization to an outbreak.

NOTE: You must purchase an Anti-Virus subscription to enforce Anti-Virus through the firewall's management interface.

Configuring Client Anti-Virus Enforcement

The Client Anti-Virus Service is enforced by enabling it on a zone. You can click the **Network > Zones** link at the top of the **Client AV Enforcement** page to jump to the **MANAGE | System Setup | Network > Zones** page, where you can edit the desired zone and enable the service.

For information on activating Network Anti-Virus Service, see Activating the Gateway Anti-Virus, Anti-Spyware, and IPS License on page 165.

The **MANAGE | Security Configuration | Security Services > Client AV Enforcement** page only displays the available settings when at least one client anti-virus service is licensed. Depending on the SonicOS version on your firewall and the licensed services, the **MANAGE | Security Configuration | Security Services > Client AV Enforcement** page will look different.

Capture Client and McAfee in Client AV Enforcement

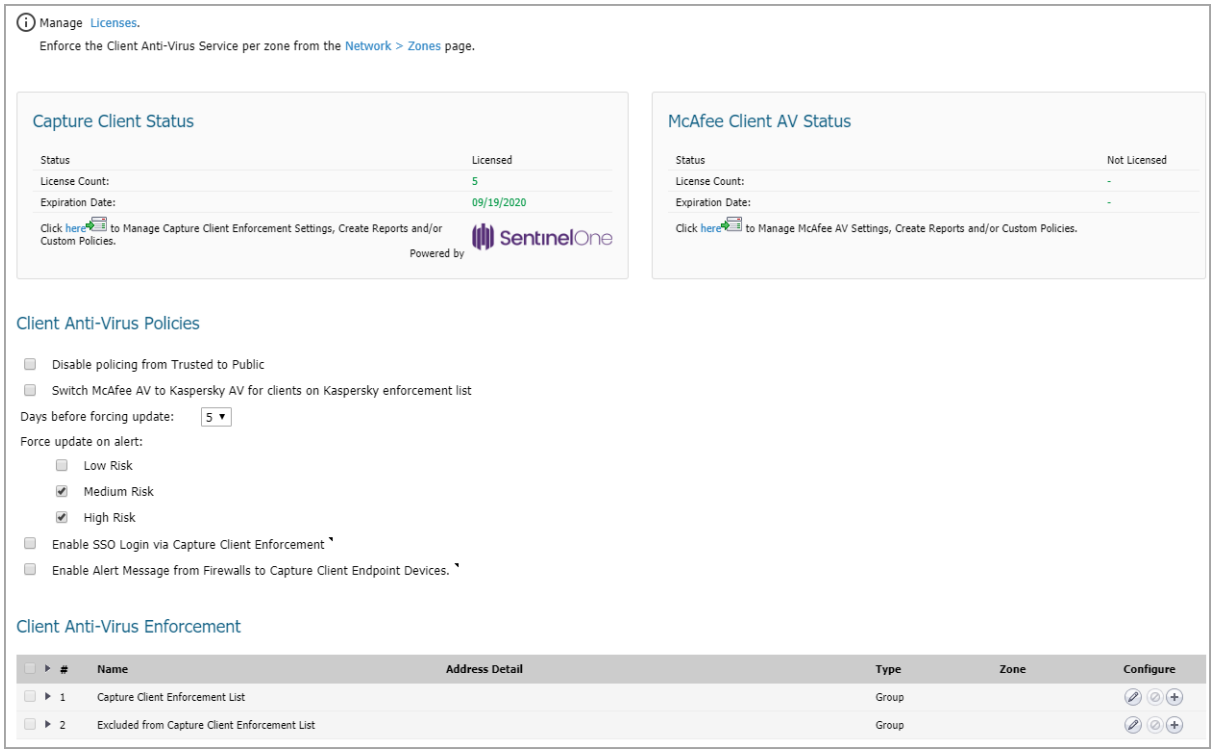

McAfee and SentinelOne in Client AV Enforcement

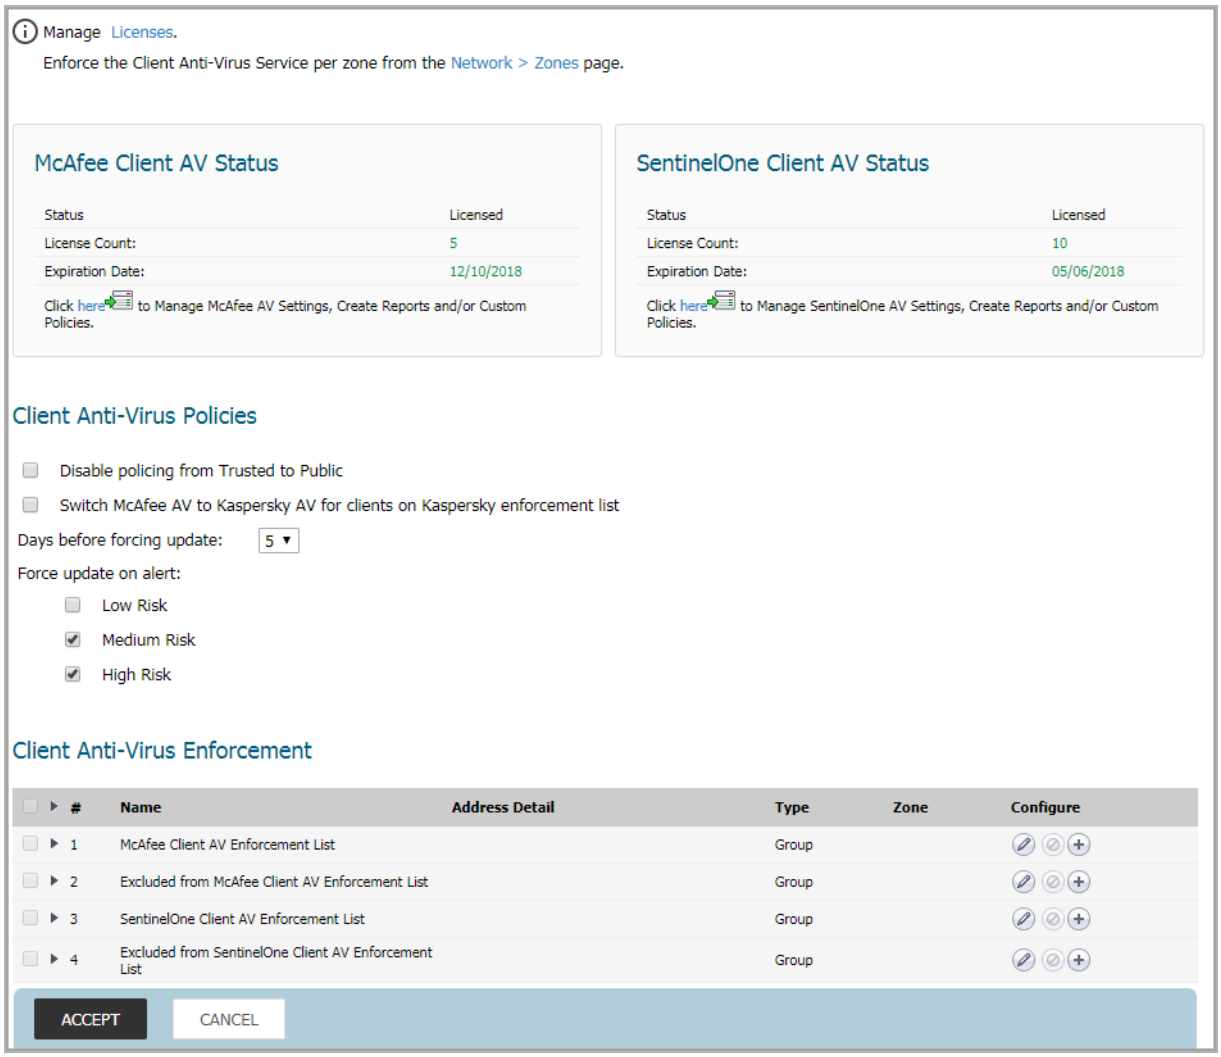

Topics:

- Capture Client Status
- McAfee / SentinelOne Client AV Status
- Client Anti-Virus Policies
- Client Anti-Virus Enforcement

Capture Client Status



The **Capture Client Status** section:

- Displays information about whether the firewall is licensed for client AV enforcement with Capture Client, the number of licenses, and the date the license expires.
- Contains a link to log into the Capture Client management interface, where you can manage and review protected users, groups and devices, configure security policies and certificates, and manage licenses, tenant settings, administrators, client installers and more.

  Clicking the link opens the Capture Client Management login page. When Capture Client is already licensed on your firewall, click the **Login with MySonicWall** link and enter your MySonicWall credentials.

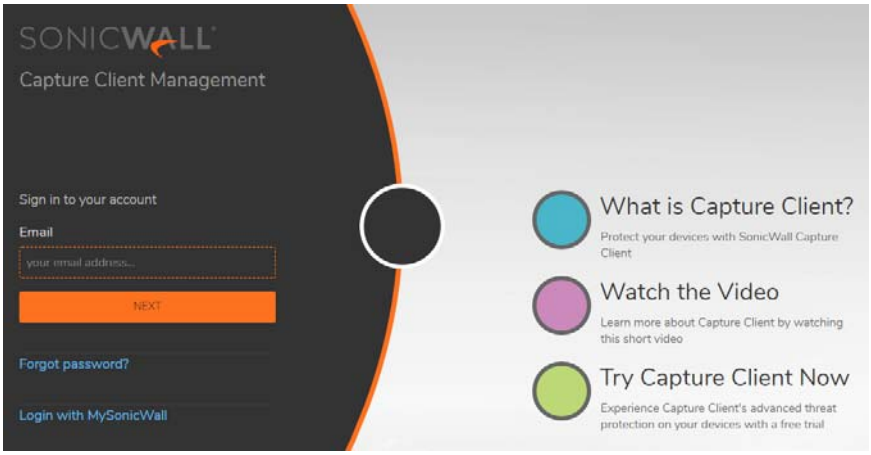

When you log in, the default view is the Dashboard.

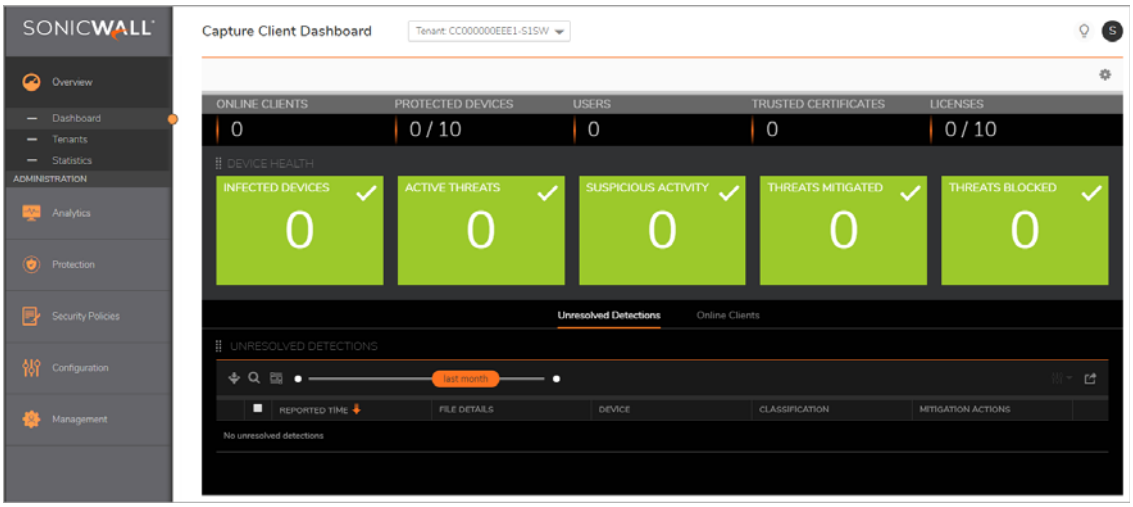

You can see at a glance the number of active endpoints, the number of active users, the status of the policies and the number of active threats. Each item on the dashboard is assigned a color based on the status of the item.

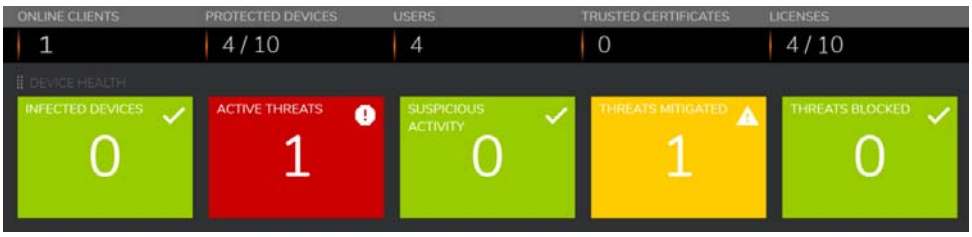

For detailed information about configuring and using Capture Client, refer to the *Capture Client Operations* guide and *Capture Client User Guide*, available on the SonicWall technical documentation portal at https://www.sonicwall.com/support/technical-documentation.

McAfee / SentinelOne Client AV Status

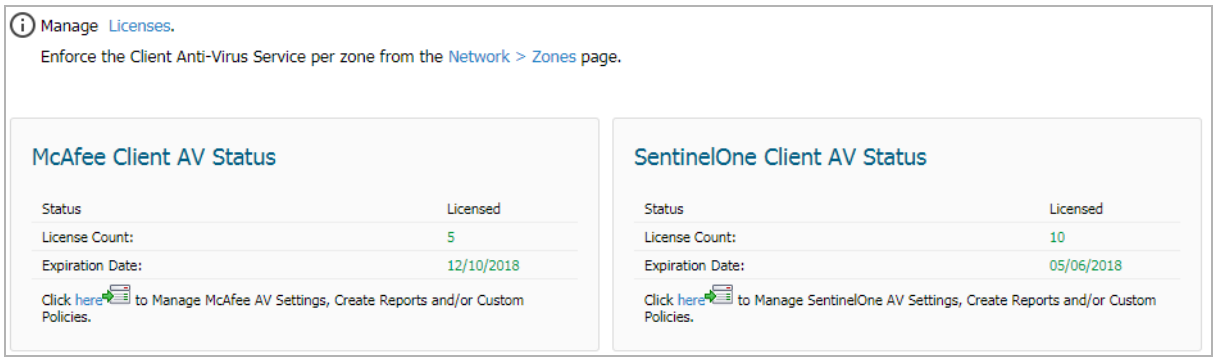

The **Client AV Status** section:

- Displays information about whether the firewall is licensed, the number of licenses, and the date the license expires.
- Contains a link to login to MySonicWall for managing and reviewing detailed system and network information. Clicking this link displays the **Licenses > License Management** page for MySonicWall login.

Client Anti-Virus Policies

Client Anti-Virus Policies - Including Capture Client Settings



The following features are available in the **Client Anti-Virus Policies** section:

- **Disable policing from Trusted to Public** - Cleared, this option enforces anti-virus policies on computers located on Trusted zones. Choosing this option allows computers on a trusted zone (such as a LAN) to access computers on public zones (such as DMZ), even if anti-virus software is not installed on the LAN computers.
- **Switch McAfee AV to Kaspersky AV for clients on Kaspersky enforcement list** - When selected, uses Kaspersky AV for clients on the Kaspersky enforcement list instead of McAfee AV.
- **Days before forcing update** - This feature defines the maximum number of days of access to the Internet before the SonicWall requires the latest virus date files to be downloaded. Select from 0 to 5 days; **5** is the default.
- **Force update on alert** - SonicWall broadcasts virus alerts to all SonicWall appliances with an Anti-Virus subscription. When an alert is received with this option selected, users are upgraded to the latest version of VirusScan before they can access the Internet. This option overrides the maximum number of days allowed before forcing update selection. In addition, every virus alert is logged, and an alert message is sent to the administrator.

  Three levels of alerts are available, and you may select more than one:

  - **Low Risk** - A virus that is not reported in the field and is considered unlikely to be found in the field in the future has a low risk. Even if such a virus includes a very serious or unforeseeable damage payload, its risk is still low. This option is not selected by default.
  - **Medium Risk** - If a virus is found in the field, and if it uses a less common infection mechanism, it is considered to be medium risk. If its prevalence stays low and its payload is not serious, it can be downgraded to a low risk. Similarly it can be upgraded to high risk if the virus becomes more and more widespread. This option is selected by default.
  - **High Risk** - To be assigned a high risk rating, it is necessary that a virus is reported frequently in the field. Additionally, the payload must have the ability to cause at least some serious damage. If it causes very serious or unforeseeable damage, high risk may be assigned even with a lower level of prevalence. This option is selected by default.

Starting in SonicOS 6.5.4.5, the **Client Anti-Virus Policies** section includes two settings that are specific to Capture Client:

- **Enable SSO Login via Capture Client Enforcement** - When selected, this option enables the periodic sharing of user login information (domain/user format) from Capture Client endpoints to SonicWall

firewalls that enforce Capture Client, when there is proper connectivity between the Capture Client endpoints and the Client Management Console (CMC). The firewalls are also notified if user information is changed or updated on the Capture Client endpoint IP address.

User and group information is identified on the firewall based on information from Capture Client. This is supported alongside other methods of obtaining user information, such as the Single Sign-On Agent. If the user information from Capture Client differs from the information provided by the SSO Agent or other authentication methods, the Capture Client user information has the lowest priority.

Traffic from a source IP address can be allowed or denied by the firewall based on the current user information for that IP address.

- **Enable Alert Message from Firewalls to Capture Client Endpoint Devices** - When selected, this option enables push notifications from the firewall to the client when a connection is blocked or traffic is dropped due to actions by other SonicWall security services. Without this option, clients have little visibility into firewall actions taken on traffic other than HTTP/HTTPS. These alert notifications provide a summary of the event containing the following information:
  - Timestamp
  - Source IP/Port
  - Destination IP/Port
  - Category - The SonicWall security service that blocked traffic or dropped the connection.

    The following Categories are included:
    - App Control
    - Botnet
    - Geo-IP Filter
    - Content Filter Service
    - Gateway Anti-Virus
    - Anti-Spyware
    - Capture ATP
  - Message (if available)

  After sending the alert notification to the Capture Client endpoint, the event is logged as usual in the SonicOS event log.

Client Anti-Virus Enforcement

Client Anti-Virus Enforcement for Capture Client

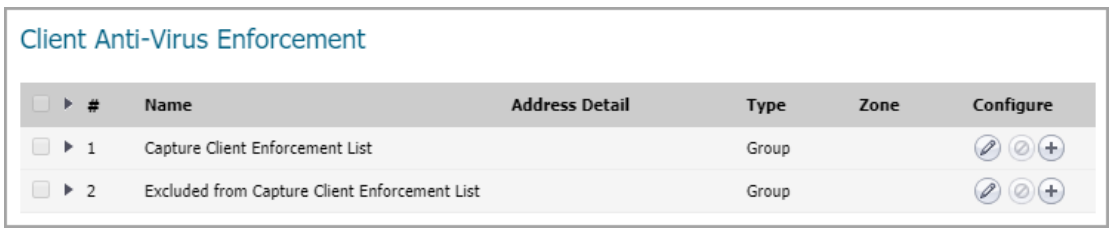

Client Anti-Virus Enforcement for McAfee / SentinelOne

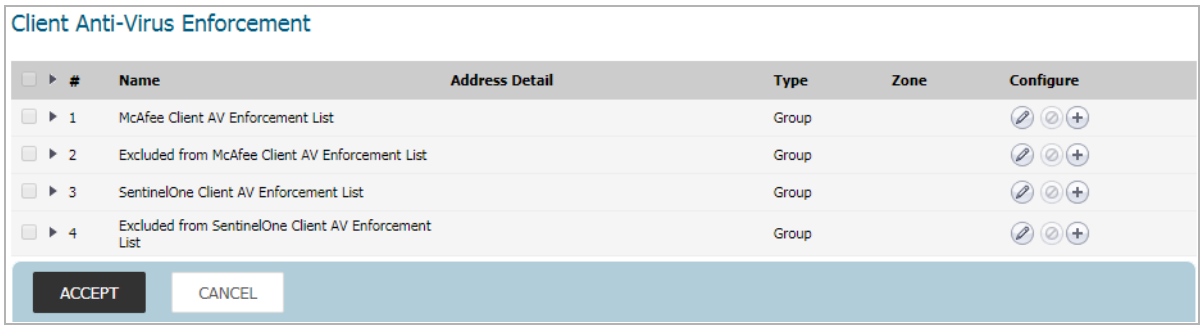

The **Client Anti-Virus Enforcement** table has two entries for each licensed client anti-virus service, both with a **Type** of **Group**:

- ***<Name>* Client AV Enforcement List** (where ***<Name>*** is **Capture Client**, **McAfee** or **SentinelOne**, depending on which you use)
- **Excluded from *<Name>* Client AV Enforcement List**

To see the IP addresses associated with each entry, click the **Expand** icon. The **Address Detail**, **Type**, and **Zone** for each entry displays. If you have not configured the enforcement list, clicking the **Expand** icon displays **No Entries**.

To hide the IP addresses, click the **Collapse** icon.

You can edit or add to these entries, but you cannot delete them.

Topics:

- Creating the Client AV Enforcement List
- Excluding Address Objects from the Client AV Enforcement List
- Protecting Computers Not In Either List

Creating the Client AV Enforcement List

NOTE: Predefined Address Objects, such as interface IPs or the Default Gateway cannot be edited or deleted individually; their **Edit** and **Delete** icons are dimmed. You remove a predefined Address Object from the **Client AV Enforcement List** through editing the List itself. You can, however, edit or delete any Address Object you have defined.

You need to configure the client AV enforcement list with the IP address of the address objects that are to have Client AV enforced.

You can define ranges of IP addresses to receive Anti-Virus enforcement by creating an Address Object containing a range of IP addresses. Any computer requiring enforcement needs a static IP address within the specified range of IP addresses. Up to 64 IP address ranges can be entered for enforcement.

To create the client AV enforcement list from existing Address Objects:

1. Navigate to the **MANAGE | Security Configuration | Security Services > Client AV Enforcement** page.
2. Scroll to the **Client Anti-Virus Enforcement** section.

3 Click the **Edit** icon for the ***<Name>*** **Enforcement List**. The **Edit Address Object Group** dialog displays.

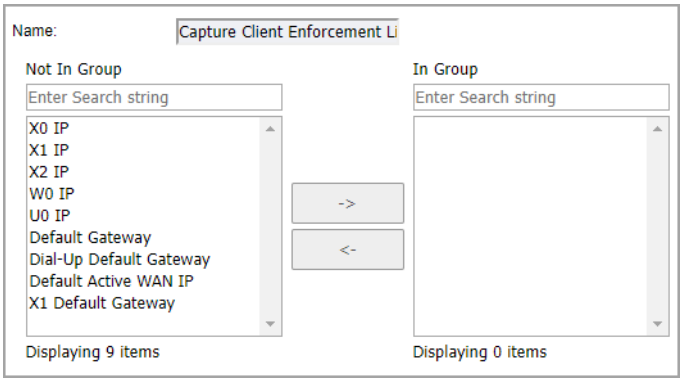

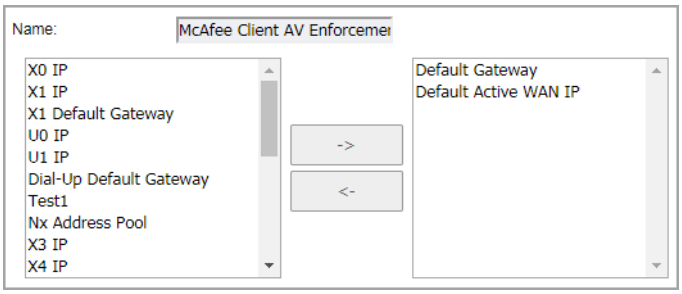

4 Select the IP address(es) to have client AV enforcement from the list on the left.

5 Click the **Right Arrow** button to move the entries to the list on the right.

6 When finished adding Address Objects, click **OK**.

To add an Address Object to the Client AV Enforcement List:

1 Navigate to the **MANAGE | Security Configuration | Security Services > Client AV Enforcement** page.

2 Scroll to the **Client Anti-Virus Enforcement** section.

3 Click the **Add** icon for the ***<Name>*** **Enforcement List**. The **Add Address Object** dialog displays.

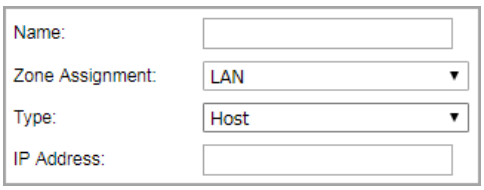

4 Enter a friendly name in the **Name** field.

5 Select the zone from the **Zone Assignment** drop-down menu.

6 Select the type from the **Type** drop-down menu.

7 Enter the IP address of the Address Object in the **IP Address** field.

8 Click **OK**.

Excluding Address Objects from the Client AV Enforcement List

SonicWall Client Anti-Virus currently supports Windows platforms. To access the internet, devices with other operating systems must be exempt from Anti-Virus policies.

CAUTION: To ensure full network protection from virus attacks, it is recommended that only servers and unsupported machines be excluded from protection and that third-party anti-virus software is installed on each machine before excluding that machine from Anti-Virus enforcement.

NOTE: Predefined Address Objects, such as interface IPs or the Default Gateway cannot be edited or deleted individually; their **Edit** and **Delete** icons are dimmed. You remove a predefined Address Object from the **Excluded from Client AV Enforcement List** through editing the List itself. You can, however, edit or delete any Address Object you have defined.

To define excluded Address Objects:

1 Navigate to the **MANAGE | Security Configuration | Security Services > Client AV Enforcement** page.

2 Scroll to the **Client Anti-Virus Enforcement** section.

3 Click the **Edit** icon for the **Excluded from *<Name>* Enforcement List**. The **Edit Address Object Group** displays.

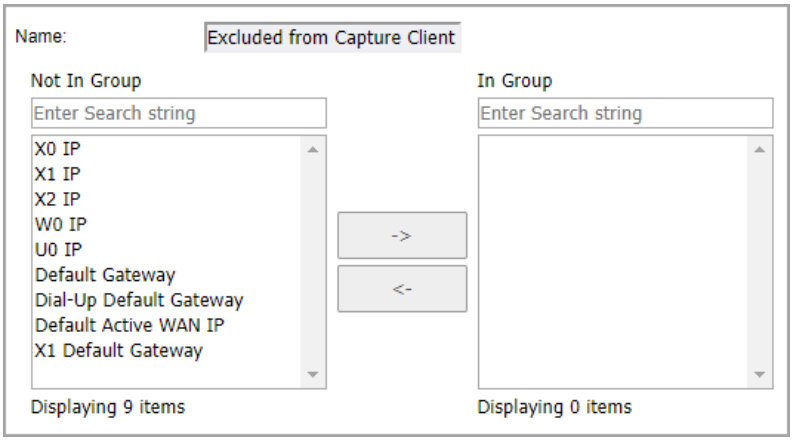

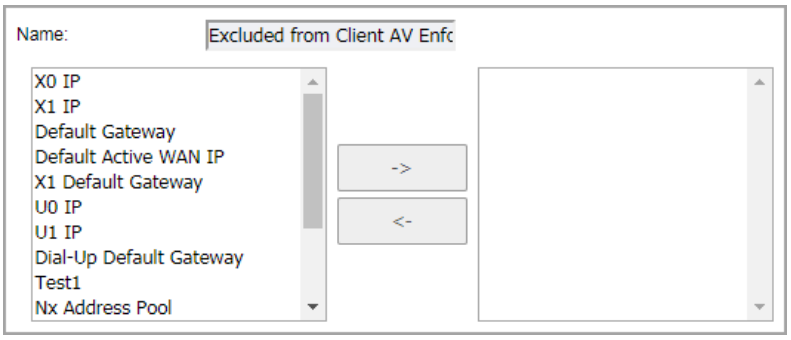

4 Select the Address Object(s) to be excluded from the list on the left.

5 Click the **Right Arrow** to move the objects to the list on the right.

6 When finished excluding Address Objects, click **OK**.

To add an Address Object to the Excluded from Client AV Enforcement List:

1 Scroll to the client **Anti-Virus Enforcement** section.

2 Click the **Add** icon for the **Excluded from *<Name>* Enforcement List**. The **Add Address Object** dialog displays.

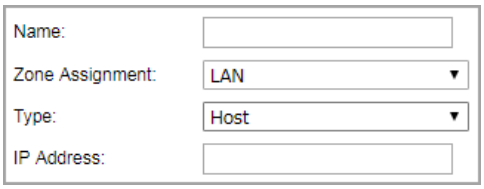

3 Enter a friendly name in the **Name** field.

4 Select the zone from the **Zone Assignment** drop-down menu.

5 Select the type from the **Type** drop-down menu.

6 Enter the IP address of the Address Object in the **IP Address** field.

7 Click **OK**.

Protecting Computers Not In Either List

For those computers not included in either enforcement list, you can specify the type of default enforcement to be applied to them.

NOTE: This feature is deprecated in SonicOS 6.5.1, and is no longer available in that and later versions.

To specify a default enforcement to computers not in an enforcement list:

1 Scroll to the client **Anti-Virus Enforcement** section.

2 Scroll to the bottom of the **Security Services > Client AV Enforcement** page.

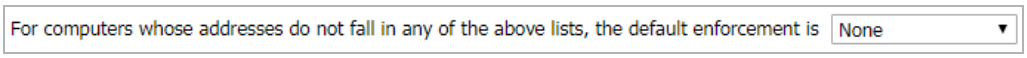

3 Select the type of default enforcement from the **For computers whose addresses do not fall in any of the above lists, the default enforcement is** drop-down menu:

- **None** (default)
- Third-party anti-virus program (McAfee or Kaspersky, depending on your system)

11

DPI-SSL Enforcement

TIP: For information about DPI-SSL, see About DPI-SSL on page 244.

- About DPI-SSL Enforcement on page 152
- Managing DPI-SSL Enforcement on page 153

About DPI-SSL Enforcement

When you enable the DPI-SSL services on your SonicWall network security appliance, the clients behind the firewalls that have no related certificates are often required to confirm by going through HTTPS web pages. Otherwise, users have to install the corresponding DPI-SSL certificates manually if they want to bypass this step. by downloading the corresponding certificate, and then installing it.

To simplify the procedure so that clients can download and install the certificates automatically, DPI-SSL enforcement is necessary.

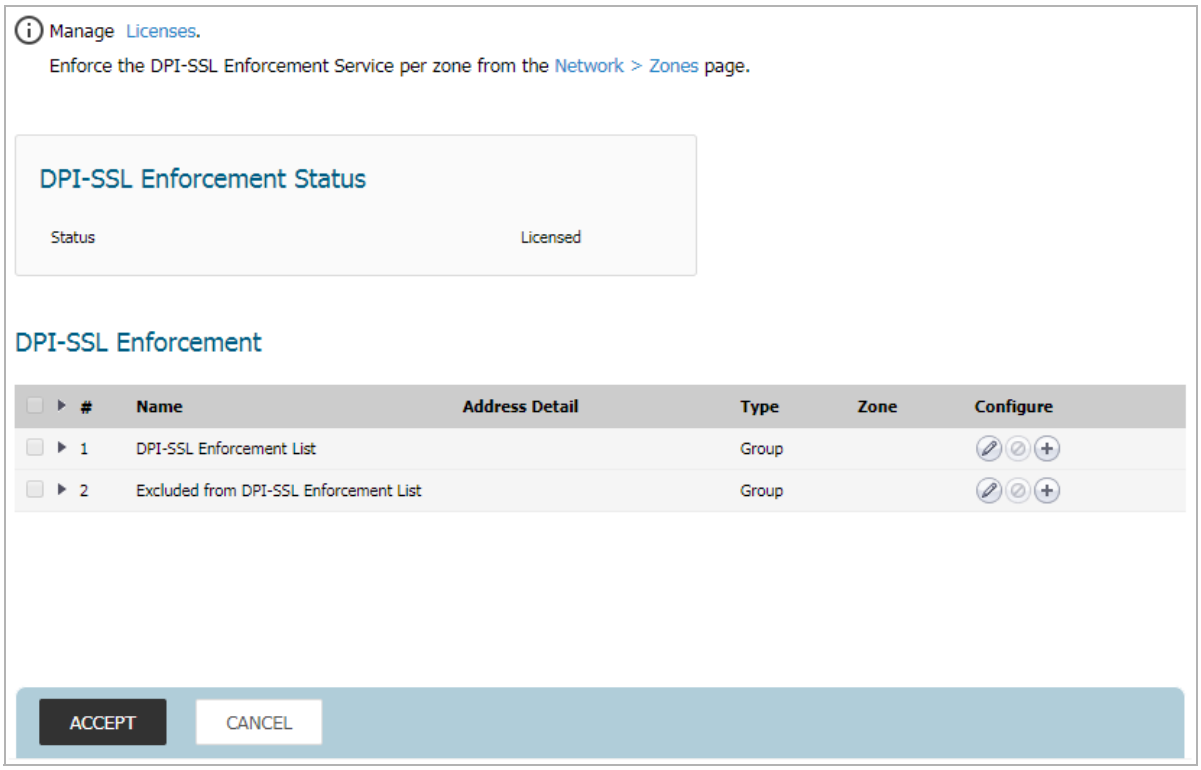

Topics:

- Links on page 153

- DPI-SSL Enforcement Status on page 153
- DPI-SSL Enforcement on page 153

Links

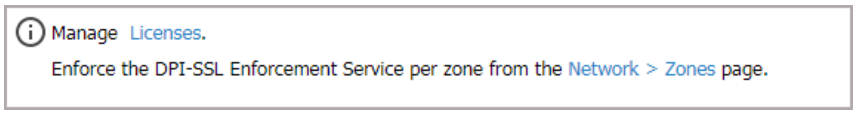

The top of the **Security Services > DPI-SSL Enforecement** page displays links for:

- Viewing and managing licenses.
- Displaying the **MANAGE | System Setup > Network > Zones** page where you can configure DPI-SSL Enforcement Service per zone.

DPI-SSL Enforcement Status

The **DPI-SSL Enforcement Status** section shows the licensing status of the DPI-SSL Enforcement Status feature.

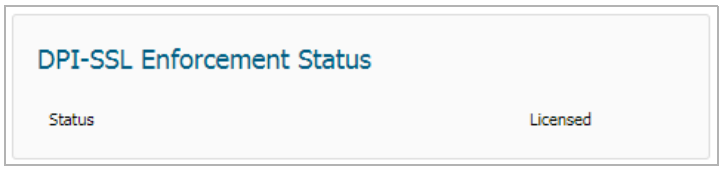

DPI-SSL Enforcement

The **DPI-SSL Enforcement** section contains the lists of addresses included in and excluded from DPI-SSL enforcement.

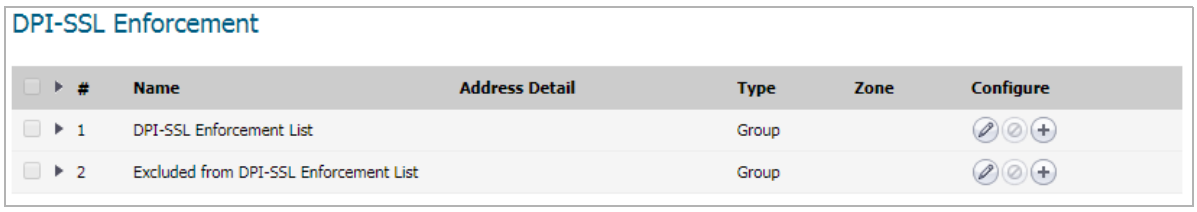

Managing DPI-SSL Enforcement

On the **MANAGE | Security Configuration > Security Services > DPI-SSL Enforecement** page, you can add, edit, and delete items on:

- DPI-SSL Enforcement List
- Excluded from DPI-SSL Enforcement List

Topics:

- Editing a DPI-SSL Enforcement List on page 154

- Adding a Policy to a DPI-SSL Enforcement List on page 154
- Editing a DPI-SSL Enforcement Policy on page 154
- Managing Zones for DPI-SSL Enforcement on page 155

Editing a DPI-SSL Enforcement List

To edit a DPI-SSL enforcement list:

1. Navigate to the **MANAGE | Security Configuration > Security Services > DPI-SSL Enforecement** page.
2. Scroll to the **DPI-SSL Enforcement** section.
3. Click the **Edit** icon next to the list to which you want to edit. The **Edit Address Object Group** dialog displays.
4. Select the address objects to be added from the left column. Multiple address objects can be selected at one time.
5. Click the **Right Arrow** button.

   To delete an address object from the group, select the address object and click the **Left Arrow** button.
6. Click **OK**.

Adding a Policy to a DPI-SSL Enforcement List

To add new policies to DPI-SSL enforcement list:

1. Navigate to the **MANAGE | Security Configuration > Security Services > DPI-SSL Enforecement** page.
2. Scroll to the **DPI-SSL Enforcement** section.
3. Click the **Add** icon next to the list to which you want to add a policy. The **Add Address Object** dialog displays.
4. Enter a friendly name for the server in the **Name** field.
5. From **Zone Assignment**, select the server's zone.
6. Select the type of host from **Type**. The following setting(s) change, depending on the host type selected.
7. If you selected:
   - **Host** (default) – Enter the IP address in the **IP Address** field.
   - **Range** – Enter the starting and ending IP addresses in the **Starting IP Address** and **Ending IP Address** fields.
8. Click **OK**.

Editing a DPI-SSL Enforcement Policy

To edit a DPI-SSL enforcement policy:

1. Navigate to the **MANAGE | Security Configuration > Security Services > DPI-SSL Enforecement** page.
2. Scroll to the **DPI-SSL Enforcement** section.

3 Cick the **Edit** icon next to the a policy you want to edit. The **Edit Address Object** dialog displays.

4 Update the values you want to change.

5 Click **OK**.

Managing Zones for DPI-SSL Enforcement

Use the **MANAGE | System Setup > Network > Zones** page to manage DPI-SSL enforcement for specific zones. For more information about zones, see *SonicWall SonicOS 6.5 System Setup*.

12

Configuring Client CF Enforcement

- Security Services > Client CF Enforcement
- Enabling and Configuring Client CF Enforcement
- Enabling Client CFS in Network Zones

Security Services > Client CF Enforcement

SonicWall Client CF Enforcement provides protection and productivity policy enforcement for businesses, schools, libraries and government agencies. SonicWall has created a revolutionary content filtering architecture, utilizing a scalable, dynamic database to block objectionable and unproductive Web content.

Client CF Enforcement provides the ideal combination of control and flexibility to ensure the highest levels of protection and productivity. Client CF Enforcement prevents individual users from accessing inappropriate content while reducing organizational liability and increasing productivity. Web sites are rated according to the type of content they contain. The Content Filtering Service (CFS) blocks or allows access to these web sites based on their ratings and the policy settings for a user or group.

Businesses can typically control web surfing behavior and content when the browsing is initiated within the perimeter of the security appliance by setting filter policies on the appliance. But when the same device exits the perimeter, the control is lost. Client CF Enforcement kicks into action to address this gap, by blocking objectionable and unproductive Web content outside the security appliance perimeter.

SonicWall security appliances working in conjunction with Client CF Enforcement automatically and consistently ensure all endpoints have the latest software updates for the ultimate network protection. The client is designed to work with both Windows and Mac PCs.

Client CF Enforcement consists of the following three main components:

- A Network Security Appliance running SonicOS whose role is to facilitate and verify licencing of CFS and to enable or disable enforcement and configure exclusions and other settings.
- Automatic triggering to install the Client CF Enforcement of any client attempting to access the Internet without the client software installed will be blocked from accessing Websites until it is installed.
- Administration of client policies and client groups using the cloud-based EPRS server accessed from MySonicWall or from SonicOS running on the appliance.

Topics:

- Enabling and Configuring Client CF Enforcement
- Enabling Client CFS in Network Zones

Enabling and Configuring Client CF Enforcement

This section describes how to enable and configure settings for Client CF Enforcement in SonicOS.

Client CF Enforcement must be enabled on the SonicWall appliance before users will be presented with a Website block page, which prompts the user to install the Client CF Enforcement.

NOTE: If the Content Filtering Client (CFS) is not activated on MySonicWall, you must activate it to enforce client content filtering policies on client systems.

Configuring Client CF Enforcement in Security Services

To configure settings for Client CF Enforcement:

1 Navigate to the **MANAGE | Security Configuration > Client CF Enforcement** page.

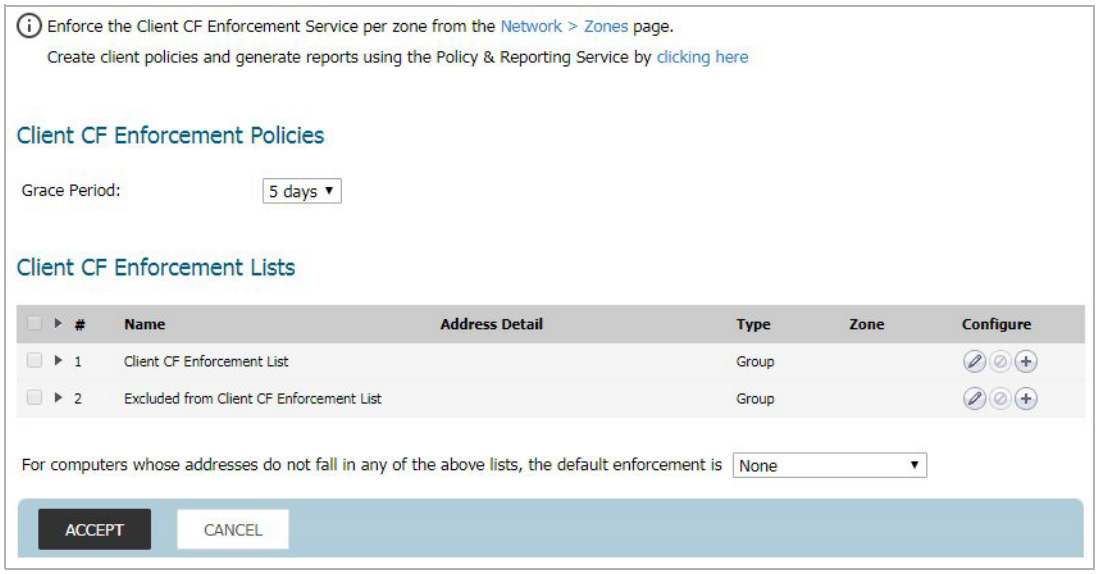

2 Under the **Client CF Enforcement Policies** section, select the number of days from the drop-down list for the **Grace Period** during which CFS enforcement policies remain valid.

The **Client CF Enforcement Lists** section contains a table including the Client CFS Enforcement List and the Excluded from Client CF Enforcement List.

To configure either of these tables, click the **Configure** icon for the list you wish to configure. The Edit Address Object Group dialog displays. Select from the available list the values to include/not include for the group.

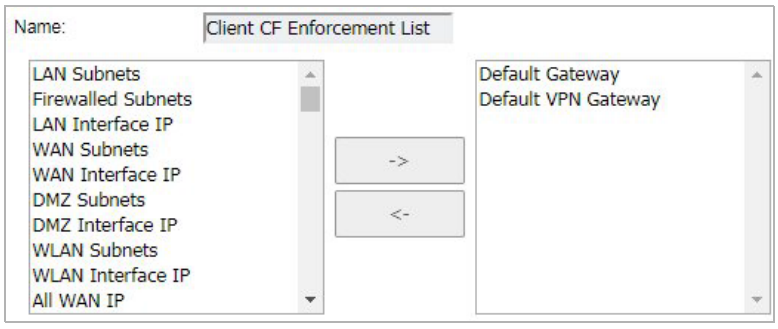

3 For the **Client CF Enforcement List** and **Excluded from Client CF Enforcement List**. If you have made any entries in these lists, you can click the arrow next to the list title to display the entries. To add entries to either list, click the Configure icon in that row.

4 For the field labeled **For computers whose addresses do not fall in any of the above lists, the default enforcement is**, select **Client CF Enforcement** from the drop-down list. This is located below the **Client CF**

Enforcement Lists section. Selecting this will prompt all other computers connecting to the Internet through the appliance to install the Enforced Client. You can select **None** from the drop-down list if you only want to enforce the service on computers that you have configured.

5 Click **ACCEPT**.

Enabling Client CFS in Network Zones

Client Content Filtering is enforced on a per-zone basis.

To enforce CFS on a per-zone basis:

1 At the top of the **Security Services > Client CF Enforcement** page, click the **Network > Zones** link in the **Note**.

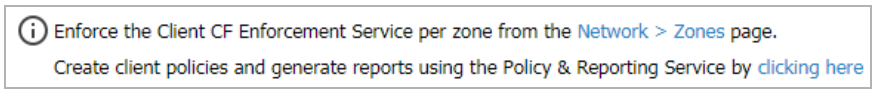

The **Network > Zones** page displays.

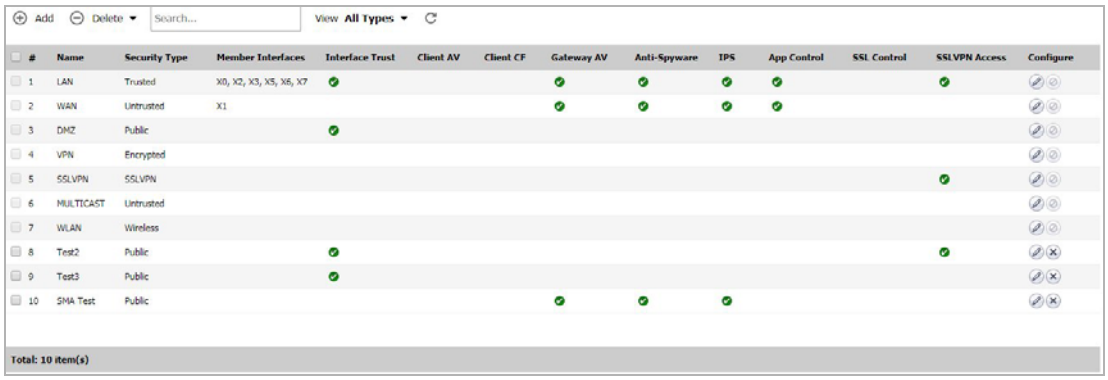

2 Click the **Configure** button for the zone on which you want to enforce the Client Content Filtering Service. The **Add Zone** dialog appears.

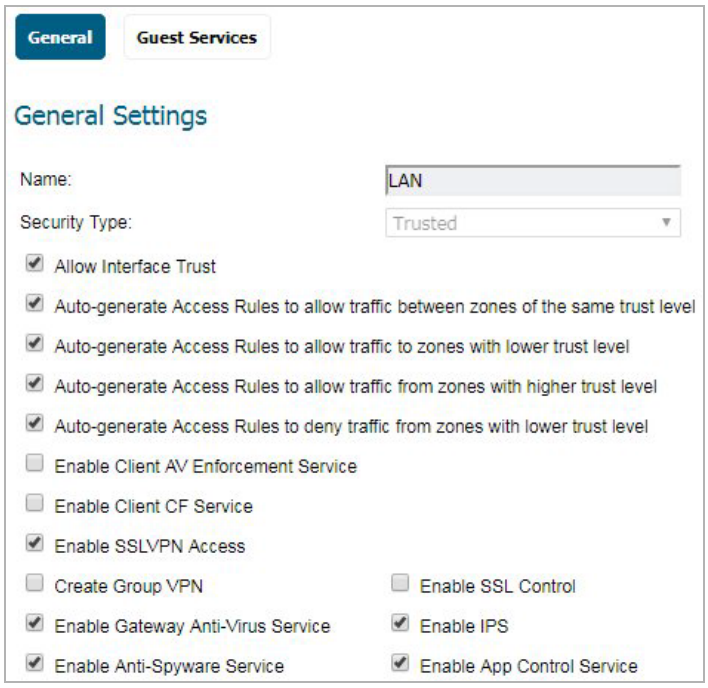

3 Select the **Enable Client CF Service** checkbox.

4 Click **OK**.

13

Managing SonicWall Gateway Anti-Virus Service

- About SonicWall Gateway Anti-Virus Service on page 160
- Setting Up SonicWall Gateway Anti-Virus Protection on page 165
- Viewing SonicWall GAV Signatures on page 175

About SonicWall Gateway Anti-Virus Service

SonicWall Gateway Anti-Virus (GAV) service delivers real-time virus protection directly on the SonicWall security appliance by using SonicWall's IPS-Deep Packet Inspection v2.0 engine to inspect all traffic that traverses the SonicWall gateway. Building on SonicWall's reassembly-free architecture, SonicWall GAV inspects multiple application protocols, as well as generic TCP streams, and compressed traffic. Because SonicWall GAV does not have to perform reassembly, there are no file-size limitations imposed by the scanning engine. Base64 decoding, ZIP, LHZ, and GZIP (LZ77) decompression are also performed on a single-pass, per-packet basis.

SonicWall GAV delivers threat protection by matching downloaded or emailed files against an extensive and dynamically updated database of threat virus signatures. Virus attacks are caught and suppressed before they travel to desktops. New signatures are created and added to the database by a combination of SonicWall's SonicAlert Team, third-party virus analysts, open source developers, and other sources.

SonicWall GAV can be configured to protect against internal threats as well as those originating outside the network. It operates over a multitude of protocols including SMTP, POP3, IMAP, HTTP, FTP, NetBIOS, instant messaging and peer-to-peer applications, and dozens of other stream-based protocols, to provide you with comprehensive network threat prevention and control. Because files containing malicious code and viruses can also be compressed and therefore inaccessible to conventional anti-virus solutions, SonicWall GAV integrates advanced decompression technology that automatically decompresses and scans files on a per-packet basis.

SonicWall GAV parses supported email protocols for the header fields `to`, `cc`, and `bcc`. The information in these fields are displayed and logged in Capture ATP for both sender and receiver.

Topics:

- SonicWall GAV Multi-Layered Approach
- SonicWall GAV Architecture
- Activating the Gateway Anti-Virus, Anti-Spyware, and IPS License
- Setting Up SonicWall Gateway Anti-Virus Protection
- Viewing SonicWall GAV Signatures

SonicWall GAV Multi-Layered Approach

SonicWall GAV delivers comprehensive, multi-layered anti-virus protection for networks at the desktop, the network, and at remote sites; see SonicWall GAV multi-layer approach. SonicWall GAV enforces anti-virus policies at the gateway to ensure all users have the latest updates and monitors files as they come into the network.

SonicWall GAV multi-layer approach

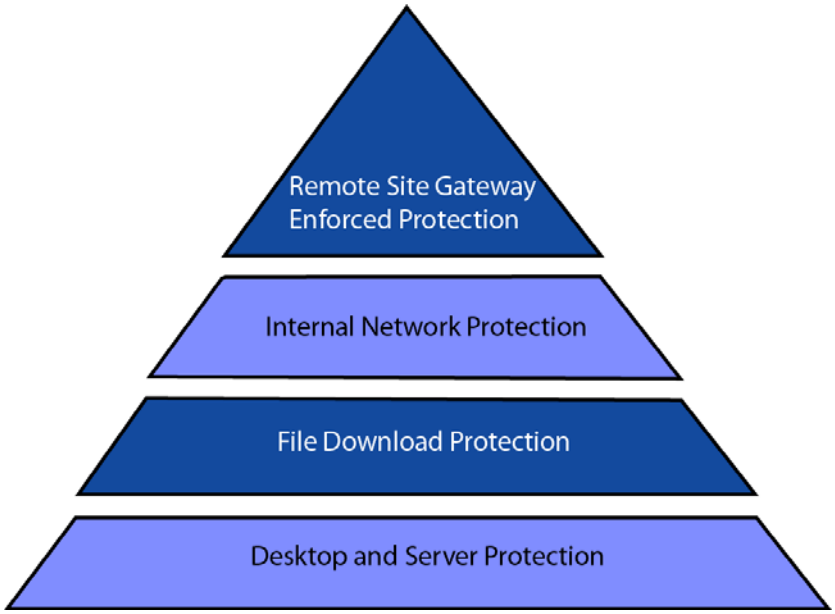

Topics:

- Remote Site Protection
- Internal Network Protection
- HTTP File Downloads
- Server Protection
- Cloud Anti-Virus Database

Remote Site Protection

Remove site protection

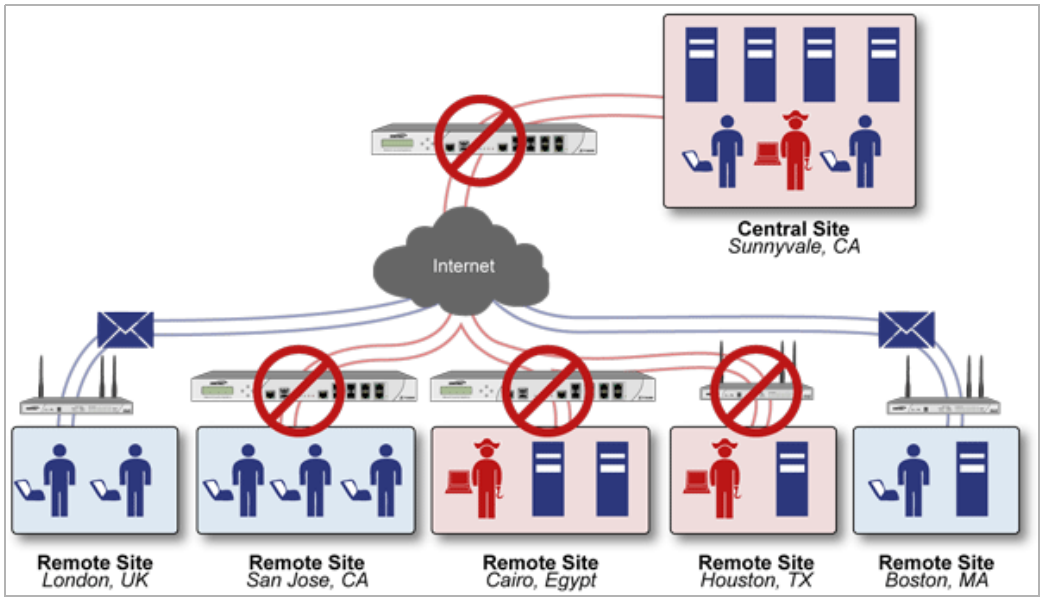

1 Users send typical e-mail and files between remote sites and the corporate office.

2 SonicWall GAV scans and analyses files and e-mail messages on the SonicWall security appliance.

3 Viruses are found and blocked before infecting remote desktop.

4 The virus is logged, and an alert is sent to the administrator.

Internal Network Protection

Internal network protection

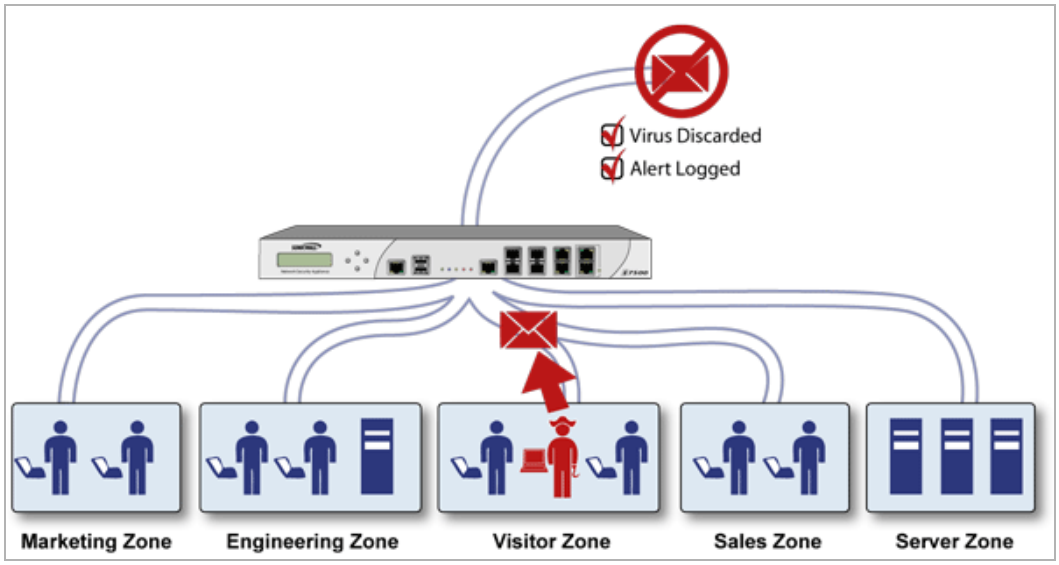

1 Internal user contracts a virus and releases it internally.

2 All files are scanned at the gateway before being received by other network users.

3 If a virus is found, the file is discarded.

4 The virus is logged, and an alert is sent to the administrator.

HTTP File Downloads

HTTP file downloads

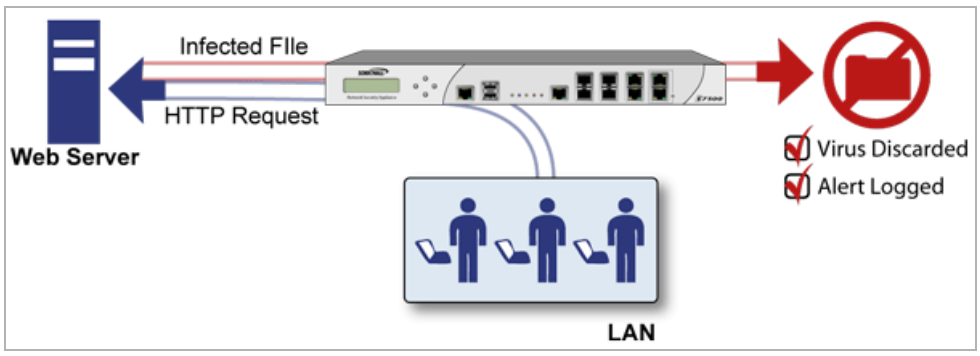

1 Client makes a request to download a file from the Web.

2 The file is downloaded through the Internet.

3 The file is analyzed the SonicWall GAV engine for malicious code and viruses.

4 If a virus is found, the file 8is discarded.

5 The virus is logged, and an alert is sent to the administrator.

Server Protection

Server protection

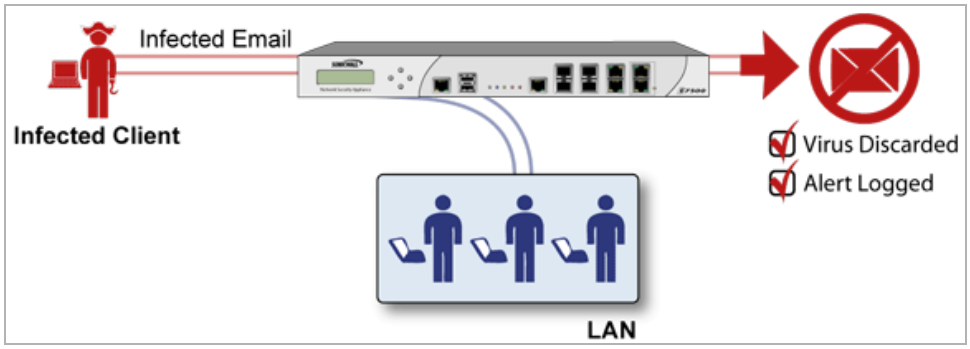

1 Outside user sends an incoming email.

2 The email is analyzed by the SonicWall GAV engine for malicious code and viruses before being received by the email server.

3 If a virus is found, the threat is prevented.

4 The email is returned to the sender, the virus is logged, and an alert sent to the administrator.

Cloud Anti-Virus Database

The Cloud Gateway Anti-Virus feature introduces an advanced malware scanning solution that compliments and extends the existing Gateway Anti-Virus scanning mechanisms present on SonicWall firewalls to counter the continued growth in the number of malware samples in the wild.

Cloud Gateway Anti-Virus expands the Reassembly Free Deep Packet Inspection engine capabilities by consulting with the datacenter-based malware analysis servers. This approach keeps the foundation of RFDPI-based malware detection by providing a low-latency, real-time solution that is capable of scanning unlimited numbers of files of unlimited size on all protocols that are presently supported without adding any significant incremental processing overhead to the appliances themselves. With this additional layer of security, SonicWall's Next Generation Firewalls are able to extend their current protection to cover multiple millions of pieces of malware.

SonicWall GAV Architecture

SonicWall GAV is based on SonicWall's high performance DPIv2.0 engine (Deep Packet Inspection version 2.0) engine, which performs all scanning directly on the SonicWall security appliance. SonicWall GAV includes advanced decompression technology that can automatically decompress and scan files on a per-packet basis to search for viruses and malware; see SonicWall GAV architecture. The SonicWall GAV engine can perform base64 decoding without ever reassembling the entire base64 encoded mail stream. Because SonicWall's GAV does not have to perform reassembly, there are no file-size limitations imposed by the scanning engine. Base64 decoding and ZIP, LHZ, and GZIP (LZ77) decompression are also performed on a single-pass, per-packet basis. Reassembly free virus scanning functionality of the SonicWall GAV engine is inherited from the Deep Packet Inspection engine, which is capable of scanning streams without ever buffering any of the bytes within the stream.

SonicWall GAV architecture

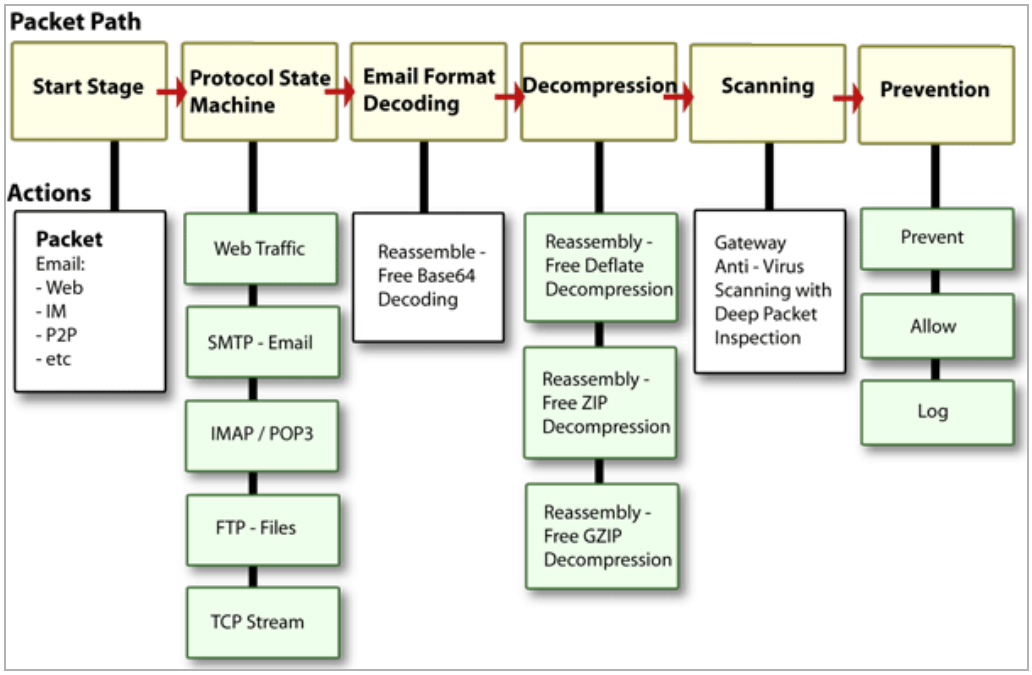

Building on SonicWall's reassembly-free architecture, GAV has the ability to inspect multiple application protocols, as well as generic TCP streams, and compressed traffic. SonicWall GAV protocol inspection is based on high performance state machines which are specific to each supported protocol. SonicWall GAV delivers protection by inspecting over the most common protocols used in today's networked environments, including SMTP, POP3, IMAP, HTTP, FTP, NetBIOS, instant messaging and peer-to-peer applications and dozens of other

stream-based protocols. This closes potential backdoors that can be used to compromise the network while also improving employee productivity and conserving Internet bandwidth.

TIP: If your SonicWall security appliance is connected to the Internet and registered at mySonicWall.com, you can activate a 30-day FREE TRIAL of SonicWall Gateway Anti-Virus, SonicWall Anti-Virus, and SonicWall Intrusion Prevention Service separately from the **Security Services > Gateway Anti-Virus, Security Services > Anti-Spyware**, and **Security Services > Intrusion Prevention** pages in the management interface.

Activating the Gateway Anti-Virus, Anti-Spyware, and IPS License

Your appliance must be registered on MySonicWall to use these security services. See your *Quick Start Guide* for information on creating a MySonicWall account and registering your appliance. For information about upgrading the services in a closed environment, see *SonicWall SonicOS 6.5 Updates*.

Because SonicWall Anti-Spyware is part of SonicWall Gateway Anti-Virus, Anti-Spyware, and Intrusion Prevention Service, the Activation Key you receive is for all three services on your SonicWall security appliance.

If you do not have a SonicWall Gateway Anti-Virus, Anti-Spyware, and Intrusion Prevention Service. license activated on your SonicWall security appliance, you must purchase it from a SonicWall reseller or through your mySonicWall.com account (limited to customers in the USA and Canada).

Activating FREE TRIAL Versions

You can try FREE TRIAL versions of SonicWall Gateway Anti-Virus, SonicWall Anti-Spyware, and SonicWall Intrusion Prevention Service. For information about activating a free trial of any or all of the Security Services, see the *Quick Start Guide* for your appliance.

Setting Up SonicWall Gateway Anti-Virus Protection

Activating the SonicWall Gateway Anti-Virus license on your SonicWall security appliance does not automatically enable the protection.

To configure SonicWall Gateway Anti-Virus:

1. Enable SonicWall Gateway Anti-Virus.
2. Apply SonicWall Gateway Anti-Virus Protection to zones.

Topics:

- Security Services > Gateway Anti-Virus Page
- Enabling SonicWall GAV
- Applying SonicWall GAV Protection on Zones
- Viewing SonicWall GAV Status Information
- Specifying Protocol Filtering

- Configuring Gateway AV Settings
- Configuring Cloud Gateway AV

Security Services > Gateway Anti-Virus Page

The **Security Services > Gateway Anti-Virus** page provides the settings for configuring SonicWall GAV on your SonicWall security appliance as well as displays both the anti-virus status and the anti-virus signatures.

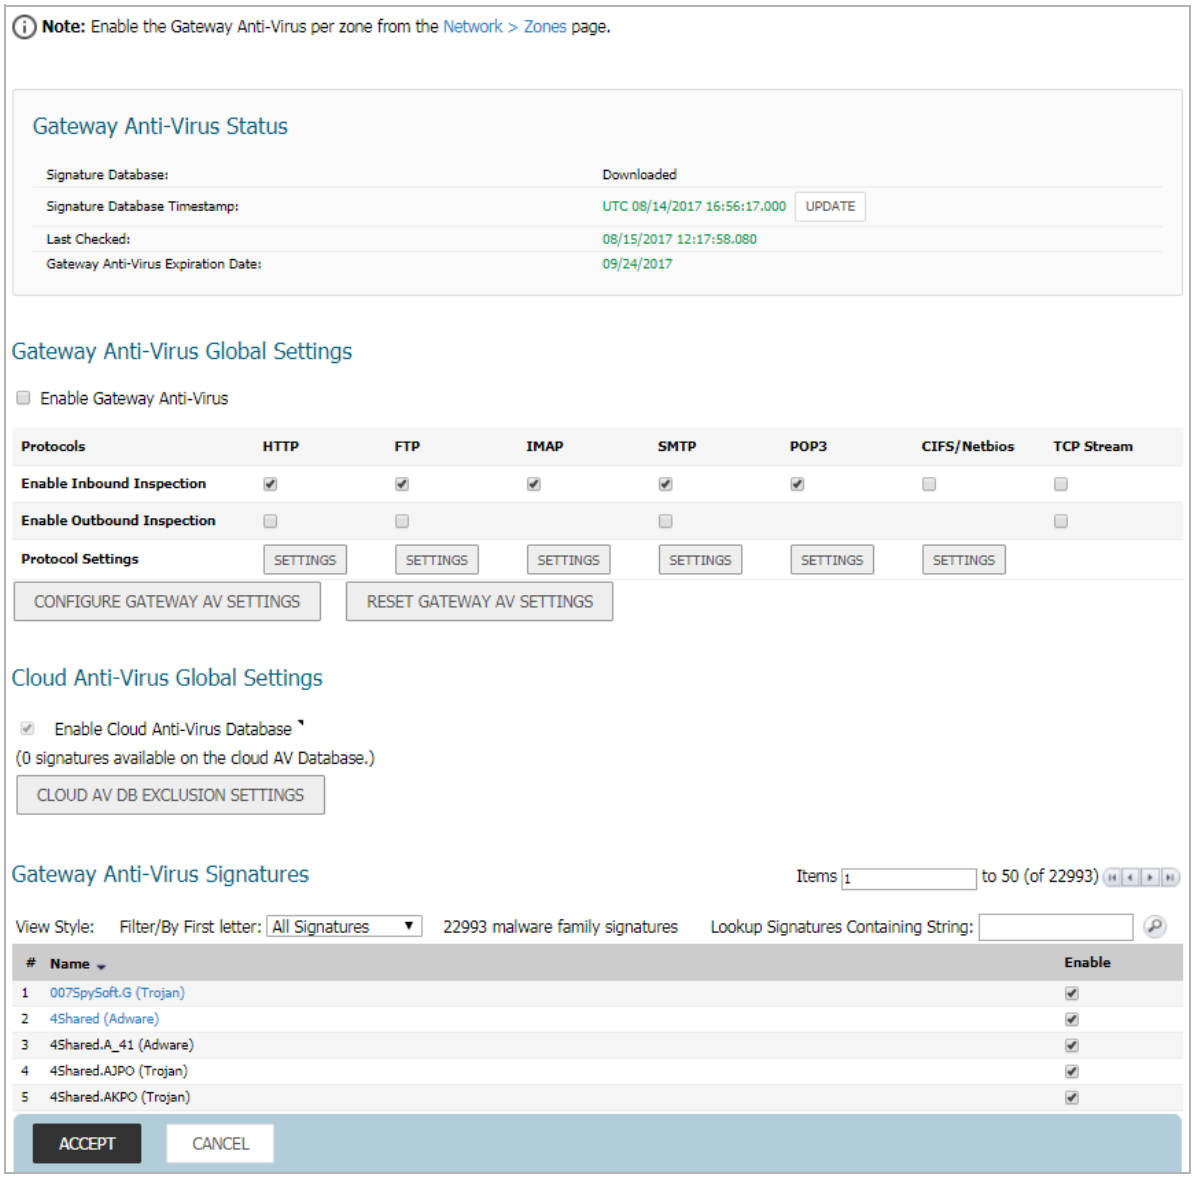

Enabling SonicWall GAV

You must select the **Enable Gateway Anti-Virus** checkbox in the **Gateway Anti-Virus Global Settings** section to enable SonicWall GAV on your SonicWall security appliance.



You must specify the zones you want SonicWall GAV protection on the **System Setup | Network > Zones** page.

Applying SonicWall GAV Protection on Zones

You apply SonicWall GAV to zones when you add or edit a zone on the **Network > Zones** page. From the **Security Services > Gateway Anti-Virus** page, you can quickly display the **Network > Zones** page by clicking the link in the `Note: Enable the Gateway Anti-Virus per zone from the Network > Zones page.` in the **Gateway Anti-Virus Status** section.

Viewing SonicWall GAV Status Information

The **Gateway Anti-Virus Status** section shows the state of the anti-virus signature database, including the database's timestamp, and the time the SonicWall signature servers were last checked for the most current database version. The SonicWall security appliance automatically attempts to synchronize the database on startup, and once every hour.



Topics:

- Checking the SonicWall GAV Signature Database Status
- Updating SonicWall GAV Signatures

Checking the SonicWall GAV Signature Database Status

The **Gateway Anti-Virus Status** section displays the following information:

- **Signature Database** indicates whether the signature database needs to be downloaded or has been downloaded.
- **Signature Database Timestamp** displays the last update to the SonicWall GAV signature database, not the last update to your SonicWall security appliance.
- **Last Checked** indicates the last time the SonicWall security appliance checked the signature database for updates. The SonicWall security appliance automatically attempts to synchronize the database on startup, and once every hour.
- **Gateway Anti-Virus Expiration Date** indicates the date when the SonicWall GAV service expires. If your SonicWall GAV subscription expires, the SonicWall IPS inspection is stopped and the SonicWall GAV configuration settings are removed from the SonicWall security appliance. These settings are automatically restored after renewing your SonicWall GAV license to the previously configured state.

The **Gateway Anti-Virus Status** section displays `Note: Enable the Gateway Anti-Virus per zone from the Network > Zones page.` Clicking on the `Network > Zones` link displays the **Network > Zones** page for applying SonicWall GAV on zones.

Updating SonicWall GAV Signatures

By default, the SonicWall security appliance running SonicWall GAV automatically checks the SonicWall signature servers once an hour. There is no need for an administrator to constantly check for new signature updates. You can also manually update your SonicWall GAV database at any time by clicking the **Update** button located in the **Gateway Anti-Virus Status** section.

SonicWall GAV signature updates are secured. The SonicWall security appliance must first authenticate itself with a pre-shared secret, created during the SonicWall Distributed Enforcement Architecture licensing registration. The signature request is transported through HTTPS, along with full server certificate verification.

Specifying Protocol Filtering

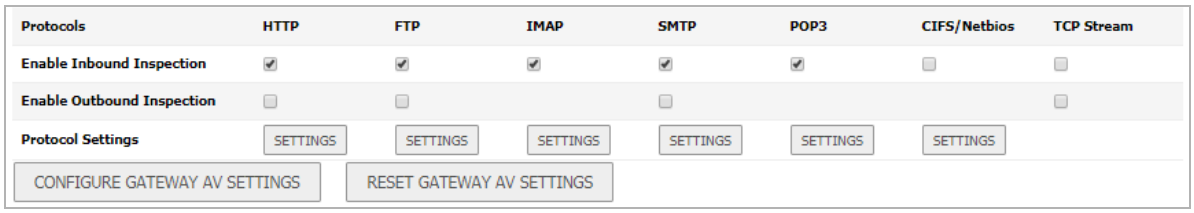

Application-level awareness of the type of protocol that is transporting the violation allows SonicWall GAV to perform specific actions within the context of the application to gracefully handle the rejection of the payload.

Topics:

- Enabling Inbound Inspection
- Enabling Outbound Inspection
- Restricting File Transfers
- Resetting Gateway AV Settings

Enabling Inbound Inspection

By default, SonicWall GAV inspects all inbound **HTTP, FTP, IMAP, SMTP** and **POP3** traffic. Generic **TCP Stream** can optionally be enabled to inspect all other TCP based traffic, such as non-standard ports of operation for SMTP and POP3, and IM and P2P protocols.

Within the context of SonicWall GAV, the **Enable Inbound Inspection** protocol traffic handling refers to the following; see Inspection of inbound traffic: SMTP vs. all other traffic:

- Non-SMTP traffic initiating from a Trusted, Wireless, or Encrypted zone destined to any zone.
- Non-SMTP traffic from a Public zone destined to an Untrusted zone.
- SMTP traffic initiating from a non-Trusted zone destined to a Trusted, Wireless, Encrypted, or Public zone.
- SMTP traffic initiating from a Trusted, Wireless, or Encrypted zone destined to a Trusted, Wireless, or Encrypted zone.

Inspection of inbound traffic: SMTP *vs.* all other traffic

SMTP traffic

To / From	Trusted	Encrypted	Wireless	Public	Untrusted
Trusted	√	√	√		
Encrypted	√	√	√		
Wireless	√	√	√		
Public	√	√	√	√	√
Untrusted	√	√	√	√	√

All other traffic

To / From	Trusted	Encrypted	Wireless	Public	Untrusted
Trusted	√	√	√	√	√
Encrypted	√	√	√	√	√
Wireless	√	√	√	√	√
Public					√
Untrusted					

Enabling Outbound Inspection

The **Enable Outbound Inspection** feature is available for HTTP, FTP, SMTP, and TCP traffic.

Restricting File Transfers

For each protocol, except TCP Stream, you can restrict the transfer of files with specific attributes by clicking on the **Settings** button under the protocol in the **Gateway Anti-Virus Global Settings** section.

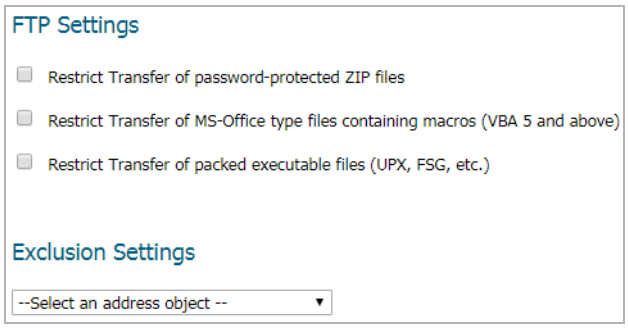

Topics:

- FTP Settings
- Exclusion Settings

FTP Settings

These restrict-transfer **FTP Settings** include:

- **Restrict Transfer of password-protected Zip files** - Disables the transfer of password protected ZIP files over any enabled protocol. This option only functions on protocols (for example, HTTP, FTP, SMTP) that are enabled for inspection.
- **Restrict Transfer of MS-Office type files containing macros (VBA 5 and above)** - Disables the transfers of any MS Office 97 and above files that contain VBA macros.
- **Restrict Transfer of packed executable files (UPX, FSG, etc.)** - Disables the transfer of packed executable files.

  Packers are utilities that compress and sometimes encrypt executables. Although there are legitimate applications for these, they are also sometimes used with the intent of obfuscation, so as to make the executables less detectable by anti-virus applications. The packer adds a header that expands the file in memory, and then executes that file.

  SonicWall Gateway Anti-Virus currently recognizes the most common packed formats: UPX, FSG, PKLite32, Petite, and ASPack. Additional formats are dynamically added along with SonicWall GAV signature updates.

Exclusion Settings

- Drop-down menu – Excludes the selected address object from the restrict-transfer FTP settings.

Resetting Gateway AV Settings

1 To reset all Gateway Anti-Virus (AV) settings to factory default values, click the **Reset Gateway AV Settings** button. A confirmation message displays.

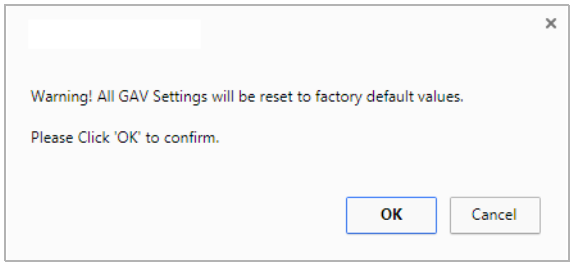

2 Click **OK**.

Configuring Gateway AV Settings

Clicking the **Configure Gateway AV Settings** button at the bottom of the **Gateway Anti-Virus Global Settings** section displays the **Gateway AV Configuration View** dialog, which allows you to configure clientless notification alerts and create a SonicWall GAV exclusion list.

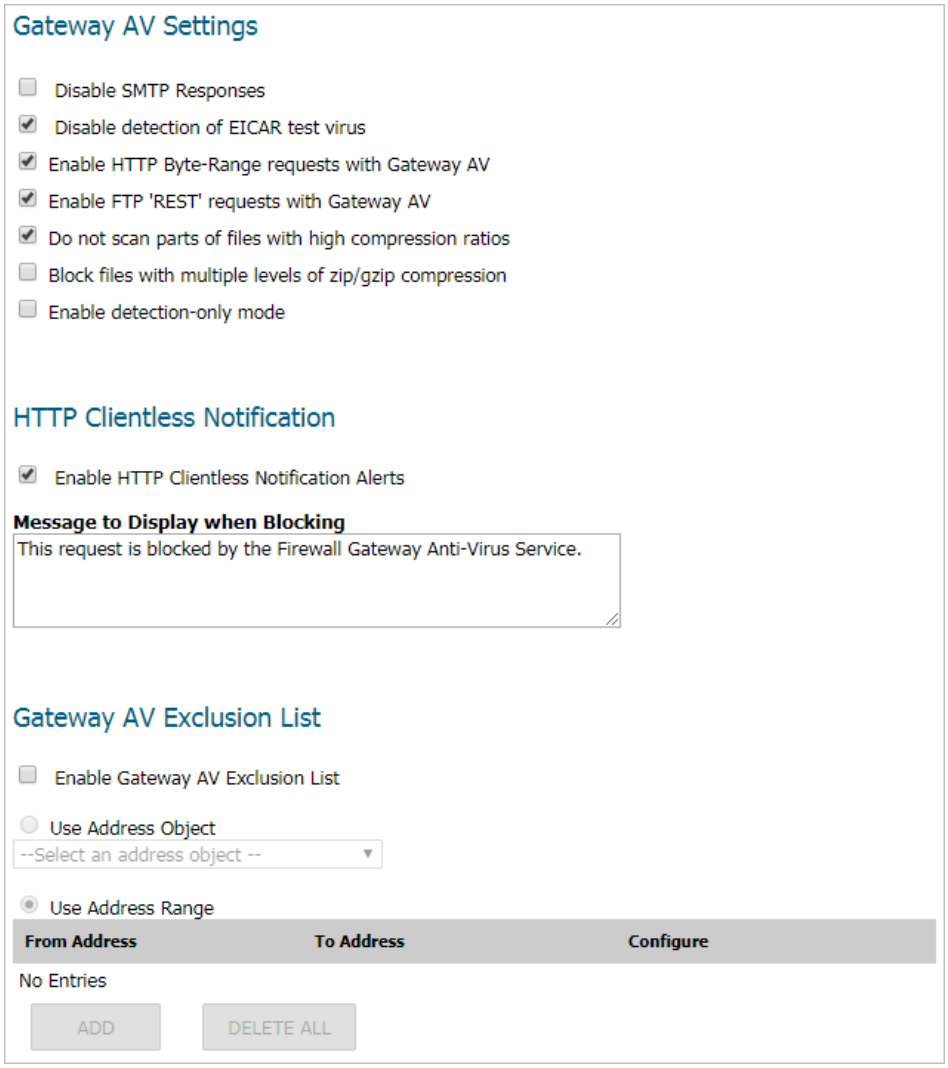

Topics:

- Configuring Gateway AV Settings
- Configuring HTTP Clientless Notification
- Configuring a SonicWall GAV Exclusion List

Configuring Gateway AV Settings

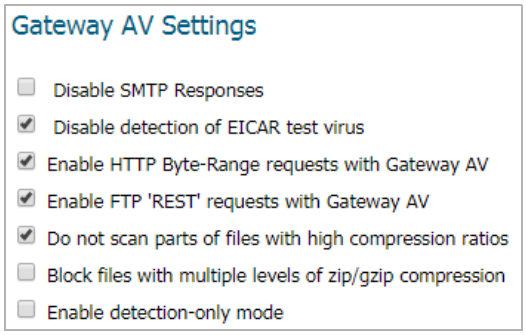

To configure Gateway AV options:

1. To suppress the sending of e-mail messages (SMTP) to clients from SonicWall GAV when a virus is detected in an e-mail or attachment, select the **Disable SMTP Responses** checkbox. This option is not selected by default.
2. The EICAR Standard Anti-Virus Test file is a special virus simulator file that checks and confirms the correct operation of the SonicWall Gateway AV service. To suppresses the detection of the EICAR, select the **Disable detection of EICAR test virus** checkbox. This option is selected by default.
3. To allow the sending of byte serving, the process of sending only a portion of an HTTP message or file, select the **Enable HTTP Byte-Range requests with Gateway AV** checkbox. This option is selected by default.

   The SonicWall Gateway Anti-Virus (GAV) security service, by default, suppresses the use of HTTP Byte-Range requests to prevent the sectional retrieval and reassembly of potentially malicious content. This is done by terminating the connection and thus preventing the user from receiving the malicious payload. By enabling this setting you override this default behavior.
4. To allow the use of the FTP REST request to retrieve and reassemble sectional messages and files, select the **Enable FTP 'REST' requests with Gateway AV** checkbox. This option is selected by default.

   The SonicWall GAV, by default, suppresses the use of the FTP 'REST' (restart) request to prevent the sectional retrieval and reassembly of potentially malicious content. This is done by terminating the connection and thus preventing the user from receiving the malicious payload. By enabling this setting you override this default behavior.
5. To suppresses the scanning of files, or parts of files, that have high compression rates, select the **Do not scan parts of files with high compression rates** checkbox. This option is selected by default.
6. To block files containing multiple levels of zip and/or gzip compression, select the **Block files with multiple levels of zip/gzip compression** checkbox. This option is selected by default.
7. To have the Gateway AV service in detection-only mode, which only detects and logs virus traffic without stopping such traffic, select the **Enable detection-only mode** checkbox. This option is not selected by default.

Configuring HTTP Clientless Notification

The HTTP Clientless Notification feature notifies users when GAV detects an incoming threat from an HTTP server.

If this feature is disabled, when GAV detects an incoming threat from an HTTP server, GAV blocks the threat and the user receives a blank HTTP page. Typically, users will attempt to reload the page because they are not aware of the threat. The HTTP Clientless Notification feature informs the user that GAV detected a threat from the HTTP server.

TIP: The HTTP Clientless Notification feature is also available for SonicWall Anti-Spyware.

To configure this feature.

1 Select the **Enable HTTP Clientless Notification Alerts** checkbox. This option is selected by default.



2 Optionally, enter a message in the **Message to Display when Blocking** field. The default message is `This request is blocked by the Firewall Gateway Anti-Virus Service.`

TIP: You can configure a timeout for the HTTP Clientless Notification on the **Security Services > Base Setup** page under the **Security Services Settings** heading.

Configuring a SonicWall GAV Exclusion List

Any IP addresses listed in the exclusion list bypass virus scanning on their traffic.The **Gateway AV Exclusion List** section provides the ability to either select an Address Object or define a range of IP addresses whose traffic will be excluded from SonicWall GAV scanning.

CAUTION: Use caution when specifying exclusions to SonicWall GAV protection.

To add an IP address range for exclusion:

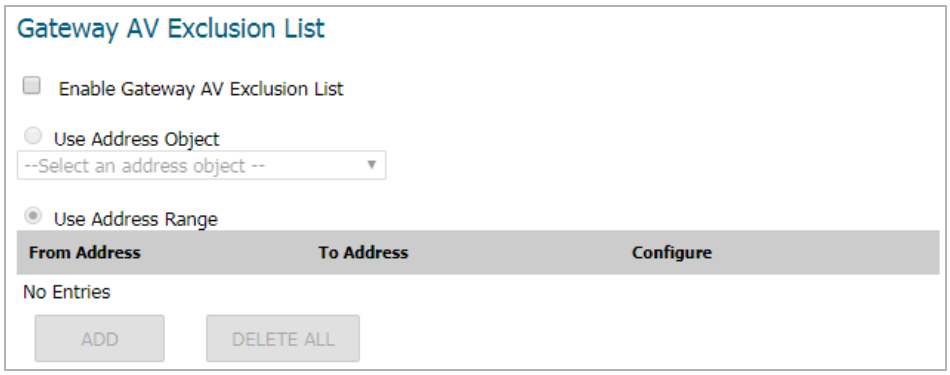

1 Navigate to **MANAGE | Security Configuration > Security Services > Gateway Anti-Virus**.

2 Scroll to the **Gateway Anti-Virus Global Settings** section.

3 Click the **CONFIGURE GATEWAY AV SETTINGS** button.

4 Select the **Enable Gateway AV Exclusion List** checkbox in the **Gateway AV Exclusion List** section to enable the exclusion list.

5 Select one of these:

- **Use Address Object** radio button
  a) Select an address object from the drop-down menu.
  b) Go to Step 6.
- **Use Address Range** radio button.
  a) Click the **Add** button. The **Add GAV Range Entry** dialog displays.

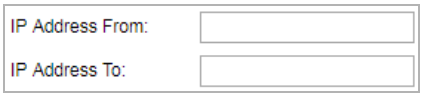

  b) Enter the IP address range in the **IP Address From** and **IP Address To** fields.
  c) Click **OK**. Your IP address range appears in the **Gateway AV Exclusion List** table.

NOTE: To change an entry, click the **Edit** icon in the **Configure** column or to delete an entry, click the **Delete** icon. To delete all entries in the exclusion list, click the **Delete All** button.

6 Click **OK**.

Configuring Cloud Gateway AV

To enable the Cloud Gateway Anti-Virus feature:

1 Navigate to the **Security Services > Gateway Anti-Virus > Cloud Anti-Virus Global Settings** section.

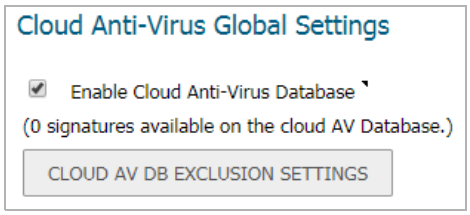

2 Select the **Enable Cloud Anti-Virus Database** checkbox. (This option is selected by default.)

Optionally, certain cloud-signatures can be excluded from being enforced to alleviate false positive problems or to enable downloading specific virus files as necessary.

To configure the exclusion list:

1 Click **CLOUD AV DB EXCLUSION SETTINGS**. The **Add Cloud AV Exclusion** dialog displays.

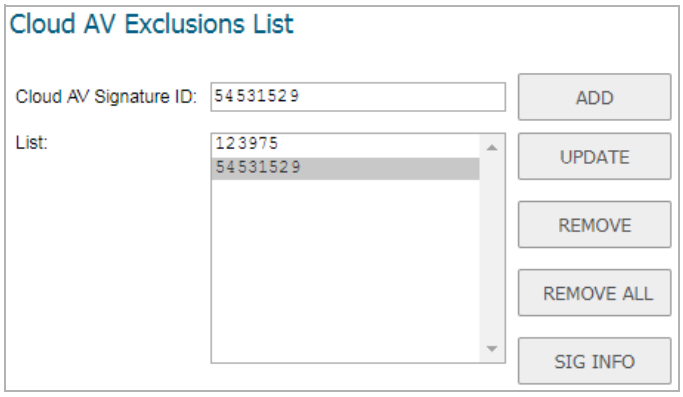

2 Enter the signature ID in the **Cloud AV Signature ID** field. The ID must be a numeric value.

3 Click **ADD**.

4 Repeat Step 2 and Step 3 for each signature ID to be added.

5 Optionally, to update a signature ID:

   a Select the signature ID in the **List** field.

   b Enter the updated signature in the **Cloud AV Signature ID** field.

   c Click **UPDATE**.

6 Optionally, to delete:

   - A signature ID, select the ID in the **List** field, and then click **REMOVE**.
   - All signatures, click **REMOVE ALL**.

7 Optionally, to view the latest information on a signature, select the signature ID in the list and click the **Sig Info** button. The information for the signature is displayed on the SonicALERT website.

8 Click **OK** when you have finished configuring the Cloud AV exclusion list.

Viewing SonicWall GAV Signatures

The **Gateway Anti-Virus Signatures** section allows you to view the contents of the SonicWall GAV signature database. All the entries displayed in the **Gateway Anti-Virus Signatures** table are from the SonicWall GAV

signature database downloaded to your SonicWall security appliance. The number of malware family signatures is displayed above the table.

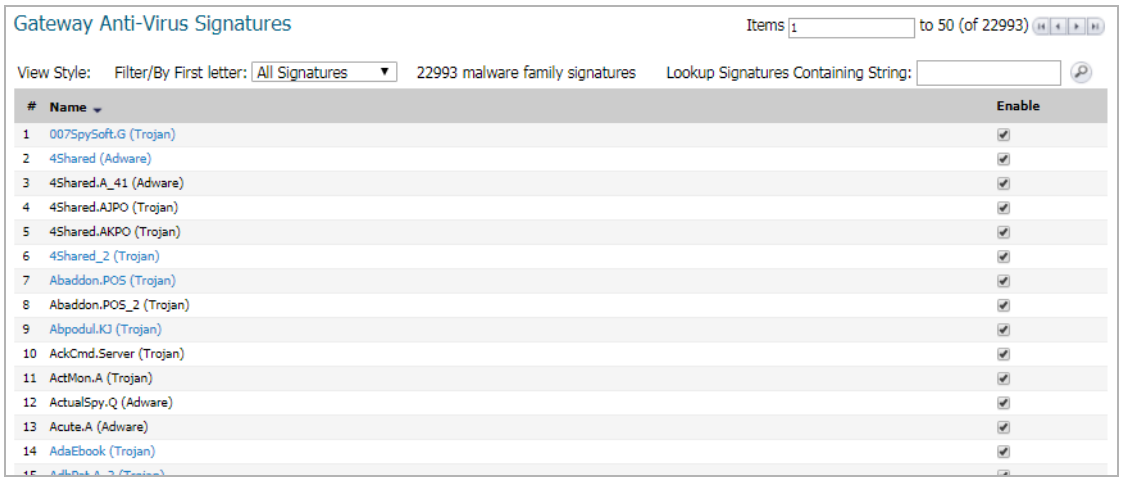

NOTE: Signature entries in the database change over time in response to new threats.

Topics:

- Displaying Signatures
- Navigating the Gateway Anti-Virus Signatures Table
- Searching the Gateway Anti-Virus Signature Database

Displaying Signatures

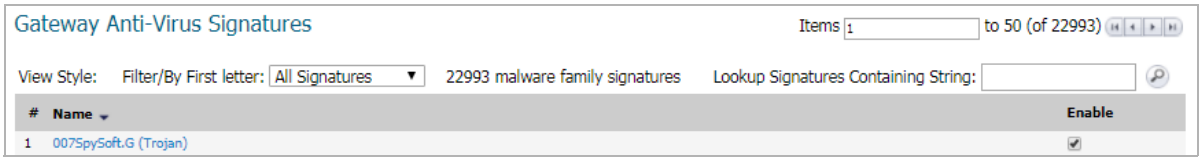

You can display the signatures in a variety of views:

TIP: When you filter the signature, the number of signatures found is displayed along with the total number of signatures in the database.

- **View Style** – Select one of these from the **First Letter** drop-down menu:
  - **All Signatures** - Displays all the signatures in the table, 50 to a page.
  - **0 – 9** - Displays signature names beginning with the number you select from the menu.
  - **A – Z** - Displays signature names beginning with the letter you select from menu.
- **Search String** - Displays signatures containing a specific string:
  a Enter the string in the **Lookup Signatures Containing String** field.
  b Click the **Magnifying Glass** icon.

Navigating the Gateway Anti-Virus Signatures Table

The SonicWall GAV signatures are displayed fifty to a page in the **Gateway Anti-Virus Signatures** table. The **Items** field displays the table number of the first signature. For information about navigating through the table, see *SonicWall SonicOS 6.5 About SonicOS*.

Searching the Gateway Anti-Virus Signature Database

You can search the signature database by entering a search string in the **Lookup Signatures Containing String** field, then clicking the **Search** icon.

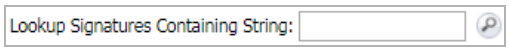

Only the signatures that match the specified string are displayed in the **Gateway Anti-Virus Signatures** table.

14

Activating Intrusion Prevention Service

- About Intrusion Prevention Service on page 178
- Configuring Intrusion Prevention Service on page 180

About Intrusion Prevention Service

Intrusion Prevention Service (IPS) delivers a configurable, high performance Deep Packet Inspection engine for extended protection of key network services such as Web, Email, file transfer, Windows services and DNS. SonicWall IPS is designed to protect against application vulnerabilities as well as worms, Trojans, and peer-to-peer, spyware and back-door exploits. The extensible signature language used in SonicWall's Deep Packet Inspection engine also provides proactive defense against newly discovered application and protocol vulnerabilities. SonicWall IPS off loads the costly and time-consuming burden of maintaining and updating signatures for new hacker attacks through SonicWall's industry-leading Distributed Enforcement Architecture (DEA). 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.

Topics:

- SonicWall Deep Packet Inspection on page 178
- How SonicWall's Deep Packet Inspection Works on page 179
- Glossary on page 179
- IPS Status on page 181
- IPS Global Settings on page 181

SonicWall Deep Packet Inspection

Deep Packet Inspection looks at the data portion of the packet. The Deep Packet Inspection technology includes intrusion detection and intrusion prevention. Intrusion detection finds anomalies in the traffic and alerts the administrator. Intrusion prevention finds the anomalies in the traffic and reacts to it, preventing the traffic from passing through.

Deep Packet Inspection is a technology that allows a firewall to classify passing traffic based on rules. These rules include information about layer 3 and layer 4 content of the packet as well as the information that describes the contents of the packet's payload, including the application data (for example, an FTP session, an HTTP Web browser session, or even a middleware database connection). This technology allows the administrator to detect and log intrusions that pass through the firewall, as well as prevent them (i.e. dropping the packet or resetting the TCP connection). SonicWall's Deep Packet Inspection technology also correctly handles TCP fragmented byte stream inspection as if no TCP fragmentation has occurred.

How SonicWall's Deep Packet Inspection Works

Deep Packet Inspection technology enables the firewall to investigate farther into the protocol to examine information at the application layer and defend against attacks targeting application vulnerabilities. This is the technology behind SonicWall Intrusion Prevention Service. SonicWall's Deep Packet Inspection technology enables dynamic signature updates pushed from the SonicWall Distributed Enforcement Architecture.

The following steps describe how the SonicWall Deep Packet Inspection Architecture works; see SonicWall deep packet inspection architecture:

1. Pattern Definition Language Interpreter uses signatures that can be written to detect and prevent against known and unknown protocols, applications and exploits.
2. TCP packets arriving out-of-order are reassembled by the Deep Packet Inspection framework.
3. Deep Packet Inspection engine preprocessing involves normalization of the packet's payload. For example, a HTTP request may be URL encoded and thus the request is URL decoded in order to perform correct pattern matching on the payload.
4. Deep Packet Inspection engine postprocessors perform actions which may either simply pass the packet without modification, or could drop a packet or could even reset a TCP connection.
5. SonicWall's Deep Packet Inspection framework supports complete signature matching across the TCP fragments without performing any reassembly (unless the packets are out of order). This results in more efficient use of processor and memory for greater performance.

SonicWall deep packet inspection architecture

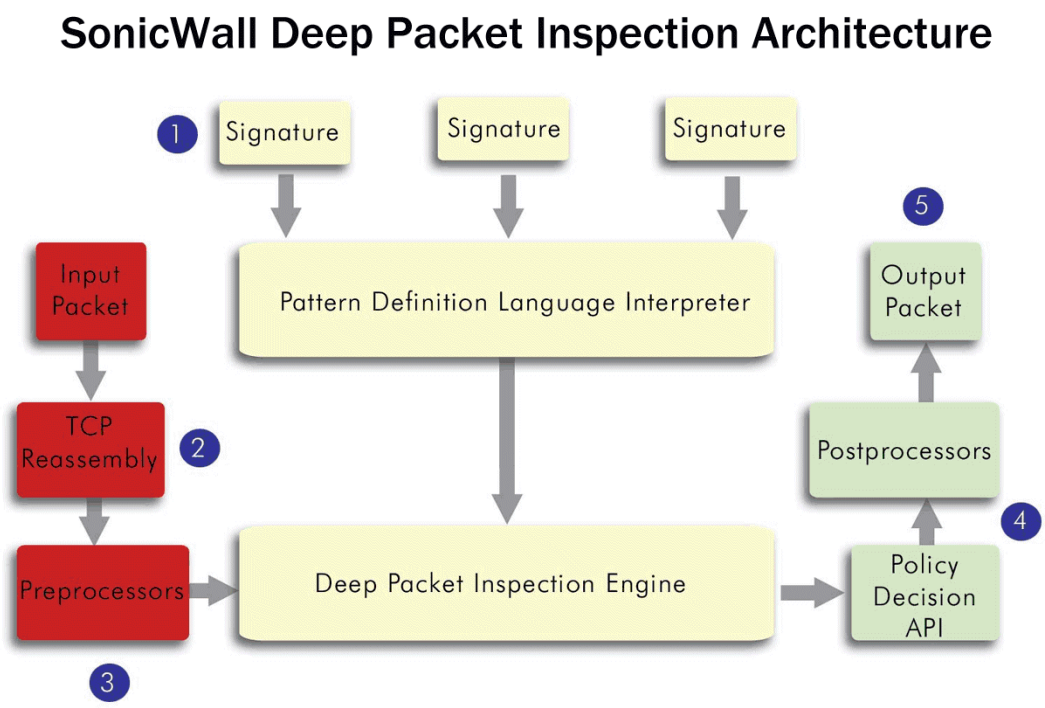

Glossary

- **Stateful Packet Inspection** - looking at the header of the packet to control access based on port, protocol, and IP address.

- **Deep Packet Inspection** - looking at the data portion of the packet. Enables the firewall to investigate farther into the protocol to examine information at the application layer and defend against attacks targeting application vulnerabilities.
- **Intrusion Detection** - a process of identifying and flagging malicious activity aimed at information technology.
- **False Positive** - a falsely identified attack traffic pattern.
- **Intrusion Prevention** - finding anomalies and malicious activity in traffic and reacting to it.
- **Signature** - code written to detect and prevent intrusions, worms, application exploits, and Peer-to-Peer and Instant Messaging traffic.

Configuring Intrusion Prevention Service

Intrusion Prevention Service (IPS) is configured on the **MANAGE | Security Configuration > Security Services > Intrusion Prevention** page, which is divided into panels:

- IPS Status
- IPS Global Settings
- IPS Policies

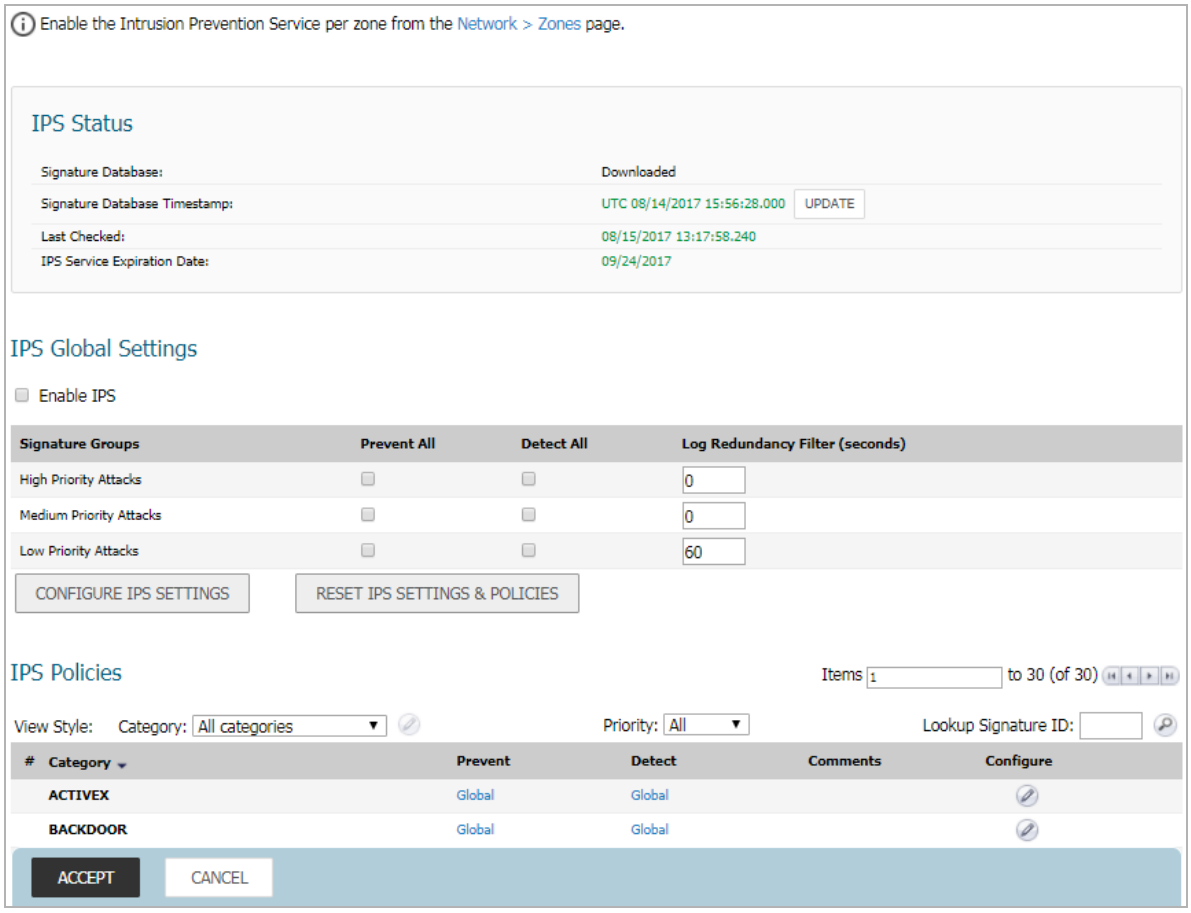

Topics:

- IPS Status on page 181

- IPS Global Settings on page 181
- Configuring IPS Protection on Zones on page 184
- IPS Policies on page 185

IPS Status

The **IPS Status** panel displays status information for the signature database and your SonicWall IPS license.



The **IPS Status** panel displays the following information:

- **Signature Database** indicates whether the signature database is being downloaded, has been downloaded, or needs to be downloaded. The signature database is updated automatically about once an hour. You can also manually update your IPS database at any time by clicking the **Update** button located in the **IPS Status** section.
- **Signature Database Timestamp** displays the last update to the IPS signature database, not the last update to your SonicWall security appliance.
- **Last Checked** indicates the last time the SonicWall security appliance checked the signature database for updates. The SonicWall security appliance automatically attempts to synchronize the database on startup, and once every hour.
- **IPS Service Expiration Date** indicates the date when the IPS service expires. If your IPS subscription expires, the SonicWall IPS inspection is stopped and the IPS configuration settings are removed from the SonicWall security appliance. After renewing your IPS license, these settings are automatically restored to the previously configured state.
- **Note:** Enable the Intrusion Prevention Service per zone from the Network > Zones page.

  If you click on Network > Zones in this note, it displays the **MANAGE | System Setup > Network > Zones** page where you can configure IPS on zones. See Configuring IPS Protection on Zones.

IPS Global Settings

The **IPS Global Settings** panel provides the key settings for enabling SonicWall IPS on your firewall.

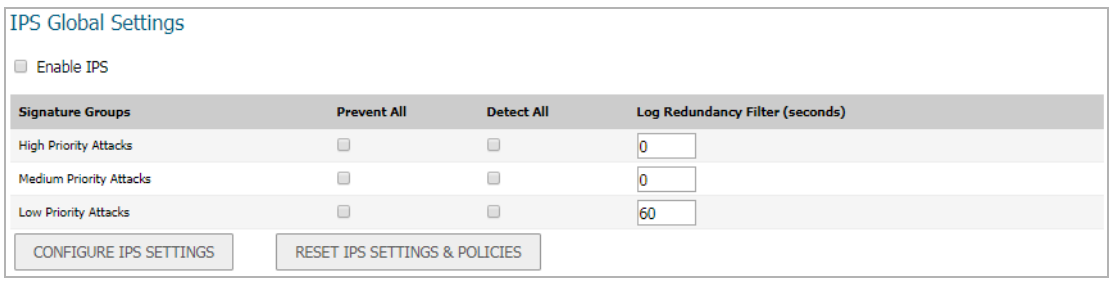

SonicWall IPS is activated by globally enabling IPS on your firewall and selecting the class of attacks. Optionally, you can configure an **IPS Exclusion List** as well.

Topics:

- Enabling IPS on page 182
- Configuring an IPS Exclusion List on page 183
- Resetting the IPS Settings and Policies on page 184
- Configuring IPS Protection on Zones on page 184

Enabling IPS

To enable IPS on your firewall:

1 Navigate to the **Security Configuration | Security Services > Intrusion Prevention** page.

2 Scroll down to the **IPS Global Settings** section.

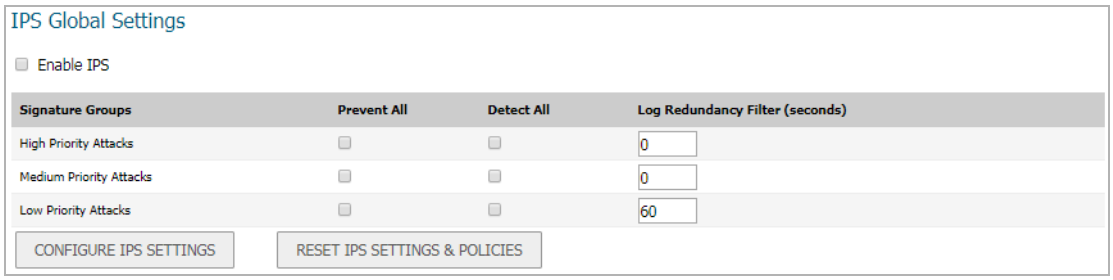

3 Select **Enable IPS**.

4 Select the action that you want (**Prevent All**, **Detect All**, or both) for each of the **Signature Groups**:

- **High Priority Attacks**
- **Medium Priority Attack**
- **Low Priority Attacks**

NOTE: To activate intrusion prevention on the firewall, you must specify a **Prevent All** action for at least one of the **Signature Groups**. If no **Prevent All** actions are checked, no intrusion prevention occurs on the firewall.

NOTE: Selecting both **Prevent All** and **Detect All** for all of the **Signature Groups** protects your network against the most dangerous and disruptive attacks.

Various attacks are often rapidly repeated, which can quickly fill up a log if each attack is logged. To reduce the duplicate number of logged attacks, enter the time, in seconds, in the **Log Redundancy Filter (seconds)** field, that the same attack is logged on the **INVESTIGATE | Logs > Event Logs** page as a single entry. The range for these intervals is 0 to 86400 seconds. The defaults for the various priorities of attacks are:

- **High Priority Attacks**: **0** seconds
- **Medium Priority Attacks**: **0** seconds
- **Low Priority Attacks**: **60** seconds

5 Click **ACCEPT**.

Configuring an IPS Exclusion List

(Optional) To configure an IPS Exclusion List:

1 Navigate to the **MANAGE | Security Configuration > Security Services > Intrusion Prevention** page.

2 Scroll down to the **IPS Global Settings** section.

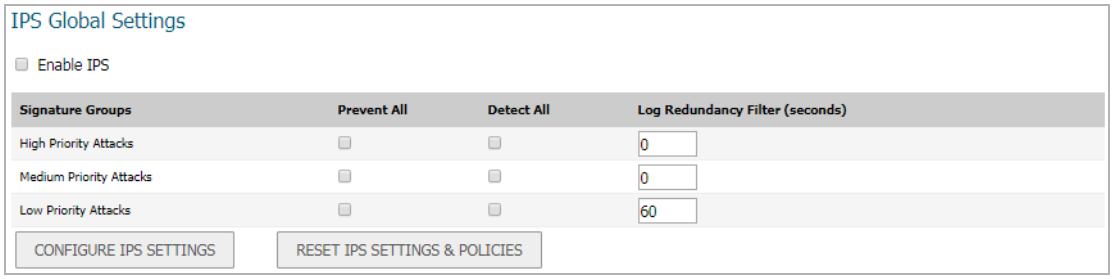

3 Select **Enable IPS**.

4 Click the **Configure IPS Settings** button.

The **IPS Exclusion List** dialog appears.

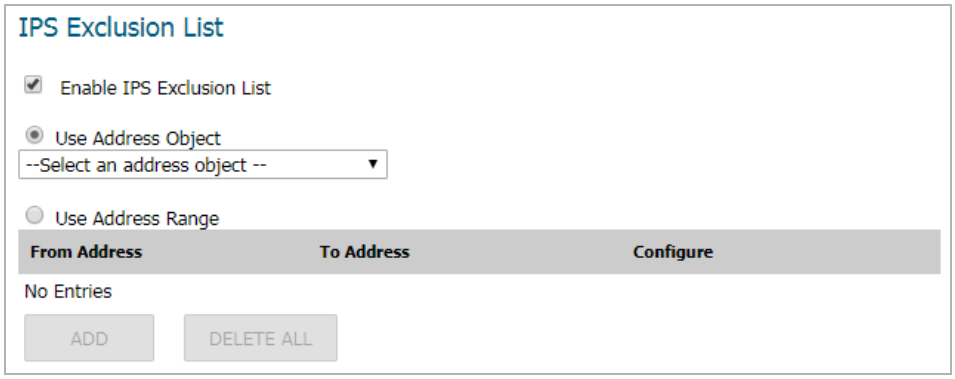

5 Select **Enable IPS Exclusion List**.

6 Select either the **Use Address Object** option or the **Use Address Range** option.

7 If you selected the **Use Address Object** option, select the address object you want to exclude from the menu.

8 If you selected the **Use Address Range** option, click the **Add** button.

The **Add IPS Range Entry** dialog appears.



9 Enter the IP address range to exclude in the **IP Address From** and the **IP Address To** boxes.

10 Click **OK**.

Resetting the IPS Settings and Policies

To reset the IPS Settings and Policies:

1 Navigate to the **MANAGE | Security Configuration > Security Services > Intrusion Prevention** page.

2 Scroll down to the **IPS Global Settings** section.

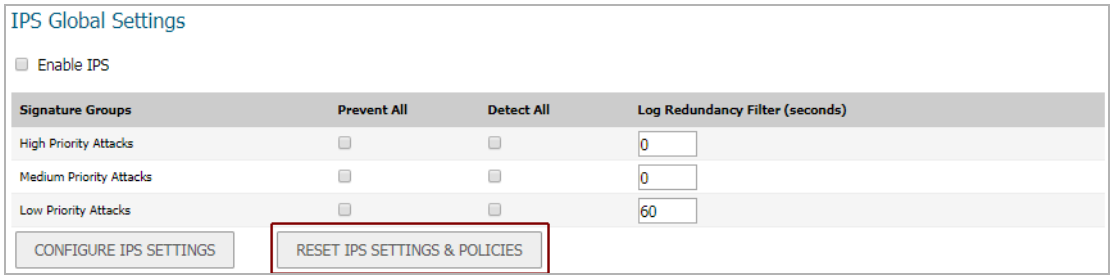

3 Click **RESET IPS SETTINGS & POLICIES**. The following message is displayed.

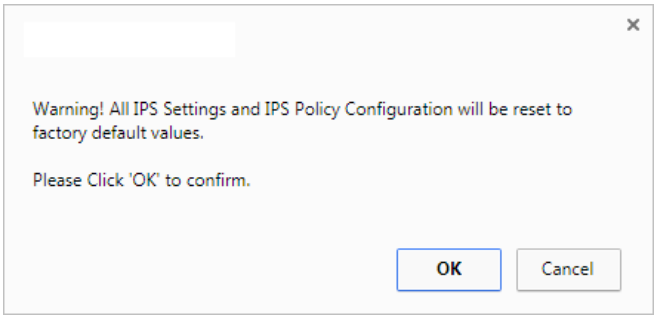

4 Click **OK**.

The following message appears at the bottom of the screen: `Status: The configuration has been updated.`

Configuring IPS Protection on Zones

You apply SonicWall IPS to zones on the **MANAGE | System Setup > Network > Zones** page to enforce SonicWall IPS not only between each network zone and the WAN, but also between internal zones. For example, enabling SonicWall IPS on the LAN zone enforces SonicWall IPS on all incoming and outgoing LAN traffic.

In the **IPS Status** section of the **Security Services > Intrusion Prevention** page, click the **Network > Zones** link to access the **MANAGE | System Setup > Network > Zones** page. You apply SonicWall IPS to a zone listed on the **Network > Zones** page.

To enable SonicWall on a zone:

1 Navigate to the **MANAGE | System Setup > Network > Zones** page or from the **IPS Status** section on the **MANAGE | Security Configuration > Security Services > Intrusion Prevention** page, click the **Network > Zones** link. The **Network > Zones** page is displayed.

2 In the **Configure** column in the **Zone Settings** table, click the **Edit** icon for the zone you want to apply SonicWall IPS. The **Edit Zone** dialog is displayed.

3 Click **Enable IPS**. A checkmark appears. To disable SonicWall IPS, clear the option.

4 Click **OK**.

You also enable SonicWall IPS protection for new zones you create on the **Network > Zones** page. Clicking the **Add** icon displays the **Add Zone** dialog, which includes the same settings as the **Edit Zone** dialog.

IPS Policies

The **IPS Policies** panel allows you to view SonicWall IPS signatures and configure the handling of signatures by category groups or on a signature by signature basis. Categories are signatures grouped together based on the type of attack.

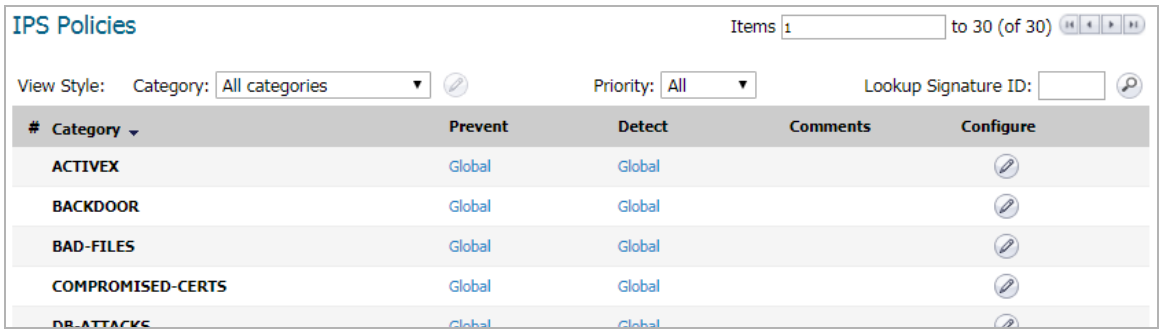

You can view the signatures in these ways:

- Viewing and Configuring Category Settings on page 185
- Viewing and Configuring Signature Settings on page 186
- Viewing and Configuring Signatures for Specific Categories on page 187
- Priority Menu on page 187
- Lookup Signature ID on page 187

Viewing and Configuring Category Settings

In the **View Style** row, the **Category** menu lets you choose the categories or signatures you want to display in the **Category** column. You can choose **All categories**, **All signatures**, or an individual category, such as **ACTIVEX** or **DNS**. If you choose an individual category, the signatures for that category are displayed.

The **Category** column allows you to sort categories and signatures in ascending or descending order by clicking the up or down arrow next to the column heading.

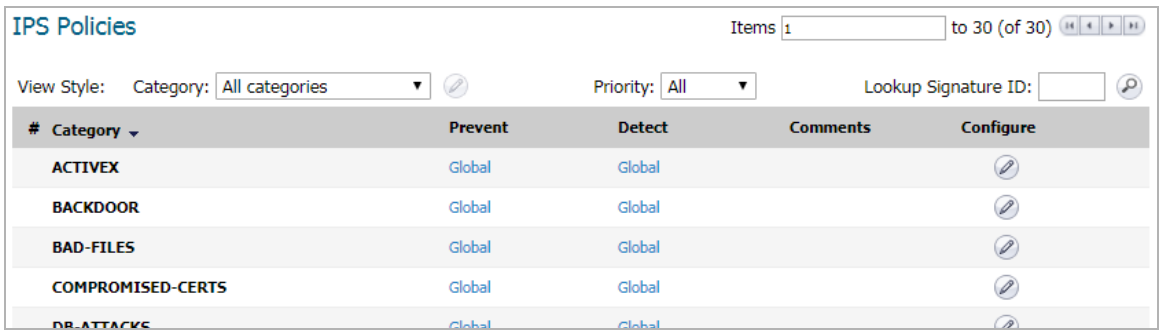

To view or change the IPS category settings for a particular category:

1 Select **All categories** from the **Category** menu.

2 Click the **Edit** icon in the **Configure** column for that category. The **Edit IPS Category** dialog appears.

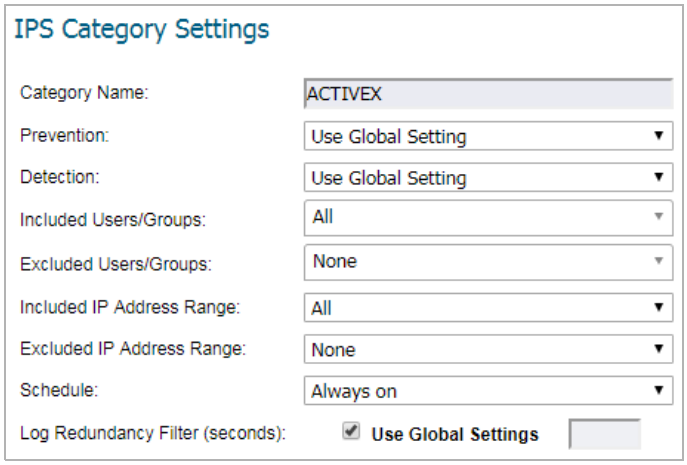

3 From the **Prevention** and **Detection** menus, select **Use Global Setting**, **Enable**, or **Disable**. If you select **Use Global Setting**, the values configured in the **IPS Global Settings** section are used, but you can override the **IPS Global Settings** by selecting **Enable** or **Disable** from these menus.

4 From the remaining menus, select the values that you want.

5 For the **Log Redundancy Filter (seconds)** option, if you want to use the values that you configured in the **IPS Global Settings** section, select **Use Global Settings.**

6 Click **OK**.

Viewing and Configuring Signature Settings

To view or change the IPS signature settings for a particular signature:

1 Select **All signatures** from the **Category** menu.

2 Click the **Edit** icon in the **Configure** column for that signature. The **Edit IPS Signature** dialog appears.

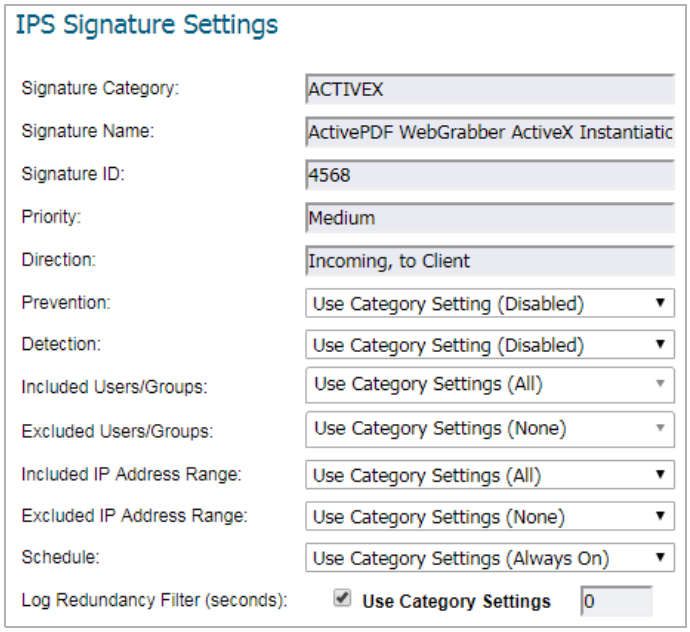

The first five boxes are grayed and contain non-configurable data for that signature.

3 From the **Prevention** and **Detection** menus, select **Enable** or **Disable**. The **Use Category Setting option** is disabled.

4 From the remaining menus, select the values that you want.

5 For the **Log Redundancy Filter (seconds)** option, if you want to use the values that you configured in the **IPS Global Settings** section, select **Use Category Settings.**

6 Click **OK**.

Viewing and Configuring Signatures for Specific Categories

To view and configure signatures for specific categories:

1 Select one of the individual categories from the **Category** menu. The signatures for that category are displayed.

2 Click the **Edit** icon in the **Configure** column for that signature. The **Edit IPS Signature** dialog appears.

The first five boxes are grayed and contain non-configurable data for that signature.

3 From the **Prevention** and **Detection** menus, select **Enable** or **Disable**. The **Use Category Setting option** is disabled.

4 From the remaining menus, select the values that you want.

5 For the **Log Redundancy Filter (seconds)** option, if you want to use the values that you configured in the **IPS Global Settings** section, select **Use Category Settings.**

6 Click **OK**.

Priority Menu

The **Priority** menu lets you specify the priority of the signatures you want to display.

To specify the priority of the signatures you want to display:

- Select one of the following priorities from the **Priority** menu:
  - **All**
  - **High**
  - **Medium**
  - **Low**

Lookup Signature ID

You can use the **Lookup Signature ID** box to view or change the IPS signature settings for a particular signature.

To view or change the IPS signature settings for a particular signature:

1 Enter the signature ID in the **Lookup Signature ID** box.

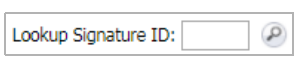

2 Click the **Lookup** icon next to the field. The **Edit IPS Signature** dialog appears.

The first five fields are grayed and contain non-configurable data for that signature.

3 From the **Prevention** and **Detection** menus, select **Enable** or **Disable**. The **Use Category Setting option** is disabled.

4 From the remaining menus, select the values that you want.

5 For the **Log Redundancy Filter (seconds)** option, if you want to use the values that you configured in the **IPS Global Settings** section, select **Use Category Settings.**

6 Click **OK**.

15

Configuring Capture ATP

- Security Services > Capture ATP
- About Capture ATP
- Enabling Capture ATP
- About the Security Services > Capture ATP Page
- Configuring Capture ATP
- Disabling GAV or Cloud Anti-Virus

Security Services > Capture ATP

IMPORTANT: Capture Advanced Threat Protection (ATP) is an add-on security service to the firewall, similar to Gateway Anti-Virus (GAV), that helps a firewall identify whether a file is malicious.

Capture ATP is supported on all SuperMassive Series, NS*a* Series, NSA Series, TZ600/TZ600P, and TZ500/TZ500W firewalls running SonicOS 6.5 or higher. Capture functionality, however, is not supported in Active/Active DPI mode.

Before you can enable Capture ATP you must first get a license, and you must enable the Gateway Anti-Virus (GAV) and Cloud Anti-Virus Database services. After Capture ATP is licensed, you can view Capture ATP status in your MySonicWall account as well as configure and receive alerts and notifications.

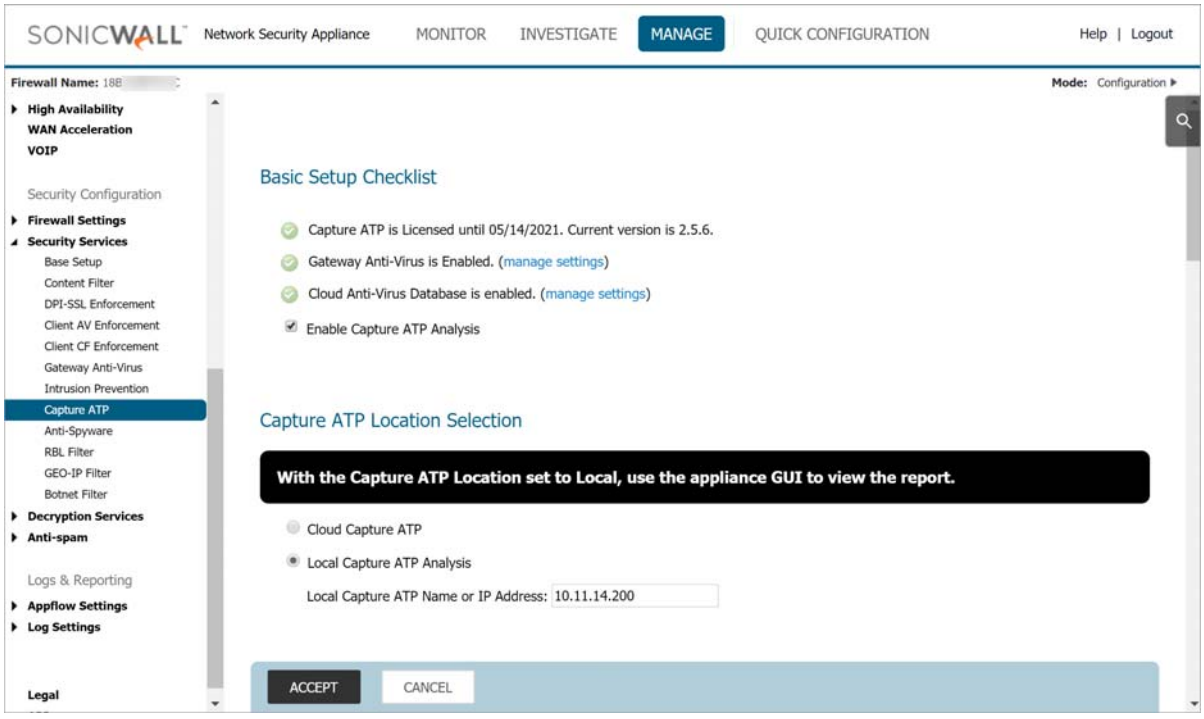

About Capture ATP

Starting in SonicOS 6.5.4.6, Capture Advanced Threat Protection (Capture ATP) supports Real-Time Deep Malware Inspection (RTDMI) in two forms:

- **SonicWall Capture ATP service**

  SonicWall's cloud-based, multi-engine, real-time file analysis service.

- **SonicWall Capture Security Appliance (CS*a*)**

  Brings the power of RTDMI into an appliance form factor to serve customers who, due to geographical, regulatory or organizational requirements, cannot send files to the cloud for ATP analysis. The CS*a* connects to your firewall or to your local network. See the *Capture Security Appliance Getting Started Guide* on the SonicWall technical documentation portal for information about connecting, licensing, and registering your CS*a*.

Capture ATP helps a firewall identify whether a file is malicious by transmitting the file to the cloud or Capture Security Appliance where SonicWall Capture ATP analyzes the file to determine if it contains a virus or other malicious elements. Capture ATP then sends the results to the firewall. The analysis and reporting are done in real time while the file is being processed by the firewall.

All files are sent to the Capture ATP cloud over an encrypted connection. Files are analyzed and deleted within minutes of a verdict being determined, unless a file is found to be malicious. Malicious files are submitted via an encrypted HTTPS connection to the SonicWall threat research team for further analysis and to harvest threat information. Files are not transferred to any other location for analysis. Malicious files are deleted after harvesting threat information within 30 days of receipt.

Capture ATP provides a file analysis report (threat report) with detailed threat behavior information.

The firewall and Capture Security Appliance are typically located on your premises, while the Capture ATP cloud service server and database are located at a SonicWall facility. The firewall creates a secure connection with the Capture ATP cloud service before transmitting data.

Capture ATP works in conjunction with the Gateway Anti-Virus (GAV) and Cloud Anti-Virus services. Capture ATP also logs/displays email header information (`to`, `cc`, `bcc`) parsed by GAV.

Topics:

- Files are Preprocessed on page 190
- Files Blocked Until Completely Analyzed on page 191
- Files are Sent over an Encrypted Connection
- Capture ATP Friendly Filename Display on page 191
- Activating the Capture ATP License

Files are Preprocessed

All files submitted to Capture ATP for analysis are first preprocessed by the GAV service to determine if a file is malicious or benign. You can also use GAV settings to select or define address objects to exclude from GAV and Capture ATP scanning.

Preprocessed files determined to be malicious or benign are not analyzed by Capture ATP. If a file is not determined to be malicious or benign during preprocessing, the file is submitted to Capture ATP for analysis.

Files Blocked Until Completely Analyzed

For HTTP/HTTPS downloads, Capture ATP has an option, **Block file download until a verdict is returned**, that ensures no packets get through until the file is completely analyzed and determined to be either malicious or benign. The file is held until the last packet is analyzed. If the file has malware, the last packet is dropped, and the file is blocked. The threat report provides information necessary to respond to a threat or infection.

Files are Sent over an Encrypted Connection

All files are sent to the Capture ATP cloud over an encrypted connection. SonicWall does not keep the files. All file types, whether they are malicious or benign are removed from the Capture ATP server after a certain time period.

The SonicWall privacy policy can be accessed at https://www.mysonicwall.com/privacypolicy.aspx.

Capture ATP Friendly Filename Display

SonicWall Capture Advanced Threat Protection logs the friendly filename of scanned files for the following non-HTTP protocols:

- SMTP
- POP3
- FTP
- IMAP
- NetBIOS

With this feature, you can easily identify the files being scanned by Capture ATP and their status displayed for filenames of these protocol types in the **MONITOR | Event Summaries > Capture ATP |Status** table and in log messages. Friendly filenames can be up to a maximum of 256 characters.

This feature cannot parse:

- Filename information for TCP protocol streams.
- A filename if it is not part of a single network packet.

No SonicOS configuration is required.

Activating the Capture ATP License

IMPORTANT: Capture ATP requires the Gateway Anti-Virus service, which must also be licensed.

To activate the Capture ATP service license, go to the **Updates > Licenses** page where you can view all service licenses and initiate licensing for Capture ATP. For more information about service licensing, see the *SonicOS 6.5 Updates* administration guide.

The Capture Security Appliance must be registered and licensed in MySonicWall and then synchronized with MySonicWall from the CS*a* local management interface. Refer to the *Capture Security Appliance Getting Started Guide* for detailed information.

After the Capture ATP service license or the CS*a* license is activated, **Capture ATP** appears in the SonicOS left-hand navigation (left nav) panel below DPI-SSL. If Capture ATP is not licensed, it does not appear in the left nav at all.

NOTE: Click on the **Synchronize** button on the **MANAGE | Updates > Licenses** page if **Capture ATP** does not appear in the left nav shortly after the Capture ATP service license or CS*a* license is activated.

Enabling Capture ATP

IMPORTANT: You must enable Gateway Anti-Virus and Cloud Anti-Virus before you can enable Capture ATP.

When Capture ATP is licensed but not enabled, the banner displays this message:

```
Capture ATP is not currently running. Please see the Basic Setup
Checklist below for troubleshooting.
```

In disabled mode, the **Basic Setup Checklist** section is visible, but the other sections are dimmed.

To enable Capture ATP:

1. Navigate to **MANAGE | Security Configuration > Security Services > Gateway Anti-Virus**.
2. Enable both Gateway Anti-Virus (GAV) and Cloud Anti-Virus as described in Managing SonicWall Gateway Anti-Virus Service on page 160.
3. Optionally, you can configure GAV and Cloud Anti-Virus settings, which also apply to Capture ATP.
4. Navigate to **MANAGE | Security Configuration > Security Services > Capture ATP**. If Capture ATP is not enabled, a warning message displays:

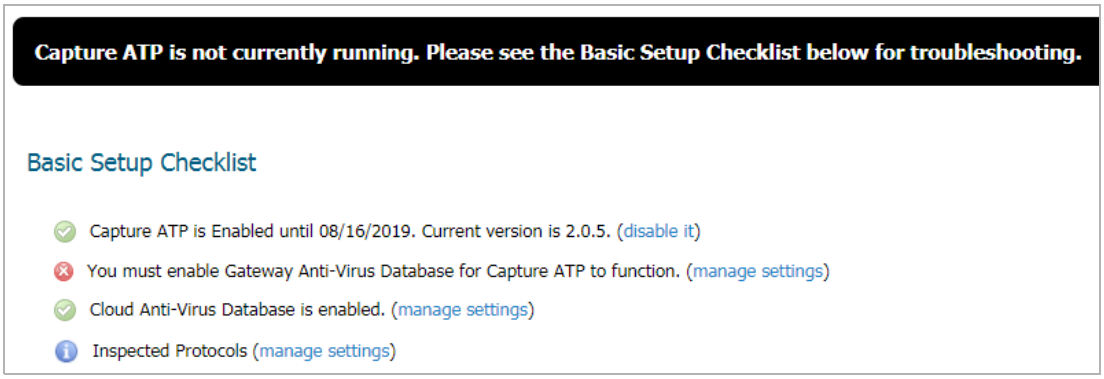

5. In the **Basic Setup Checklist** section, click **(enable it)** in **Capture ATP subscription is valid until *date* but the service is not currently enabled. (enable it)**. The warning message disappears, and the status indicator becomes a green checkmark.

About the Security Services > Capture ATP Page

Topics:

- Basic Setup Checklist
- Capture ATP Location Selection and Diagnostics
- Inspected Protocols
- Bandwidth Management
- Exclusions
- Custom Blocking Behavior

Basic Setup Checklist

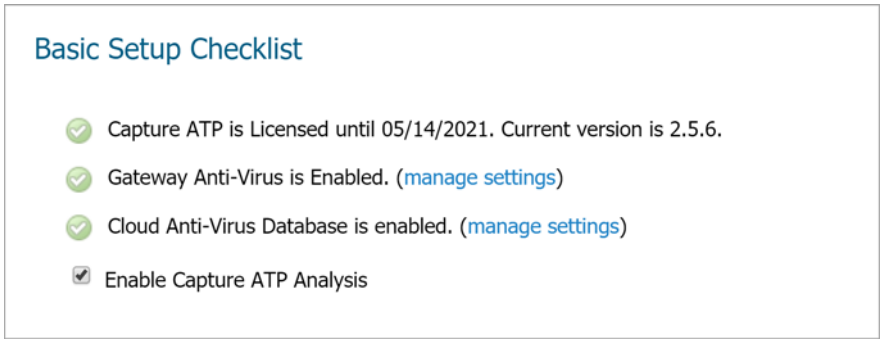

The **Basic Setup Checklist**:

- Displays the status of Capture ATP and its components, GAV and Cloud Anti-Virus.
- Displays any error states that may be present.
- Allows enabling or disabling of Capture ATP Analysis.
- Provides links to the **MANAGE | Security Configuration > Security Services > Gateway Anti-Virus** page for the GAV and Cloud Anti-Virus settings.

NOTE: For messages that display in this section, see Capture ATP status, Gateway Anti-Virus status, and Cloud Anti-Virus database status tables. **Enabled** corresponds to a green checkmark, and **Disabled** corresponds to a red X.

Capture ATP status

Icon	Message	Link	Action
Enabled	Capture ATP service is enabled until *renewal_date*.	`disable it`	Click the link to turn off Capture ATP and put the service in disabled mode. You do not need to click **ACCEPT** to apply this change.
Disabled	Capture ATP subscription is valid until *renewal_date* but the service is not currently enabled.	`enable it`	Click the link to turn on Capture ATP and put the service in enabled mode. You do not need to click **ACCEPT** to apply this change.
Disabled	Capture ATP subscription expired on *renewal_date*.	`renew it`	Click the link to go to MySonicWall to renew the service.

Gateway Anti-Virus status

Icon	Message	Link	Action
Enabled	Gateway Anti-Virus is Enabled.	`manage settings`	Click the link to display the **Security Services > Gateway Anti-Virus** page.
Disabled	You must enable Gateway Anti-Virus for Capture ATP to function.	`manage settings`	Click the link to display the **Security Services > Gateway Anti-Virus** page.

Cloud Anti-Virus database status

Icon	Message	Link	Action
Enabled	Cloud Anti-Virus Database is enabled.	`manage settings`	Click the link to display the **Security Services > Gateway Anti-Virus** page.
Disabled	You must enable the Cloud Anti-Virus Database for Capture ATP to function.	`manage settings`	Click the link to display the **Security Services > Gateway Anti-Virus** page.

Capture ATP Location Selection and Diagnostics

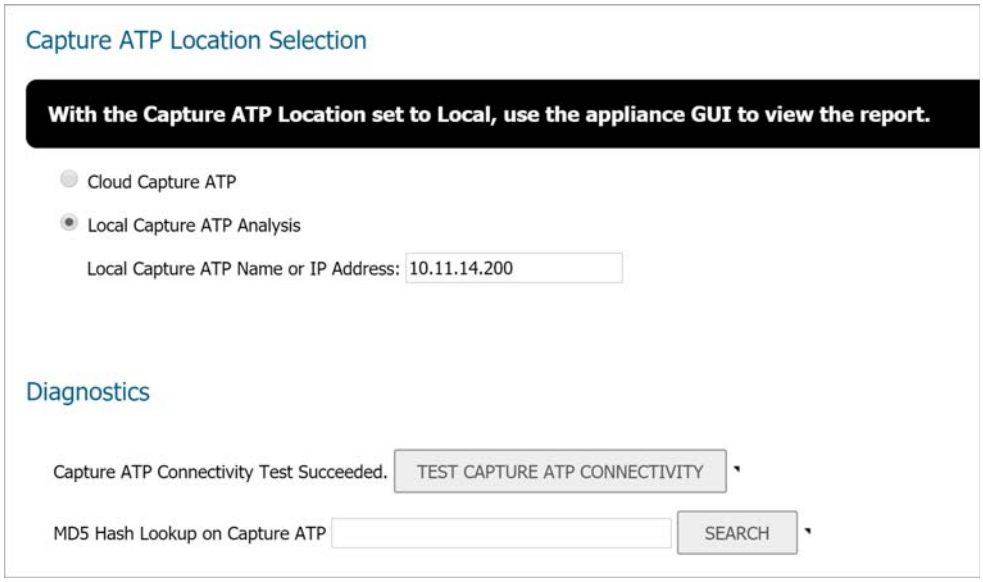

The **Capture ATP Location Selection** section provides options to select either:

- **Cloud Capture ATP** – the Capture ATP service
- **Local Capture ATP Analysis** – the Capture Security Appliance

Diagnostics is available only when deployed in Local mode. The **Diagnostics** section provides connectivity testing and lookup options:

- **TEST CAPTURE ATP CONNECTIVITY** button – Click the button to display the status of connectivity to the Capture ATP cloud server or the Capture Security Appliance. The result is displayed next to the button, such as “Capture ATP Connectivity Test Succeeded” or “Capture ATP Connectivity Test Failed”.

  The test fails if no response is received after one minute. If the test fails, the button changes to **CLEAR RESULTS**. Click it to clear the status and change the button back to **TEST CAPTURE ATP CONNECTIVITY**.

- **MD5 Hash Lookup on Capture ATP** – Enter the MD5 Hash value into the field and click the **SEARCH** button to perform a Capture ATP sandbox or Capture Security Appliance MD5 query.

  If a file with this MD5 hash value was previously submitted to Capture ATP, the result is displayed as “GOOD” if this MD5 is known to Capture ATP to be of a non-malicious file, or “BAD” if known to be a malicious file. “PENDING” is also a possible result if the file is still being analyzed. If no such MD5 hash value was previously submitted, the test times out after ten seconds with the message, “MD5 query result unavailable.”

NOTE: UDP port 2259 is used for all Capture ATP communications, file submissions, and reporting. All Capture ATP traffic sent or received is encrypted.

Inspected Protocols



The **Inspected Protocols** table provides a `manage settings` link that takes you to the **Security Services > Gateway Anti-Virus** page.

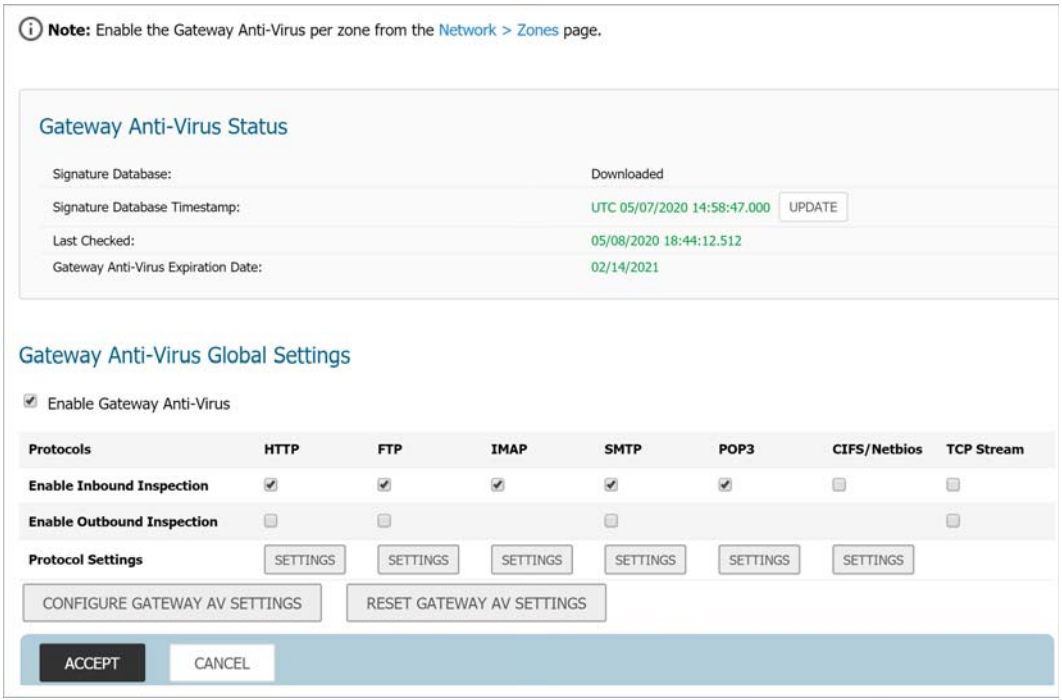

There, you can enable or disable inspection of specific network traffic protocols, including HTTP, FTP, IMAP, SMTP, POP, CIFS, and TCP Stream. Each protocol can be managed separately for inbound and outbound traffic.

The table in the **Inspected Protocols** section displays a matrix of the current protocol inspection settings for each protocol, and whether the inbound and outbound directions have been enabled. **Enabled** corresponds to a green checkmark, and **Disabled** corresponds to a red X. See the Protocols inspection settings table.

Protocols inspection settings

Icon	Message
Enabled	Protocol is inspected.
Disabled	Protocol is not inspected.
n/a	Inspection is not applicable to this protocol in this direction.

Bandwidth Management

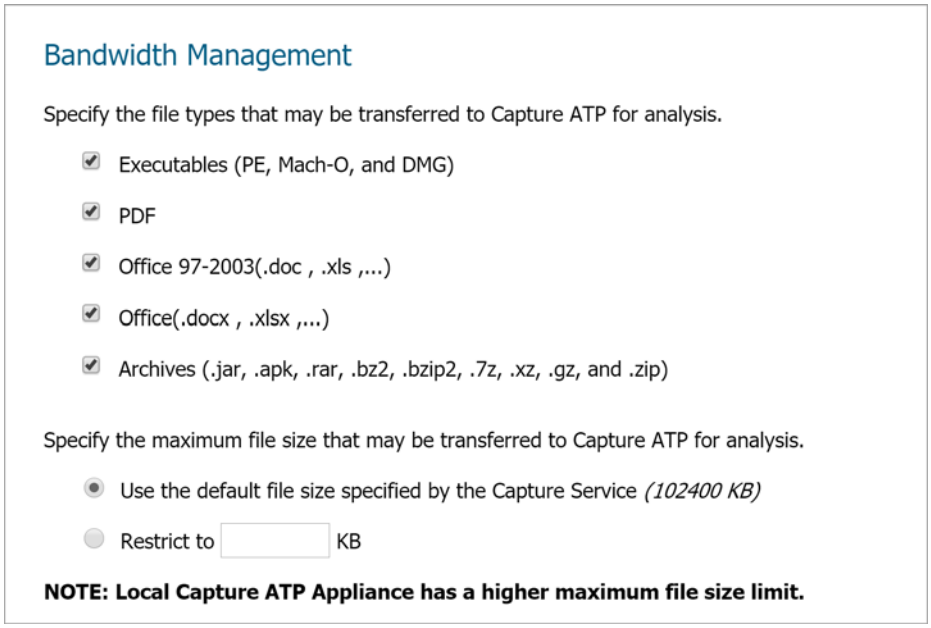

The **Bandwidth Management** section enables you to select the types of files to be submitted to Capture ATP and to specify the maximum size of submitted files. You can also specify an address object to be excluded from inspection.

By default, only the **Executables (PE, Mach-O, and DMG)** file type is enabled.

The default option for the maximum file size is **Use the default file size specified by the Capture Service (10240 KB)**. This specifies a file size limit of 10 megabytes (10 MB).

If you select **Restrict to KB**, you can enter your own custom value. This value must be a non-zero value and must not be greater than the default limit. The Capture Security Appliance has a 100MB maximum file size limit, higher than the limit for the cloud Capture ATP service.

For **Choose an Address Object to exclude from Capture ATP**, optionally select an address object from the drop-down list, or select the option to create a new address object. Members of the selected address object will be excluded from inspection by the Capture ATP service.

Exclusions

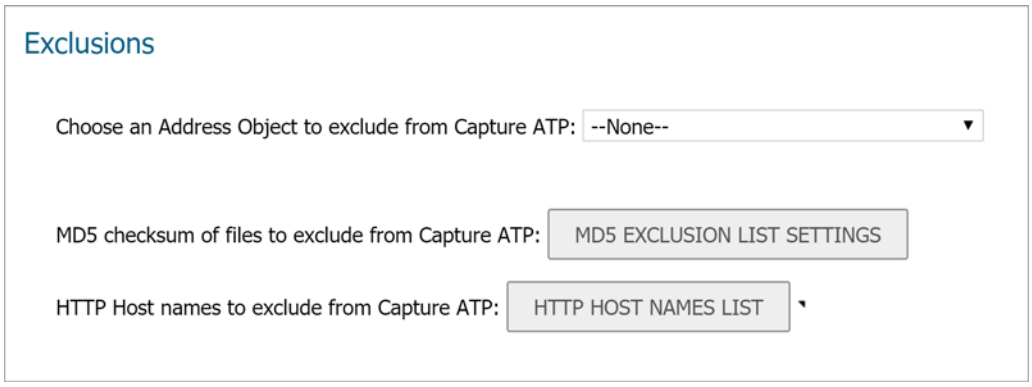

The **Exclusions** section allows you to exclude an Address Object or Group, MD5 hash/checksum values, or HTTP host names (FQDN or IP addresses) from Capture ATP.

Custom Blocking Behavior

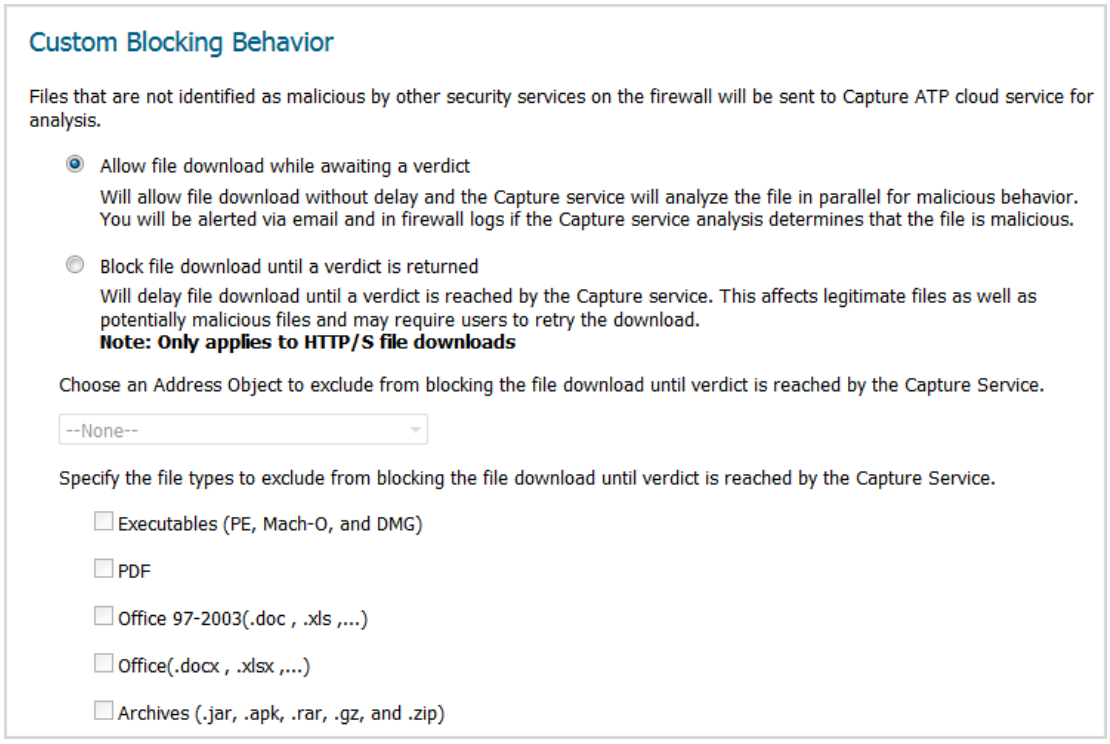

The **Custom Blocking Behavior** section allows you to select the **Block file download until a verdict is returned** feature.

The default option is **Allow file download while awaiting a verdict**. This setting allows a file to be downloaded without delay while the Capture service analyzes the file for malicious elements. You can set email alerts or check the firewall logs to find out if the Capture service analysis determines that the file is malicious.

The **Block file download until a verdict is returned** feature should only be enabled if the strictest controls are desired. If you select this feature, a warning dialog appears.

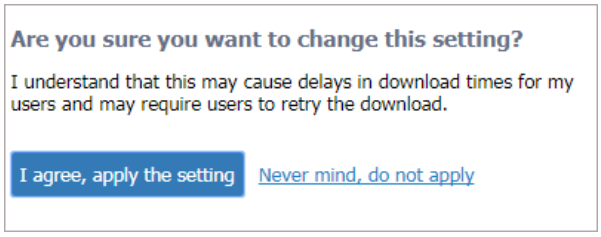

When the **Block file download until a verdict is returned** feature is enabled, the other options become available. You can:

- Select an address object from **Choose an Address Object to exclude from blocking the file download until verdict is reached by the Capture Service**. The default is **None**.
- Select one or more file types to block from **Specify the file types to exclude from blocking the file download until verdict is reached by the Capture Service**:
  - **Executables (PE, Mach-O, and DMG)**
  - **PDF**
  - **Office 97-2003(.doc , .xls ,...)**

- **Office(.docx , .xlsx ,...)**
- **Archives (.jar, .apk, .rar, .gz, and .zip)**

Configuring Capture ATP

To configure Capture ATP:

1 Navigate to **Capture ATP > Settings**.

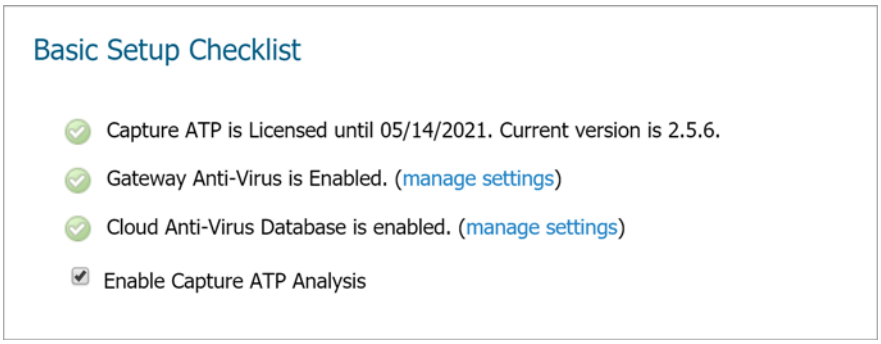

2 In the **Basic Setup Checklist** section, ensure that Capture ATP is licensed and that GAV, Cloud Anti-Virus database, and Capture ATP Analysis are enabled.

3 In the **Capture ATP Location Selection** section, select one of:

- **Cloud Capture ATP** – for the Capture ATP service
- **Local Capture ATP Analysis** – to use a Capture Security Appliance

  For **Local Capture ATP Name or IP Address**, enter the FQDN or IP address of the Capture Security Appliance your firewall is using.

4 In the **Inspected Protocols** section, ensure that the relevant protocols are enabled in the desired directions. To make changes, click the `manage settings` link to go to the **Security Services > Gateway Anti-Virus** page.

5 In the **Bandwidth Management** section, select the file types to be analyzed by Capture ATP. By default, only **Executables (PE, Mach-O, and DMG)** is selected.

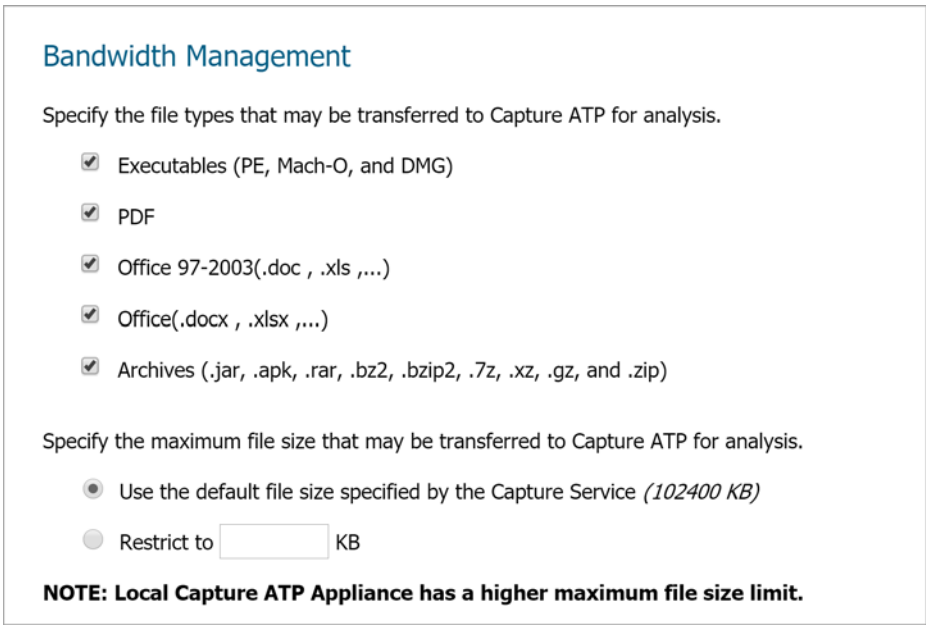

6 By default **Use the default file size specified by the Capture Service *(10240 KB)*** is selected. To specify a custom size, enter a value between 1 and 10240 in the **Restrict to KB** field.

If you are using a Capture Security Appliance, the custom size can be between 1 and 100MB (102,400KB).

7 Under **Exclusions**, to exclude an Address Object or Address Group from Capture ATP, select the Address Object/Group from the **Choose an Address Object to Exclude from Capture ATP** drop-down menu.

8 To exclude one or more files based on the MD5 checksum, click the **MD5 EXCLUSION LIST SETTINGS** button to display the **MD5 Exclusions List** dialog.

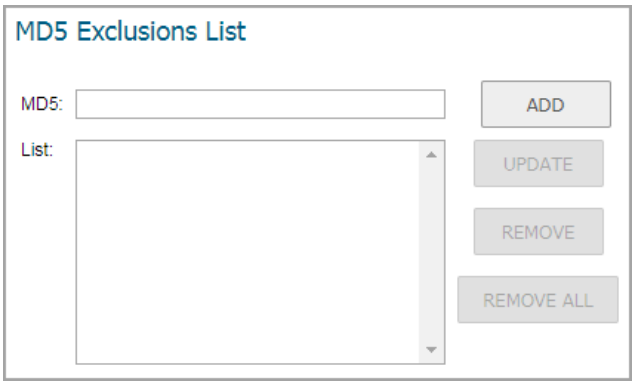

a Add the 32-digit hexadecimal hash value to the **MD5** field.

b Click **ADD** to add it to the **List** box.

c Repeat Step a and Step b for each file to exclude.

d Click **OK**.

9 To exclude one or more host names, click the **HTTP HOST NAMES LIST** button to display the **FQDN Exclusions List** dialog.

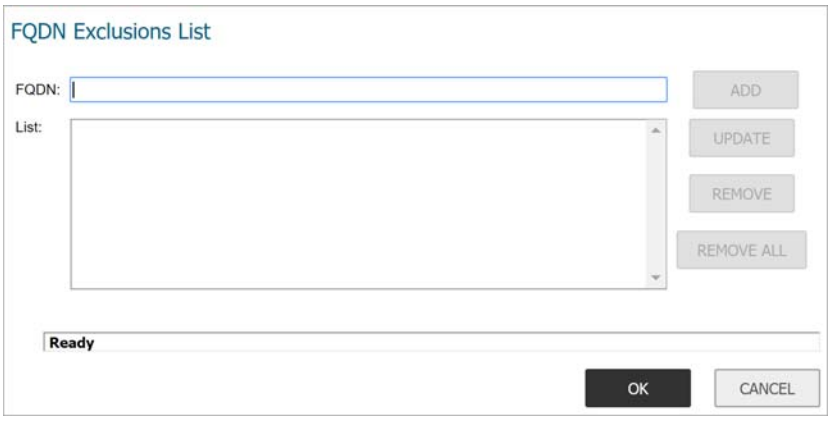

a In the **FQDN** field, enter the FQDN or IP address of the host name to be excluded. You can use an asterisk '*' as a wildcard for the prefix of a fully qualified domain name, for example: `*.google.com`.

b Click **ADD** to add it to the **List** box.

c Repeat Step a and Step b for each host name to exclude.

d Click **OK**.

10 If you are analyzing HTTP/HTTPS files, in the **Custom Blocking Behavior** section, you can specify whether all files are to be blocked until analysis is completed.

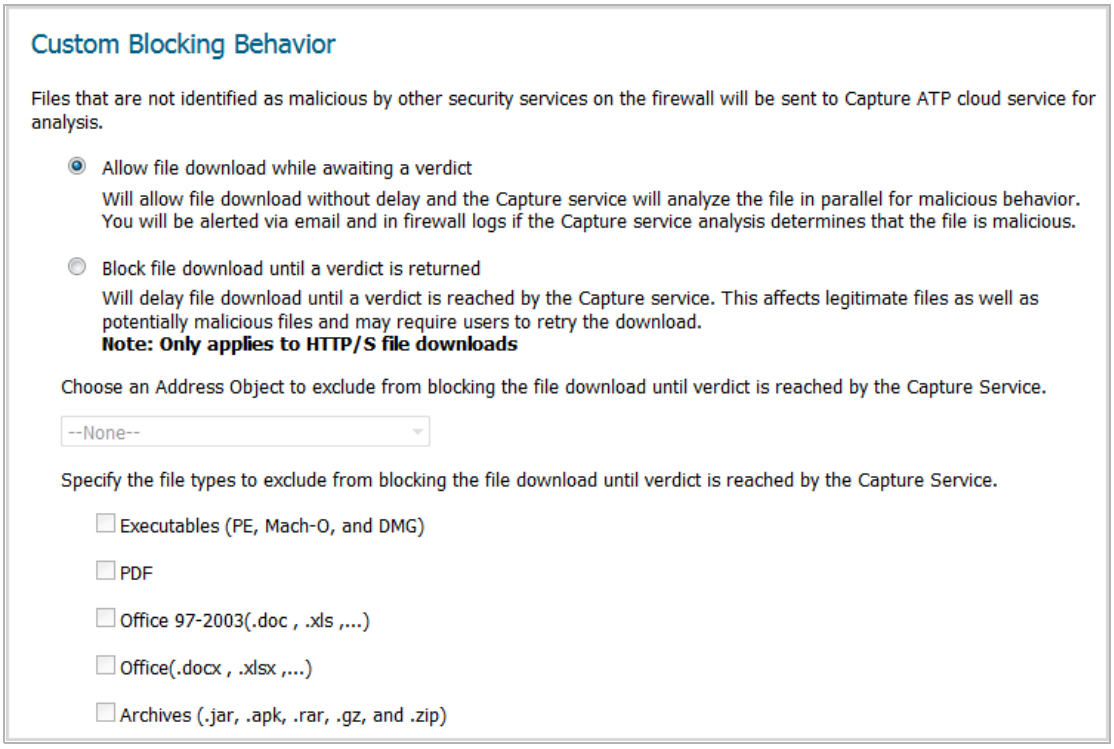

By default **Allow file download while awaiting a verdict** is selected.

IMPORTANT: The **Block file download until a verdict is returned** feature should only be enabled if the strictest controls are desired.

If you select this feature, a warning dialog appears.

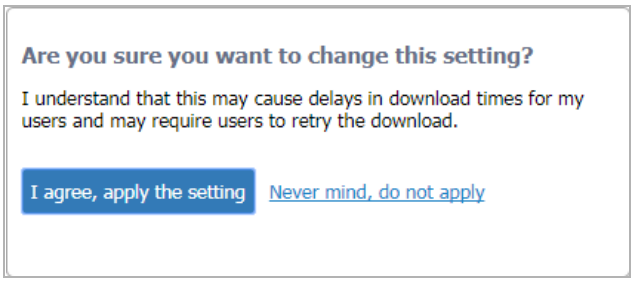

Clicking the:

- **I agree, apply the setting** button selects the **Block file download until a verdict is returned** option. You also must click the **Accept** button for the change to take effect.
- **Never mind, do not apply** link closes the dialog and leaves **Allow file download while awaiting a verdict** selected.

11 Click **ACCEPT**.

Disabling GAV or Cloud Anti-Virus

You can disable the Gateway Anti-Virus or Cloud Anti-Virus services by clearing the checkboxes for them on the **Security Configuration | Security Services > Gateway Anti-Virus** page. If you disable either service while Capture ATP is enabled, a popup message is displayed warning you that Capture ATP will also be disabled.

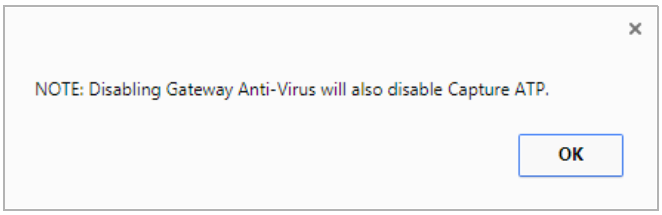

Capture ATP stops working if either Gateway Anti-Virus or Cloud Anti-Virus is disabled. For example, if Gateway Anti-Virus is not enabled, the **Capture ATP > Settings** page shows **You must enable Gateway Anti-Virus for Capture ATP to function**, along with a `manage settings` link that takes you to the **Security Services > Gateway Anti-Virus** page where you can enable it.

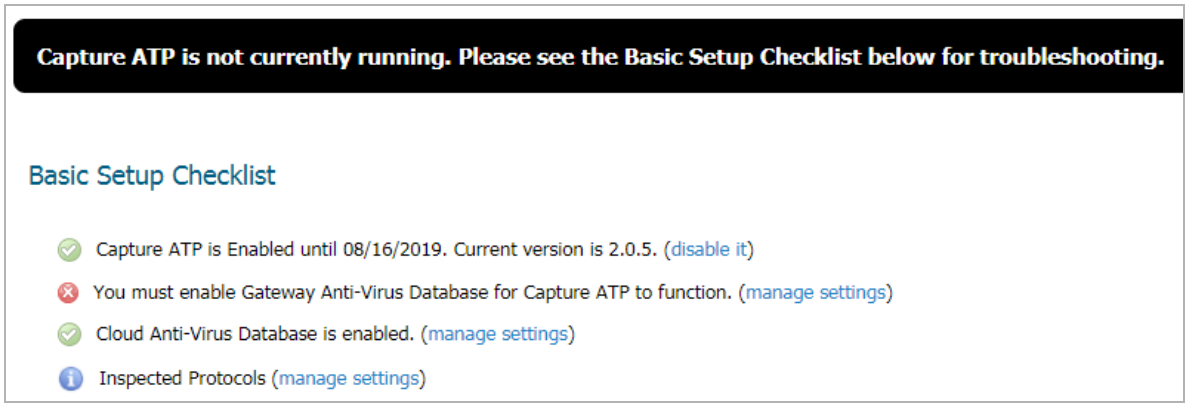

16

Activating Anti-Spyware Service

- About Anti-Spyware on page 202
- Security Services > Anti-Spyware on page 203
- Configuring Anti-Spyware Policies on page 207

About Anti-Spyware

SonicWall Anti-Spyware is part of the SonicWall Gateway Anti-Virus, Anti-Virus and Intrusion Prevention Service solution that provides comprehensive, real-time protection against viruses, worms, Trojans, spyware, and software vulnerabilities.

The SonicWall Anti-Spyware Service protects networks from intrusive spyware by cutting off spyware installations and delivery at the gateway and denying previously installed spyware from communicating collected information outbound. SonicWall Anti-Spyware works with other anti-spyware programs, such as programs that remove existing spyware applications from hosts. You are encouraged to use or install host-based anti-spyware software as an added measure of defense against spyware.

SonicWall Anti-Spyware analyzes inbound connections for the most common method of spyware delivery, ActiveX-based component installations. It also examines inbound setup executables and cabinet files crossing the gateway, and resets the connections that are streaming spyware setup files to the LAN. These file packages may be freeware bundled with adware, keyloggers, or other spyware.

If spyware has been installed on a LAN workstation prior to installing the Anti-Spyware service, the service will examine outbound traffic for streams originating at spyware infected clients and reset those connections. For example, when spyware has been profiling a user's browsing habits and attempts to send the profile information home, the firewall identifies that traffic and resets the connection.

The SonicWall Anti-Spyware Service provides the following protection:

- Blocks spyware delivered through auto-installed ActiveX components, the most common vehicle for distributing malicious spyware programs.
- Scans and logs spyware threats that are transmitted through the network and alerts administrators when new spyware is detected and/or blocked.
- Stops existing spyware programs from communicating in the background with hackers and servers on the Internet, preventing the transfer of confidential information.
- Provides granular control over networked applications by enabling administrators to selectively permit or deny the installation of spyware programs.
- Prevents emailed spyware threats by scanning and then blocking infected emails transmitted either through SMTP, IMAP or Web-based email.

Security Services > Anti-Spyware

The **MANAGE | Security Configuration > Security Services > Anti-Spyware** page displays the configuration settings for managing the service on your SonicWall security appliance.

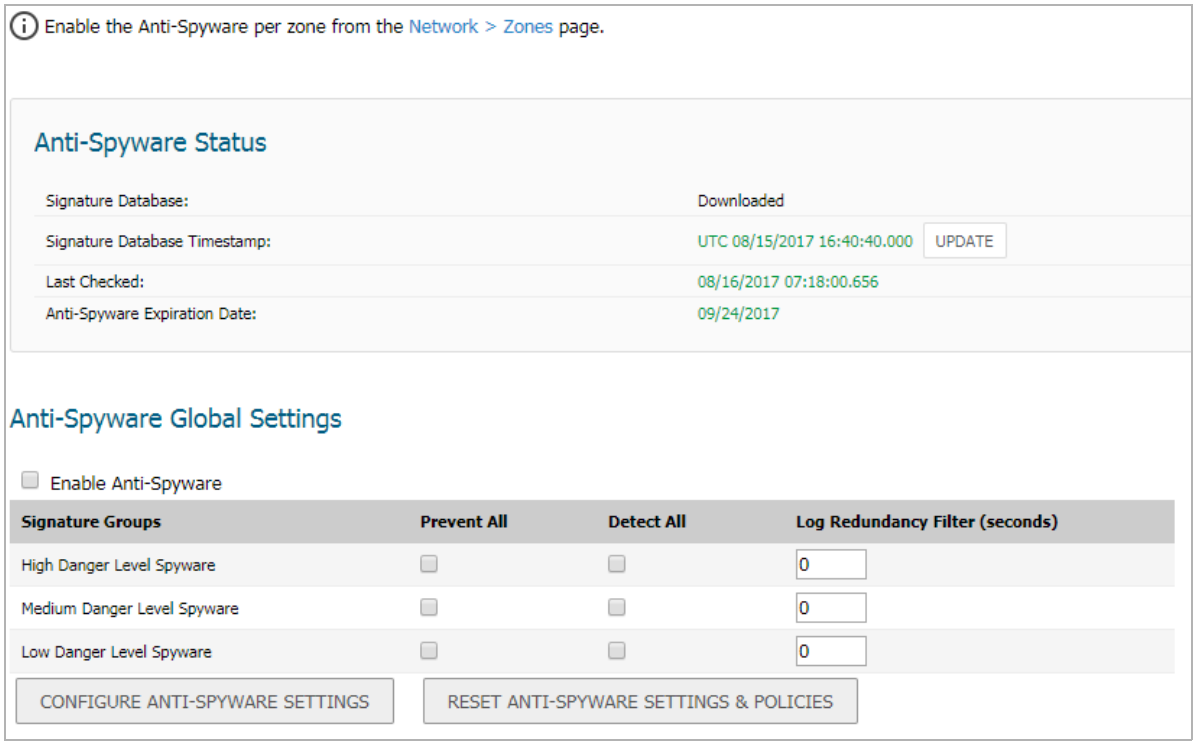

The **Security Services > Anti-Spyware** page is divided into three sections:

- **Anti-Spyware Status** – displays status information on the state of the signature database, your SonicWall Anti-Spyware license, and other information.
- **Anti-Spyware Global Settings** – provides the key settings for enabling SonicWall Anti-Spyware on your SonicWall security appliance, specifying global SonicWall Anti-Spyware protection based on three classes of spyware, and other configuration options.
- **Anti-Spyware Policies** – allows you to view SonicWall Anti-Spyware signatures and configure the handling of signatures by category groups or on a signature by signature basis. Categories are signatures grouped together based on the product or manufacturer.

NOTE: After activating your SonicWall Anti-Spyware license, you must enable and configure Anti-Spyware on the SonicWall management interface before anti-spyware policies are applied to your network traffic.

Topics:

- Anti-Spyware Status on page 204
- Anti-Spyware Global Settings on page 204
- Applying Anti-Spyware Protection on Zones on page 205
- Anti-Spyware Policies on page 206

Anti-Spyware Status

The **Anti-Spyware Status** section shows the state of the signature database, including the database's timestamp, and the time the SonicWall signature servers were last checked for the most current signatures. The SonicWall security appliance automatically attempts to synchronize the database on startup, and once every hour.

- **Signature Database** – indicates the signature database has been downloaded to the SonicWall security appliance.
- **Signature Database Timestamp** – displays the date and time the signature database was last updated. The **Signature Database Timestamp** is a timestamp for updates to the SonicWall Anti- Spyware signature database, not the last update to the SonicWall security appliance.
- **Last Checked** – displays the last time the SonicWall security appliance checked for signature updates.
- **Anti-Spyware Expiration Date** – displays your SonicWall Anti-Spyware license expiration date. If your SonicWall Anti-Spyware subscription expires, the SonicWall Anti-Spyware inspection is stopped and the SonicWall Anti-Spyware configuration settings are removed from the SonicWall security appliance. These settings are automatically restored after renewing your SonicWall Anti- Spyware license to the previously configured state.

The following note contains a link to the **MANAGE | Security Configuration > Network > Zones** page where you can configure Anti-Spyware on individual zones:

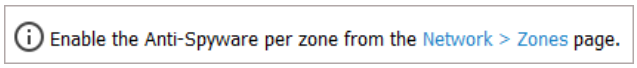

Anti-Spyware Global Settings

The **Anti-Spyware Global Settings** panel enables you to globally prevent and/or detect attacks based on the following attack levels:

- **High Danger Level Spyware** – These spyware applications are the most dangerous to your network, such as keyloggers or porn dialers, or may contain security vulnerabilities. Removal may be extremely difficult or impossible.
- **Medium Danger Level Spyware** – These spyware applications can cause disruption to your network, such as increased network traffic that slows down performance. Removal may be extremely difficult.
- **Low Danger Level Spyware** – These spyware applications are characterized by less intrusive activity and are not an immediate threat. They may profile users and usually are simple to remove.

TIP: SonicWall recommends enabling **Prevent All** for **High Danger Level Spyware** and **Medium Danger Level Spyware** to provide network protection against the most damaging spyware.

Anti-Spyware protection provides two methods for managing global spyware threats: detection (**Detect All**) and prevention (**Prevent All**). You must specify a **Prevent All** action in the Signature Groups panel for anti-spyware to occur on a global level on the SonicWall security appliance.

When **Prevent All** is enabled for a signature group in the **Signature Groups** panel, the SonicWall security appliance automatically drops and resets the connection to prevent the traffic from reaching its destination.

When **Detect All** is enabled for a signature group in the **Signature Groups** panel, the SonicWall security appliance logs and alerts any traffic that matches any signature in the group, but does not take any action against the traffic. The connection proceeds to its intended destination. You view the SonicWall log on the **Log >**

View page as well as configure how alerts are handled by the SonicWall security appliance in the **Log > Automation** page.

> **CAUTION: Be careful when selecting only Detect All. Selecting only Detect All logs and sends alerts on traffic that matches any signature in the group, but it does not take any action against the traffic. The traffic proceeds to its intended destination.**

When **Detect All** and **Prevent All** are both enabled for a signature group in the **Signature Groups** panel, the SonicOS logs and sends alerts on traffic that matches any signature in the group, and automatically drops and resets the connection to prevent the traffic from reaching its destination.

Enabling Inspection of Outbound Spyware Communication

The **Enable Inspection of Outbound Spyware Communication** option is available for scanning outbound traffic for spyware communication.

Applying Anti-Spyware Protection on Zones

If your firewall is running SonicOS, you can apply SonicWall Anti-Spyware to zones on the **Network > Zones** page to enforce Anti-Spyware not only between each network zone and the WAN, but also between internal zones. For example, enabling Anti-Spyware on the LAN zone enforces Anti-Spyware on all incoming and outgoing LAN traffic.

At the top of the **Security Services > Anti-Spyware** page, click the **Network > Zones** link to access the **MANAGE | System Setup > NNetwork > Zones** page. You apply Anti-Spyware to one of the zones listed on the **Network > Zones** page.

To enable Anti-Spyware on a zone:

1. In the firewall management interface, navigate to the **MANAGE | System Setup > Network > Zones** page. (Or from the **MANAGE | Security Configuration > Security Services > Anti-Spyware** page, click the **Network > Zones** link.) The **MANAGE | System Setup > Network > Zones** page displays.
2. In the **Configure** column in the **Zone Settings** panel, click the **Edit** icon for the zone you want to apply SonicWall Anti-Spyware. The **Edit Zone** dialog displays.
3. Click the **Enable Anti-Spyware** option. A checkmark appears. To disable SonicWall Anti-Spyware, clear the option.
4. Click **OK**.

   You can also enable SonicWall Anti-Spyware protection for new zones you create on the **MANAGE | Security Configuration > Security Services > Anti-Spyware** page. Clicking the **Add** icon displays the **Add Zone** dialog, which includes the same settings as the **Edit Zone** dialog.

Anti-Spyware Policies

The **Anti-Spyware Policies** section allows you to view and manage how SonicWall Anti-Spyware handles signatures by category groups or on a signature by signature basis. Categories are signatures grouped together by product or manufacturer, and they are listed in the **View Style** menu.

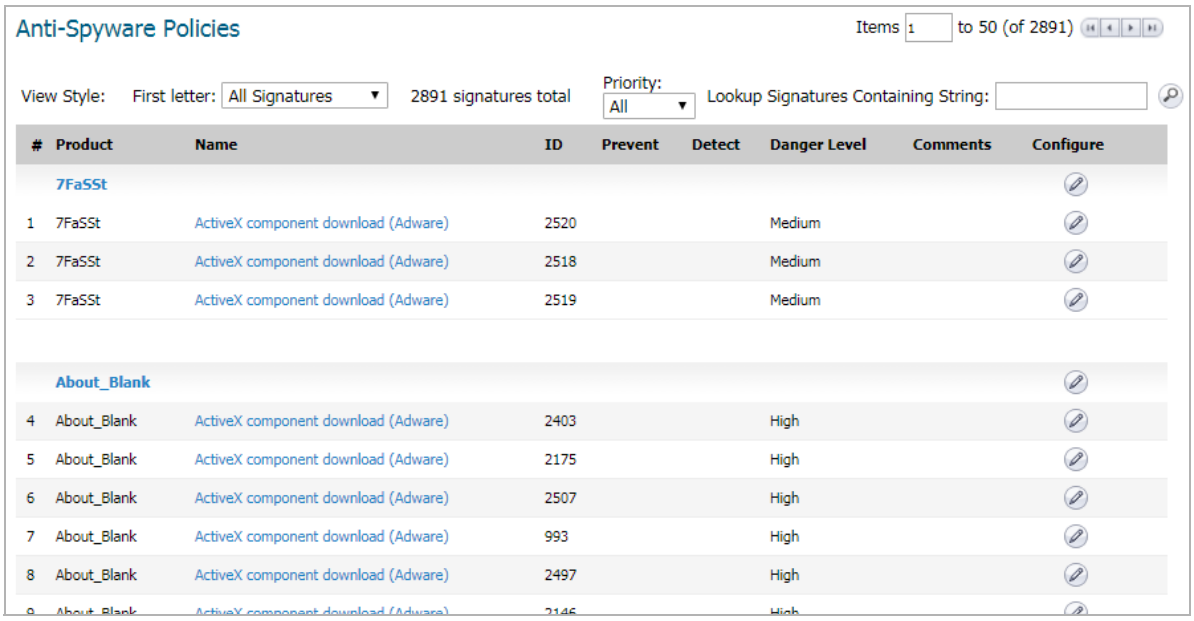

Entries listed in the **Anti-Spyware Policies** panel are from the SonicWall Anti-Spyware signature database downloaded to your firewall. Categories and signatures are dynamically updated by the Anti-Spyware Service. Categories and signatures dynamically change over time in response to new threats.

You can display the signatures in a variety of views using the **View Style** menu. This menu allows you to specify the categories or signatures to display in the **Anti-Spyware Policies** panel. You can select **All Signatures**, or you can select the first letter or number in the spyware name.



Selecting **All Signatures** from the menu displays all of the signatures by category. The **Anti-Spyware Policies** panel displays all the categories and their signatures. The category headers divide the signature entries. These headers display **Global** in the **Prevent** and **Detect** columns, indicating the global settings that you defined in the **Anti-Spyware Global Settings** section.

Topics:

- Anti-Spyware Policies Panel on page 206
- Displaying Spyware Information on page 207
- Searching the Signature Database on page 207
- Sorting Category or Signature Entries on page 207

Anti-Spyware Policies Panel

The **Anti-Spyware Policies** panel displays the following information about each signature entry:

- **Product** - Displays the spyware name or manufacturer.

- **Name** - Displays the name of the spyware as a link. Clicking the name link displays the SonicAlert information about the spyware.
- **ID** - The SonicWall database ID number of signature.
- **Prevent** - A check mark in this column indicates prevention is enabled. A green check mark appears in the **Detect** column any time you make a change from the global or category prevention settings.
- **Detect** - A check mark in this column indicates detection is enabled. A green check mark appears in the **Detect** column any time you make a change from the global or category detection settings.
- **Danger Level** - Defines the attack signature as **Low**, **Medium**, or **High** as defined for the **Signature Groups** panel.
- **Comments** - Displays a brief description of the policy.
- **Configure** - Clicking the edit icon in the **Configure** column of the category header displays the **Edit Anti-Spyware Category** window. Clicking the edit icon in the **Configure** column for an individual signature displays the **Edit Anti-Spyware Signature** window. These windows allow you to define a different action from the global settings for the specific category or signature.

Displaying Spyware Information

In the **Anti-Spyware Policies** panel, clicking on the spyware name link in **Name** column, displays a **SonicALERT** page that provides detailed information about the spyware.

Searching the Signature Database

You can search the signature database by entering a search string in the **Lookup Signatures Containing String** field, then clicking icon.

Sorting Category or Signature Entries

Clicking on the **Anti-Spyware Policies** panel headings (**Name**, **ID**, **Prevent**, **Detect**, or **Danger** Level) sorts the panel entries according to the heading. An up arrow by the column header name indicates the entries are sorted in descending order. A down arrow by the column header name indicates the entries are sorted in ascending order.

Configuring Anti-Spyware Policies

Topics:

- Configuring Category Policies on page 207
- Configuring Signature Policies on page 209

Configuring Category Policies

You can choose to override the global prevention and detection settings on a category-by-category basis. The global **Prevent All** and **Detect All** settings, which include **High Danger Level Spyware, Medium Danger Level Spyware**, and **Low Danger Level Spyware** are configured in the **Anti-Spyware Global Settings** section. Categories can include any combination of Danger Levels as defined in the **Signature Groups** panel.

The available signature categories are listed in the **View Style** menu in the **Anti-Spyware Policies** section. Configuring the prevent and detect behaviors on a category basis affects all the signatures in the category, regardless of the global attack priority settings (Low, Medium, or High)

Topics:

- Overriding Global Prevent and Detect Settings by Category on page 208
- Resetting SonicWall Anti-Spyware Configuration to Default on page 209

Overriding Global Prevent and Detect Settings by Category

1. Select **All categories** or an individual category from the **Category** menu.
2. If you select **All Categories**, click on the **Edit** icon in the **Configure** column for the category you want to change. the **Edit Anti-Spyware Category** dialog is displayed.
3. If you select an individual category, click on the **Edit** icon to the right of the **Category** menu. The **Edit Anti-Spyware Category** dialog displays.
4. If you want to change the Global Setting for **Prevention**, select **Enable** or **Disable** from the **Prevention** menu.
5. If you want to change the Global Setting for **Detection**, select **Enable** or **Disable** from the **Detection** menu.
6. If you want to change the Global Settings for both detection and prevention, select **Enable** or **Disable** from the **Detection** and **Prevention** menu.
7. The following settings allow you to select specific users/groups, IP address ranges, and schedule objects to be included or excluded from this SonicWall Anti-Spyware category:
    - **Included Users/Groups** - select the Users/Groups you want included in this SonicWall Anti-Spyware category. The default is **All**.
    - **Excluded Users/Groups** - select the Users/Groups you want excluded from this SonicWall Anti-Spyware category. The default **None**.
    - **Included IP Address Range** - select the IP address range you want included in this SonicWall Anti-Spyware category. The default **All**.
    - **Excluded IP Address Range** - select the IP address range you want excluded from this SonicWall Anti-Spyware category. The default **None**.
    - **Schedule** - select the scheduled time you want for the activation of this SonicWall Anti-Spyware category. The default **Always on**.
8. If you want to change the Log Redundancy Filter setting from the default global setting, uncheck the **Use Category Settings** box for **Log Redundancy Filter (seconds)** and enter a time value in seconds.
9. Click **OK** to save your changes.

TIP: If you select **All signatures** from the **Category** menu, all the categories and their signatures are displayed in the **Anti-Spyware Policies** panel, allowing you to configure both the category and signatures within the category.

Resetting SonicWall Anti-Spyware Configuration to Default

You can remove all custom category and signature settings you created as well as reset global **Prevent All** and **Detect All** settings and **Log Redundancy Filter (seconds)** settings by clicking the **Reset Anti-Spyware Settings & Policies** button in the **Anti-Spyware Global Settings** section.

Configuring Signature Policies

Selecting **All signatures** from the **Category** menu displays all of the signatures organized within categories. The **All signatures** option displays every signature in the Anti-Spyware database.

If global **Prevent All** and **Detect All** settings are in effect for the category, **Global** is displayed in the **Prevent** and **Detect** columns for the category and all of its signatures.

Selecting a specific signature category, displays the signatures in that category.

NOTE: You cannot import your own customized signatures into SonicWall Anti-Spyware or delete a signature entry.

CAUTION: Use caution when overriding global High Danger Level Spyware and Medium Danger Level Spyware signature behaviors because you can create vulnerabilities. If you make changes and want to restore the default global signature settings, click the Reset Anti-Spyware Settings & Policies button to restore the default settings.

Topics:

- Overriding Global Prevent and Detect Settings by Category on page 208
- Resetting SonicWall Anti-Spyware Settings to Default on page 210

Overriding Category Detect and Prevent Settings for a Signature

To override category detect and prevent attributes for signatures:

1. In the **Anti-Spyware Policies** panel, display the signature you want to change. Click the **Edit** icon in the **Configure** column for the entry to display the **Edit Anti-Spyware** dialog.
2. If you want to change the Category Setting for **Prevention**, select **Enable** or **Disable** from the **Prevention** menu.
3. If you want to change the Category Setting for **Detection**, select **Enable** or **Disable** from the **Detection** menu.
4. If you want to change the Category Setting for both detection and prevention, select **Enable** or **Disable** from the **Detection** and **Prevention** menu.
5. The following settings allow you to select specific users/groups, IP address ranges, and schedule objects to be included or excluded from this SonicWall Anti-Spyware signature:
    - **Included Users/Groups** - select the Users/Groups you want included in this SonicWall Anti-Spyware signature. The default is **All**.
    - **Excluded Users/Groups** - select the Users/Groups you want excluded from this SonicWall Anti-Spyware signature. The default **None**.

- **Included IP Address Range** - select the IP address range you want included in this SonicWall Anti-Spyware signature. The default **All**.
- **Excluded IP Address Range** - select the IP address range you want excluded from this SonicWall Anti-Spyware signature. The default **None**.
- **Schedule** - select the scheduled time you want for the activation of this SonicWall Anti-Spyware signature. The default **Always on**.

6 If you want to change the Log Redundancy Filter setting from the Category setting, uncheck the **Use Category Settings** box for **Log Redundancy Filter (seconds)** and enter a time value in seconds.

7 Click **OK** to save your changes.

Resetting SonicWall Anti-Spyware Settings to Default

You can remove all custom category and signature settings you created as well as reset global **Prevent All** and **Detect All** settings and **Log Redundancy Filter (seconds)** settings by clicking the **Reset Anti-Spyware Settings & Policies** button in the **Anti-Spyware Global Settings** section.

17

Configuring SonicWall Real-Time Blacklist

- Security Services > RBL Filter on page 211
- About Real-Time Black List Filtering on page 212
- Configuring the RBL Filter on page 212

Security Services > RBL Filter

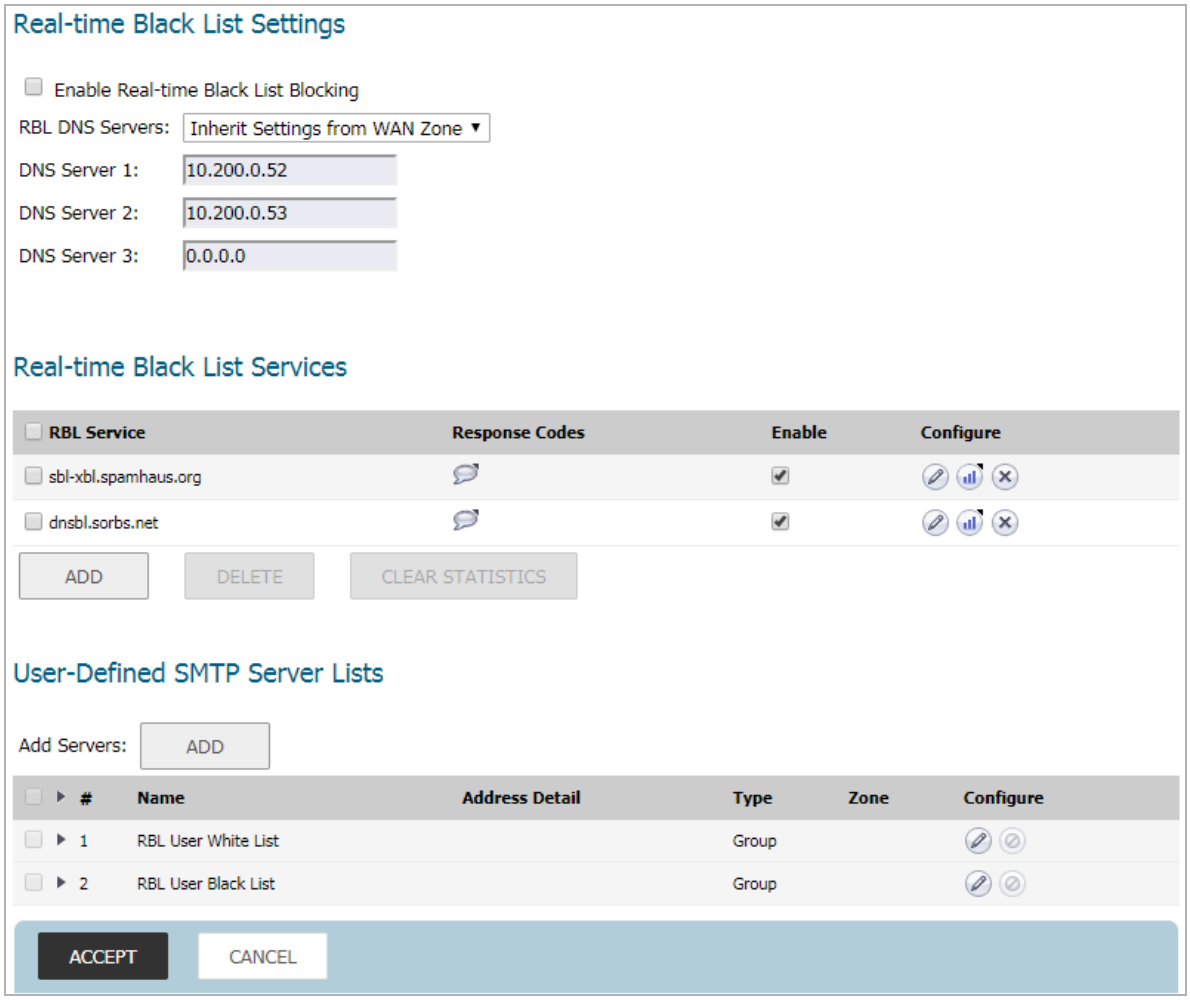

About Real-Time Black List Filtering

SMTP Real-Time Black List (RBL) is a mechanism for publishing the IP addresses of SMTP spammers use. There are a number of organizations that compile this information both for free: http://www.spamhaus.org, and for profit: https://ers.trendmicro.com/.

NOTE: SMTP RBL is an aggressive spam filtering technique that can be prone to false-positives because it is based on lists compiled from reported spam activity. The SonicOS implementation of SMTP RBL filtering provides a number of fine-tuning mechanisms to help ensure filtering accuracy.

RBL list providers publish their lists using DNS. Blacklisted IP addresses appear in the database of the list provider's DNS domain using inverted IP notation of the SMTP server in question as a prefix to the domain name. A response code from 127.0.0.2 to 127.0.0.9 indicates some type of undesirability:

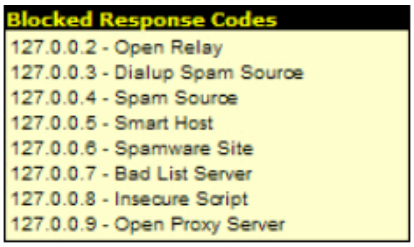

For example, if an SMTP server with IP address 1.2.3.4 has been blacklisted by RBL list provider sbl-xbl.spamhaus.org, then a DNS query to 4.3.2.1.sbl-xbl.spamhaus.org will provide a 127.0.0.4 response, indicating that the server is a known source of spam, and the connection will be dropped.

NOTE: Most spam today is known to be sent from hijacked or zombie machines running a thin SMTP server implementation.Unlike legitimate SMTP servers, these zombie machines rarely attempt to retry failed delivery attempts. Once the delivery attempt is blocked by RBL filter, no subsequent delivery attempts for that same piece of spam will be made.

Configuring the RBL Filter

Topics:

- Enabling RBL Blocking on page 212
- Adding RBL Services on page 213
- Configuring User-Defined SMTP Server Lists on page 214
- Testing SMTP IP Addresses on page 216

Enabling RBL Blocking

When **Enable Real-time Black List Blocking** is enabled in the **Real-time Black List Settings** section on the **RBL Filter** page, inbound connections from hosts on the WAN or outbound connections to hosts on the WAN are

checked against each enabled RBL service with a DNS request to the DNS servers configured under **RBL DNS Servers**.

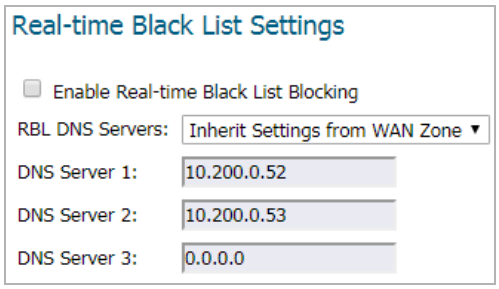

The RBL DNS Servers menu allows you to specify the DNS servers. You can choose I**nherit Settings from WAN Zone** or **Specify DNS Servers Manually**. If you select **Specify DNS Servers Manually**, enter the DNS server addresses in the **DNS Server** fields.

When you have finished, click **ACCEPT**.

The DNS responses are collected and cached. If any of the queries result in a blacklisted response, the server will be filtered. Responses are cached using TTL values, and non-blacklisted responses are assigned a cache TTL of 2 hours. If the cache fills up, then cache entries are discarded in a FIFO (first-in-first-out) fashion.

The IP address check uses the cache to determine if a connection should be dropped. Initially, IP addresses are not in the cache and a DNS request must be made. In this case the IP address is assumed innocent until proven guilty, and the check results in the allowing of the connection. A DNS request is made and results are cached in a separate task. When subsequent packets from this IP address are checked, if the IP address is blacklisted, the connection will be dropped.

Adding RBL Services

You can add additional RBL services in the **Real-time Black List Services** section.

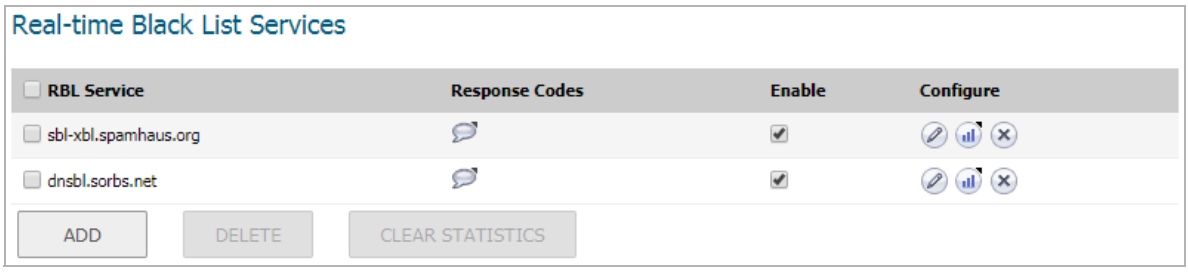

To add an RBL service, click the **ADD** button. In the **Add RBL Domain** dialog, you specify the RBL domain to be queried, enable it for use, and specify its expected response codes. Most RBL services list the responses they provide on their Web site, although selecting **Block All Responses** is generally acceptable.

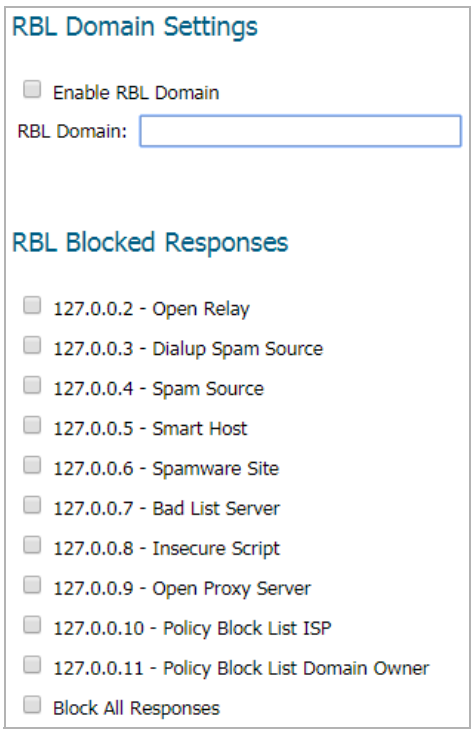

Statistics are maintained for each RBL Service in the **RBL Service** table, and can be viewed with a mouseover of the (statistics) icon to the right on the service entry.

Configuring User-Defined SMTP Server Lists

The **User Defined SMTP Server Lists** section allows for Address Objects to be used to construct a white-list (explicit allow) or black-list (explicit deny) of SMTP servers. Entries in this list will bypass the RBL querying procedure.

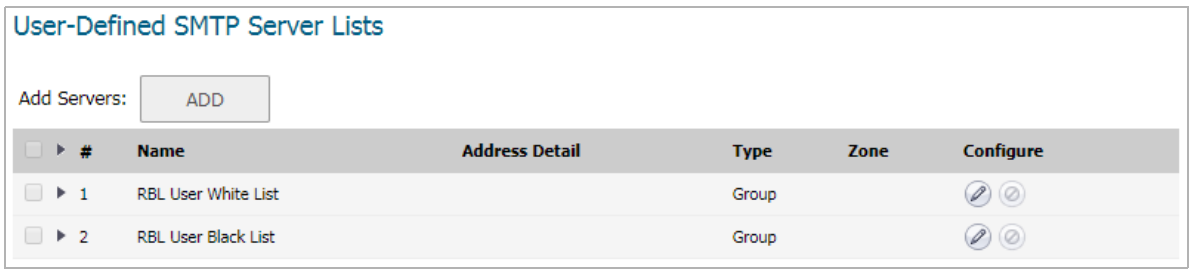

NOTE: To see entries in the RBL User White List and RBL User Black List, click the arrow to the right of the checkbox for that list.

Topics:

- Configuring a White List on page 215
- Configuring a Black List on page 215

Configuring a White List

For example, to ensure that you always receive SMTP connections from a partner site's SMTP server:

1 Create an Address Object for the server using the **Add Servers: ADD** button. the **Add Address Object** dialog appears.

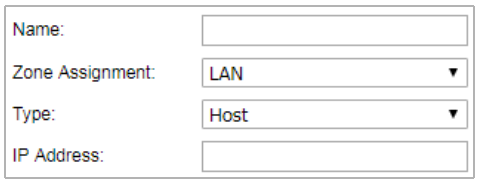

2 Configure the Address Object.

3 Click **OK**. The Address Object is added to the **RBL User White List** in the **User-Defined SMTP Server Lists** table.

4 Click the **edit** icon in the **Configure** column of the **RBL User White List** row. The **Edit Address Object** window displays.



5 Add the Address Object by selecting it and clicking the right arrow.

6 Click **OK**.

The table is updated, and that server is always allowed to make SMTP exchanges.

Configuring a Black List

1 Click the **Edit** icon in the **Configure** column of the **RBL User Black List** row. The **Edit Address Object** dialog displays.



2 Add the Address Object by selecting it and clicking the right arrow.

3 Click **OK**.

Testing SMTP IP Addresses

The **INVESTIGATE | Tools > System Diagnostics** page also provides a **Real-time Black List Lookup** feature that allows for SMTP IP addresses (or RBL services, or DNS servers) to be specifically tested. For more information about this page, see *SonicWall SonicOS 6.5 Investigate*.

For a list of known spam sources to use in testing, refer to: http://www.spamhaus.org/sbl/latest/.

18

Configuring Geo-IP Filters

NOTE: The Geo-IP Filtering feature is available on TZ300 series and above appliances.

- Security Services > Geo-IP Filter
- Configuring Geo-IP Filtering
- Creating a Custom Country List
- Customizing Web Block Page Settings
- Using Geo-IP Filter Diagnostics

Security Services > Geo-IP Filter

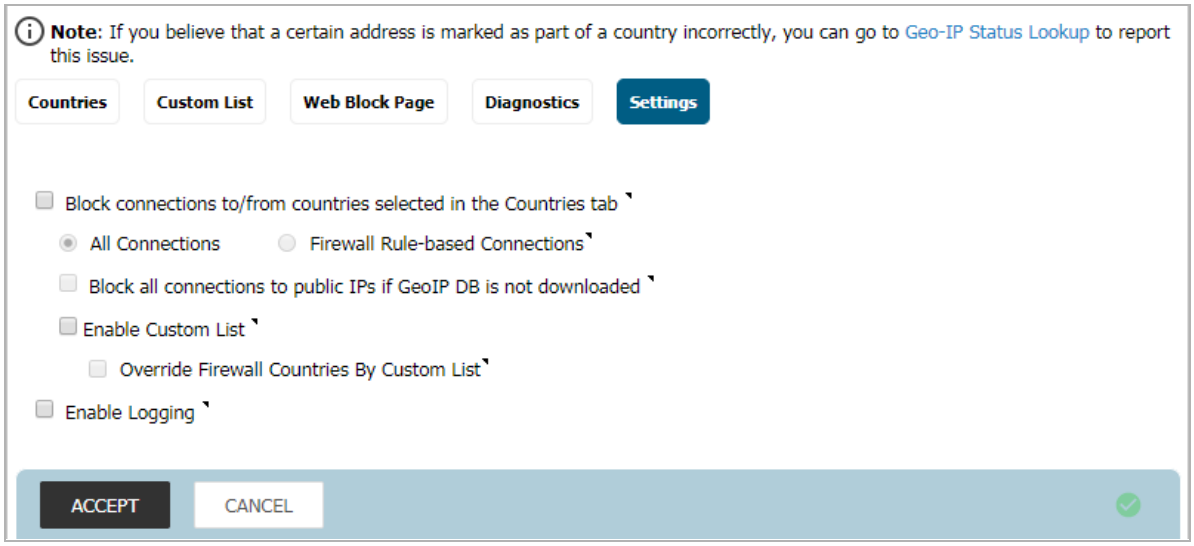

The Geo-IP Filter feature allows you to block connections to or from a geographic location. The SonicWall firewall uses the IP address to determine to the location of the connection. The GEO-IP Filter feature also allows you to create custom country lists that affect the identification of an IP address.

The Geo-IP Filter feature also allows you to create a custom message when you block a web site.

You can also use the Geo-IP Filter Diagnostics tool to show resolved locations, monitor Geo-IP cache statistics, custom countries statistics, and look up GEO-IP servers.

Configuring Geo-IP Filtering

To configure Geo-IP Filtering:

1 Navigate to **MANAGE | Security Configuration > Security Services > GEO-IP Filter** page.

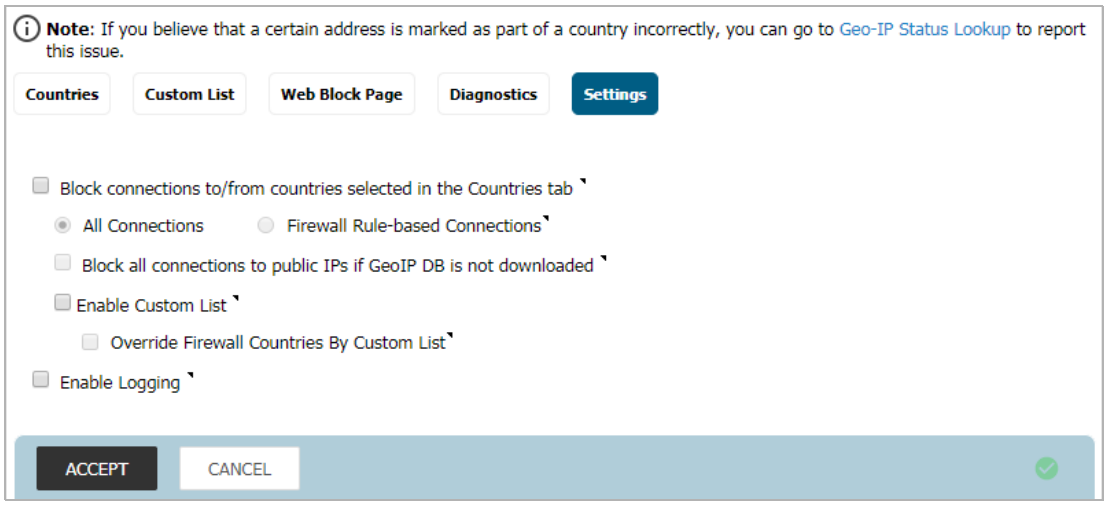

2 To block all connections to and from specific countries, select the **Block connections to/from countries listed in the table below** checkbox. This option is selected by default.

If this option is enabled, all connections to/from the selected list of countries are blocked. You can specify an exclusion list to exclude this behavior for selected IPs, as described below in Step 10.

When this option is selected, the next two options become available.

3 Select one of the following two modes for Geo-IP Filtering:

- **All Connections**: All connections to and from the firewall are filtered. This option is selected by default.
- **Firewall Rule-Based Connections**: Only connections that match an access rule configured on the firewall are filtered for blocking.

4 To block all connections to public IPs when the Geo-IP database is not downloaded, select the **Block all connections to public IPs if GeoIP DB is not downloaded** option. This option is not selected by default.

5 To enable your custom list, select the **Enable Custom List** checkbox. This option is not selected by default.

If the **Enable Custom List** checkbox is:

- Not selected, then only the firewall's country database is searched. Go to Step 6.
- Selected, the **Override Firewall Countries By Custom List** checkbox becomes available.

Enabling a custom list by selecting the **Enable Custom List** checkbox can affect country identification for an IP address. If the **Override Firewall Countries By Custom List** is:

- Not selected also, then country identification is done in this order:
  1) The firewall country database is searched. If the identification is not resolved, then:
  2) The custom country list is searched.
- Also selected, then country identification is done in this order:
  1) The custom country database is searched. If the identification is not resolved, then:

2) The firewall country list is searched.

In either case, action is taken according to the resolution.

6 To log Geo-IP Filter-related events, select **Enable logging**. This option is not selected by default.

7 Under **Countries**, in the **Selected Countries** table, select the countries to be blocked. By default, no countries are blocked.

8 Drag the selected countries in the **Available Countries** table to the **Selected Countries** table.

NOTE: Blocked countries are highlighted when selected in the **Available Countries** table.

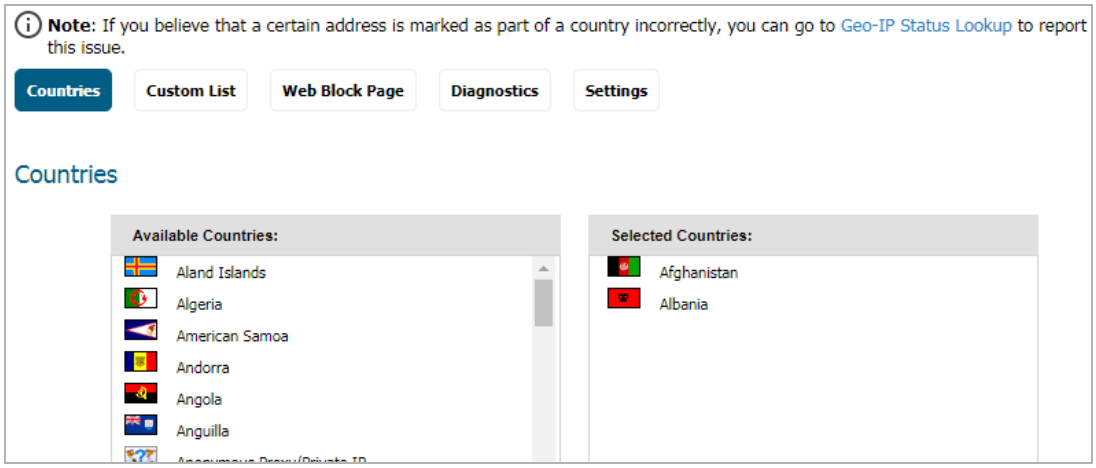

9 If you want to block any countries that are not listed, select the **Block All UNKNOWN countries** option. All connections to unknown public IPs are blocked. This option is not selected by default.

10 Optionally, you can configure an exclusion list of all connections to approved IP addresses by doing one of these:

- Select an address object or address group from the **Geo-IP Exclusion Object** drop-down menu. The default is **Default Geo-IP and Botnet Exclusion Group**.
- Create a new address object or address group by selecting **Create new address object…** or **Create new address group…** from the **Geo-IP Exclusion Object** drop-down menu.

  The **Geo-IP Exclusion Object** is a network address object group that specifies a group or a range of IP addresses to be excluded from the Geo-IP filter blocking. All IP addresses in the address object or group are allowed, even if they are from a blocked country.

  For example, if all IP addresses coming from Country A are set to be blocked and an IP address from Country A is detected, but it is in the **Geo-IP Exclusion Object** list, then traffic to and from this IP address is allowed to pass.

  For this feature to work correctly, the country database must be downloaded to the firewall. The **Status** icon at the top right of the **Custom List** page turns yellow if this download fails. Green

status indicates that the database has been successfully downloaded. Click the **Status** icon to display more information.





For the country database to be downloaded, the firewall must be able to resolve the address, `utmgbdata.global.sonicwall.com`.

When a user attempts to access a web page that is from a blocked country, a block page message is displayed on the user's web browser.

NOTE: If a connection to a blocked country is short-lived and the firewall does not have a cache for the IP address, then the connection may not be blocked immediately. As a result, connections to blocked countries may occasionally appear in the App Flow Monitor. However, additional connections to the same IP address are blocked immediately.

11 Click **ACCEPT** to enable your changes.

Creating a Custom Country List

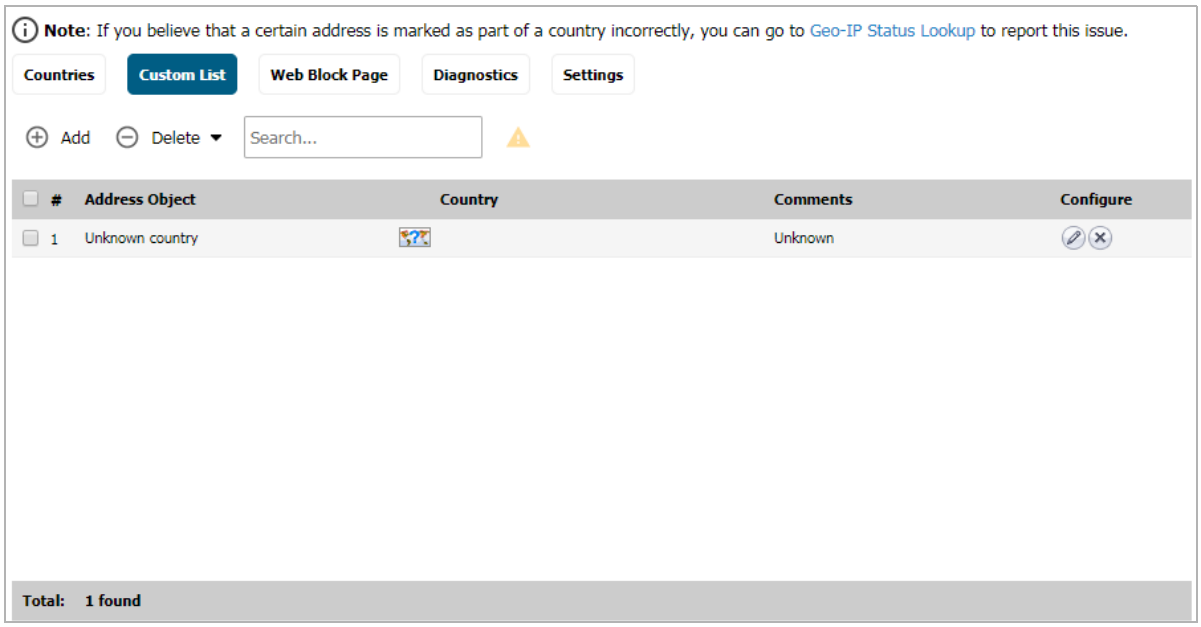

Address Object	Name given to the address object.
Country	Flag icon (if known) and name of country.

Comments	Comment made when address object was created.
Configure	Contains an **Edit** icon and a **Delete** icon.
Total	Displays the number of entries in the **Custom List**.

An IP address can be associated with a wrong country. This kind of misclassification can cause incorrect/unwanted filtering of an IP address. Having a custom country list can solve this problem by overriding the firewall country associated with a particular IP address.

Topics:

- Creating a Custom List on page 221
- Editing a Custom List Entry on page 223
- Deleting Custom List Entries on page 223

Creating a Custom List

IMPORTANT: For the firewall to use the custom country list, you must enable it as described in Configuring Geo-IP Filtering on page 218.

To create a custom country list:

1 Navigate to **MANAGE | Security Configuration > Security Services > GEO-IP Filter.**

2 Click **Settings**.

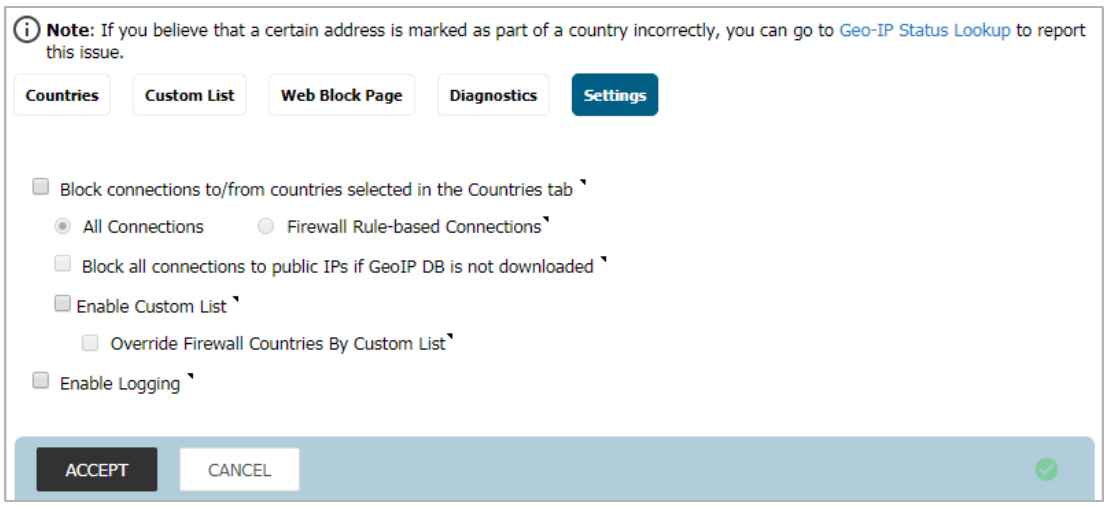

3 Select **Enable Custom List**.

4 Click **Custom List**.

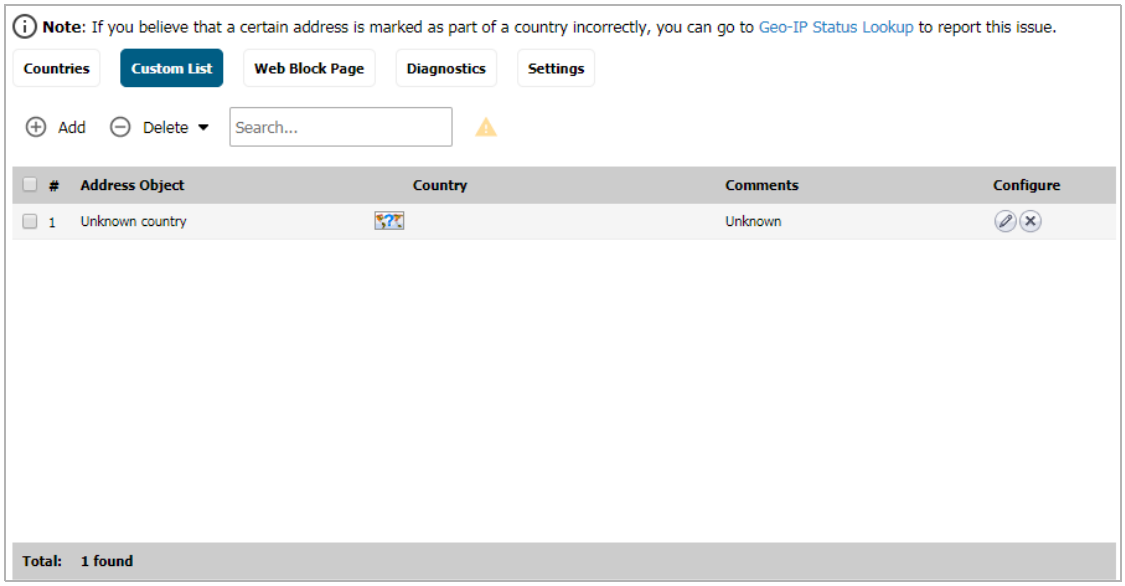

5 Click the **Add** icon. The **Add Custom List** dialog displays.

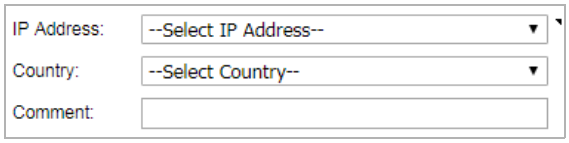

6 Select an IP address object or create a new address object from the **IP Address** drop-down menu:

IMPORTANT: An address object cannot overlap any other address objects in the custom country list. Different address objects, however, can have the same country ID.

- **Create new address object...** – the **Add Address Object** dialog displays.

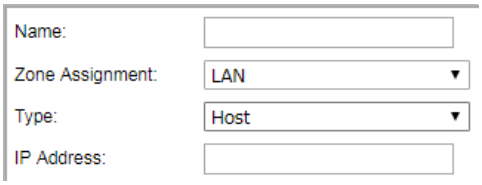

You create a new address object as described in *SonicWall SonicOS 6.5 Policies*, with these restrictions:

- Allowed types are
  - **Host**
  - **Range**
  - **Network**
  - A group of any combination of these types

All other types are disallowed types and cannot be added to the custom country list.

- **Create new address group...** – the **Add Address Object Group** dialog displays.

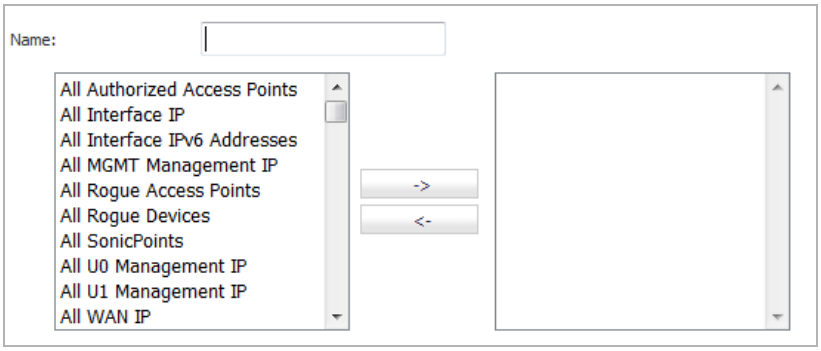

  You create a new address object as described in *SonicWall SonicOS 6.5 Policies*

- Already defined address object or address group.

7 Select a country from the **Country** drop-down menu.

8 Optionally, add a comment in the **Comment** field.

9 Click **OK**.

Editing a Custom List Entry

To edit a custom list entry:

1 Navigate to **MANAGE | Security Configuration > Security Services > GEO-IP Filter.**

2 Click **Custom List**.

3 Click the **Edit** icon in the **Configure** column for the entry to be edited. The **Add Custom List** dialog displays with the IP address and any comment about the entry.

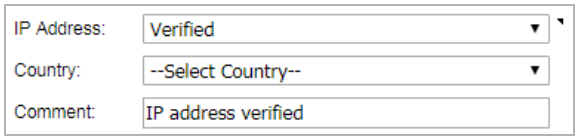

4 Select the country from the **Country** drop-down menu and make any other changes.

5 Click **OK**. The **Custom List** table is updated.

Deleting Custom List Entries

To delete a custom list entry:

1 Do one of these:

- Click the **Delete** icon in the **Configure** column for the entry.
- Select the checkbox for the entry and then click the **Delete** button.

A confirmation message displays.

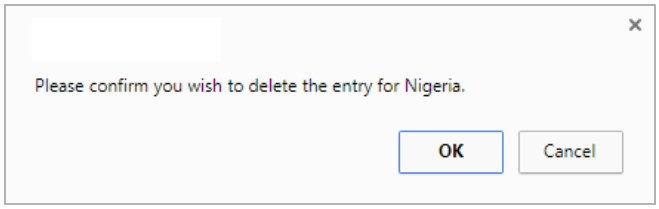

2 Click **OK**.

To delete multiple entries:

1 Select the checkboxes of the entries to be deleted. The **Delete** button becomes available.

2 Click the **Delete** button. A confirmation message displays.

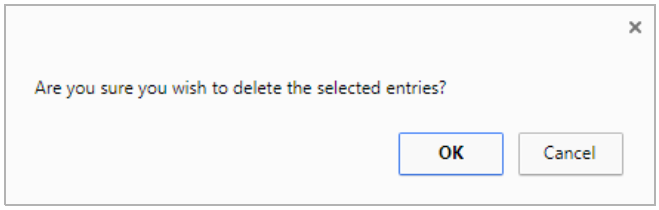

3 Click **OK**.

To delete all entries:

1 Click the checkbox in the table header.

2 Click the **Delete** button. A confirmation message displays.

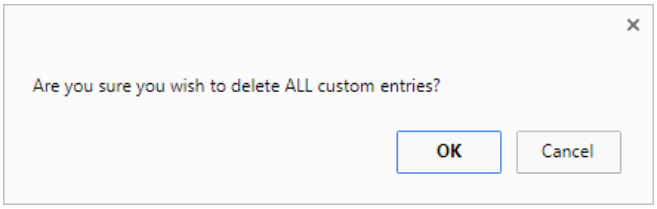

3 Click **OK**.

Customizing Web Block Page Settings

The Geo-IP Filter has a default message that is displayed when a user attempts to access a blocked page. You can have the message display detailed information, such as the reason why this IP address is blocked as well as the IP address and the country from which it was detected. You also can create a custom message and include a custom logo.

To create a custom web-block message:

1 Navigate to **MANAGE | Security Configuration > Security Services > GEO-IP Filter**.

2 Click **Settings**.

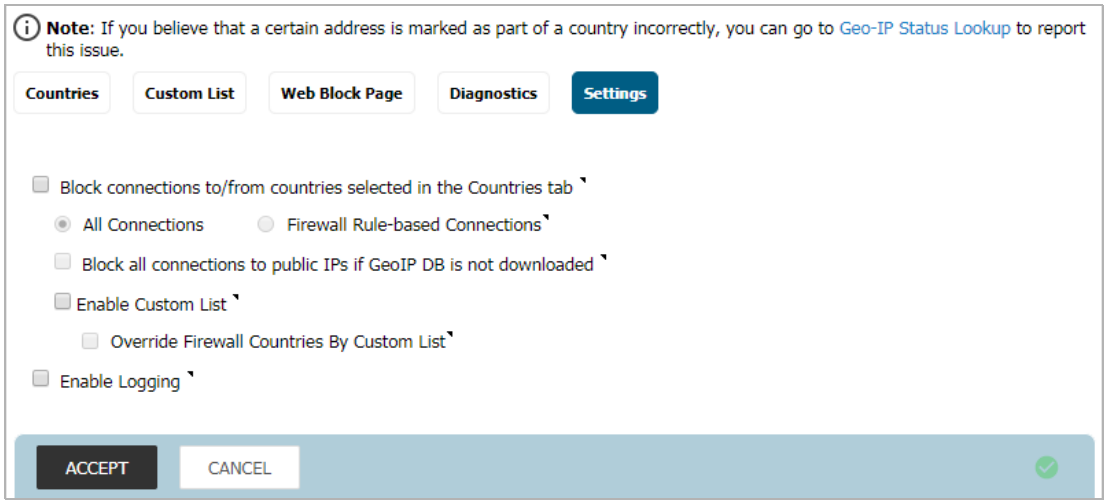

3 Click **Web Block Page**.

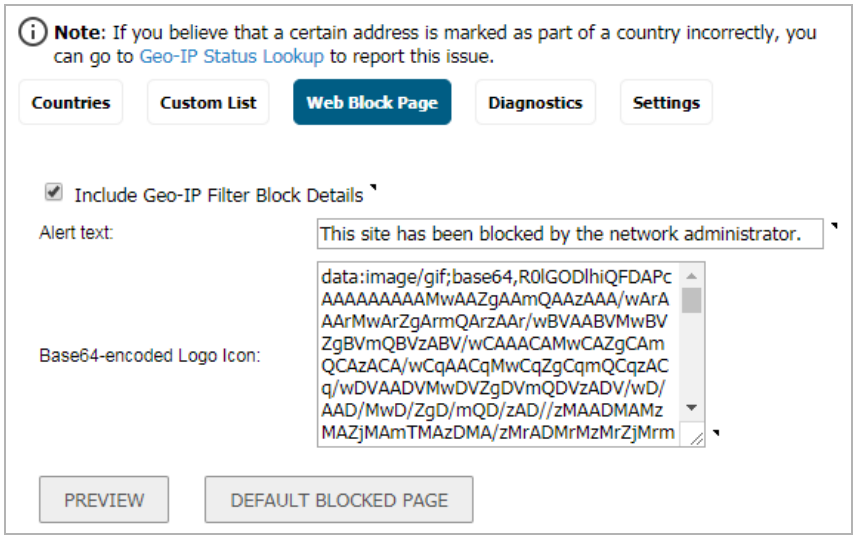

4 Ensure the **Include Geo-IP Filter Block Details** option is selected. When enabled, this option shows block details such as reason for the block, IP address, and country. When disabled, no information is displayed. By default, this option is selected. This option is selected by default.

5 Do one of the following:

- To use the default message displayed in the **Alert text** field, `This site has been blocked by the network administrator.`, click the **Default Blocked Page** button and then go to Step 7.
- Specify a custom message to be displayed in the Geo-IP Filter Block page in the **Alert text** field. Your message can be up to 100 characters long.

6 Optionally, in the **Base64-encoded Logo Icon** field, you can specify a Base 64-encoded GIF icon to be displayed instead of the default SonicWall logo.

NOTE: Ensure the icon is valid and make the size as small as possible. The recommended size is 400 x 65.

7 To see a preview of your customized message and logo (or the default message and logo), click the **Preview** button. A warning message displays.

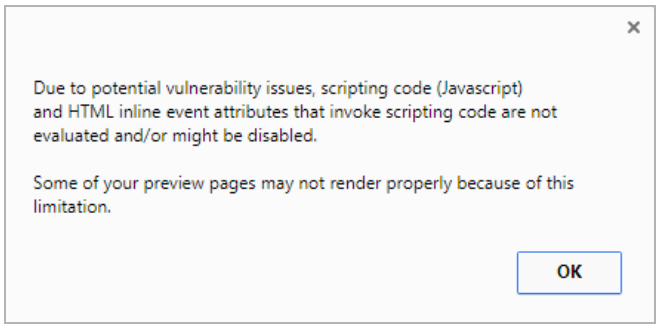

8 Click **OK**. The **Web Site Blocked** message displays.



9 Close the **Web Site Blocked** message.

10 Click **ACCEPT**.

Using Geo-IP Filter Diagnostics

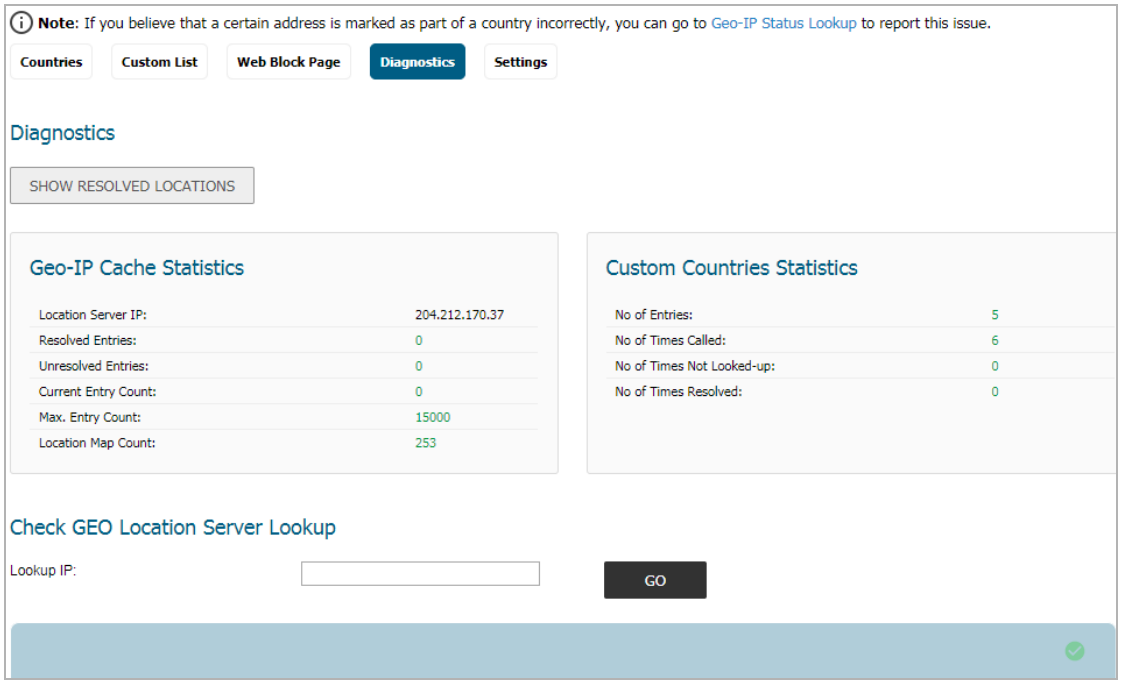

The **Security Services > GEO-IP Filter** page has a **Diagnostics** view with several tools:

- Show Resolved Locations on page 227
- Geo-IP Cache Statistics on page 227
- Custom Countries Statistics on page 228

- Check GEO Location Server Lookup on page 228
- Incorrectly Marked Address on page 229

Show Resolved Locations

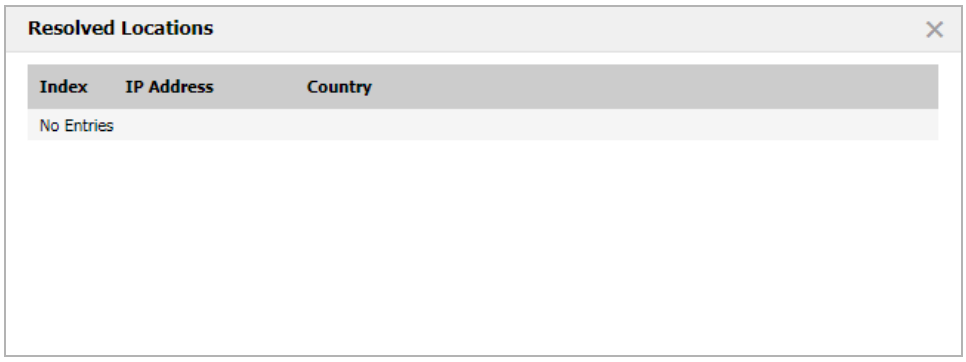

When you click the **SHOW RESOLVED LOCATIONS** button, a pop-up table of resolved IP addresses displays this information:

- **Index**
- **IP Address**
- **Country**

Geo-IP Cache Statistics



The **Geo-IP Cache Statistics** table contains this information:

- **Location Server IP**
- **Resolved Entries**
- **Unresolved Entries**
- **Current Entry Count**
- **Max. Entry Count**
- **Location Map Count**

Custom Countries Statistics



The **Custom Countries Statistics** table contains this information about the number of entries in the list and the number of times lookups have occurred for the entries:

- **No of Entries**
- **No of Times Called**
- **No of Times Not Looked-up**
- **No of Times Resolved**

Check GEO Location Server Lookup

The Geo-IP Filter also provides the ability to look up IP addresses to determine:

- Domain name or IP address
- The country of origin and whether it is classified as a Botnet server

NOTE: The similar Botnet Location Server Lookup tool can also be accessed from the **MANAGE | Security Configuration >** System Services > Botnet Filter page.

The Geo Location and Botnet Server Lookup tool can also be accessed from the **INVESTIGATE | Tools > System Diagnostics** page.

To look up a GEO server:

1. Navigate to **MANAGE | Security Configuration > Security Services > GEO-IP Filter**.
2. Click **Diagnostics**.
3. Scroll to the **Check GEO Location Server Lookup** section.



4. Enter the IP address in the **Lookup IP** field.
5. Click **Go**. Details on the IP address display below the **Result** heading.



Incorrectly Marked Address

If you think an address is marked as part of a country incorrectly, you can report the issue by clicking on the **Geo-IP Status Lookup** link in the **Note** on the **MANAGE | Security Configuration > Security Services > GEO-IP Filter** page.



The link displays the **Submit IP for Geolocation Review** page.

19

Configuring Botnet Filters

NOTE: The Botnet Filtering feature is available on TZ 300 series and above appliances.

- Security Services > Botnet Filter on page 230
- Configuring Botnet Filtering on page 231
- Creating a Custom Botnet List on page 232
- Configuring Dynamic HTTP Authentication on page 236
- Customizing Web Block Page Settings on page 237
- Using Botnet Filter Diagnostics on page 239
- Displaying the Status of the Botnet Feature and Database on page 242

Security Services > Botnet Filter

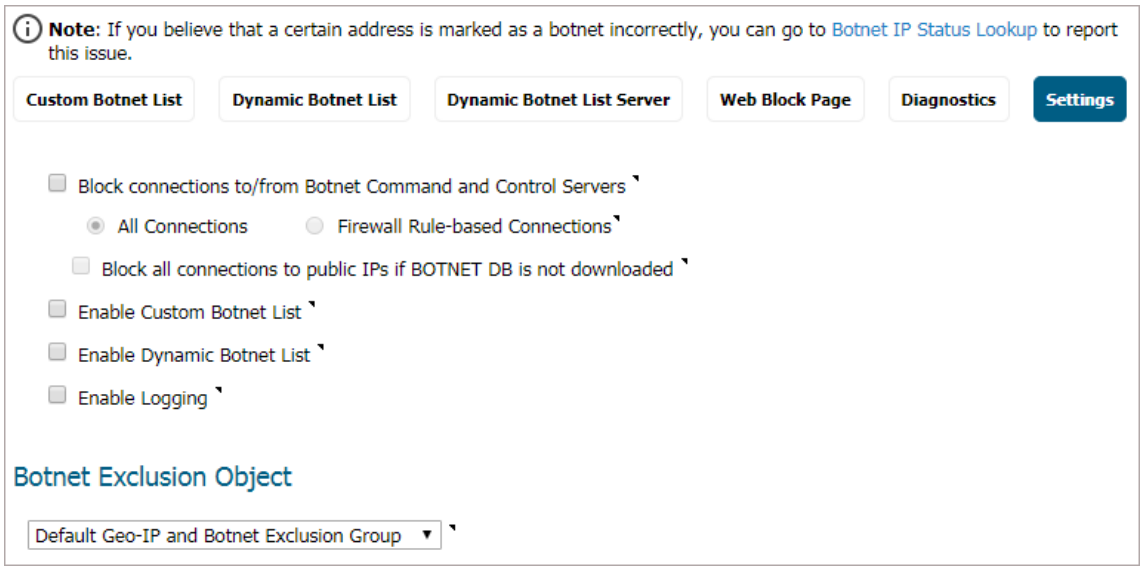

The Botnet Filtering feature allows you to block connections to or from Botnet command and control servers and to make custom Botnet lists.

The Botnet Filtering feature also allows you to create a custom message when you block a web site or to allow dynamic Botnet HTTP authentication.

You can also use the Botnet Filtering Diagnostics tool to show Botnets, monitor Botnet cache statistics, custom Botnet statistics, and look up Botnet servers.

Configuring Botnet Filtering

To configure Botnet filtering:

1 Navigate to the **MANAGE | Security Configuration > Security Services > Botnet Filter**.

2 Click **Settings**.

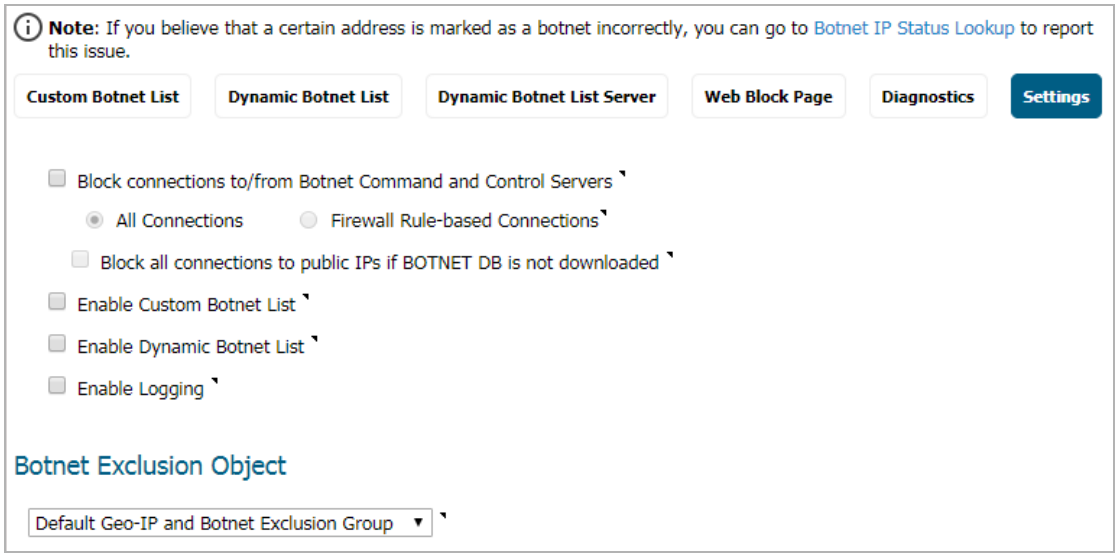

3 To block all servers that are designated as Botnet command and control servers, select the **Block connections to/from Botnet Command and Control Servers** option. All connection attempts to/from Botnet command and control servers will be blocked. This option is not selected by default.

If this option is selected, the radio buttons and the **Block all connections to public IPs if BOTNET DB is not downloaded** option become available.

To exclude selected IPs from this blocking behavior, use exclusion lists as described in the following steps and/or create a custom Botnet list as described in Creating a Custom Botnet List on page 232.

4 If **Block connections to/from Botnet Command and Control Servers** is selected, these options become available:

a Select one of the following two modes for Botnet Filtering:

- **All Connections**: All connections to and from the firewall are filtered. This is the default Botnet block mode.
- **Firewall Rule-Based Connections**: Only connections that match an access rule configured on the firewall are filtered.

b If you want to block all connections to public IPs when the Botnet database is not downloaded, select the **Block all connections to public IPs if BOTNET DB is not downloaded**. This option is not selected by default.

5 To enable the Custom Botnet List, select the **Enable Custom Botnet List** checkbox. This option is not selected by default.

If the **Enable Custom Botnet List** checkbox is not selected, then only the firewall's Botnet database is searched. Go to Step 6.

Enabling a custom list by selecting the **Enable Custom Botnet List** checkbox can affect country identification for an IP address:

a During Botnet identification, the custom Botnet list is searched first.

b If the IP address is not resolved, the firewall's Botnet database is searched.

If an IP address is resolved from the custom Botnet list, it can be identified as either a Botnet IP address or a non-Botnet IP address, and action taken accordingly.

6 Select **Enable logging** to log Botnet Filter-related events.

7 Optionally, you can configure an exclusion list of all IPs belonging to the configured address object/address group. All IPs belonging to the list are excluded from being blocked. To enable an exclusion list, select an address object or address group from the **Botnet Exclusion Object** drop-down menu.



The default exclusion object is Default Geo-IP and Botnet Exclusion Group. You can create your own address object or address group object. as described in *SonicWall SonicOS 6.5 Policies*.

8 Click **ACCEPT**.

Creating a Custom Botnet List

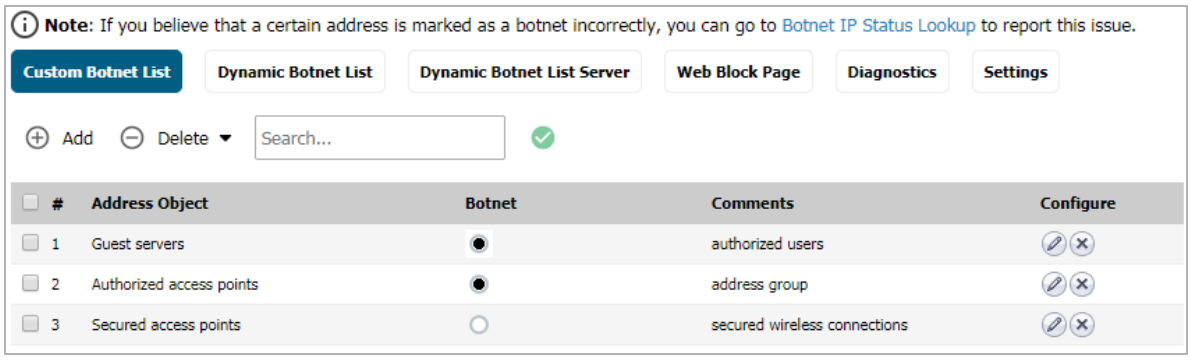

Address Object	Name of the address object or address group object.
Botnet	Icon indicating whether the entry was defined as a Botnet when created. A black circle indicates a Botnet, a white circle a non-Botnet.
Comments	Any comments you added about the entry.
Configure	Contains **Edit** and **Delete** icons for the entry.
Total	Displays the number of entries in the **Custom Botnet List**.

An IP address can be wrongly marked as Botnet. This kind of misclassification can cause incorrect/unwanted filtering of an IP address. Having a custom Botnet list can solve this problem by overriding the Botnet tag for a particular IP address.

Topics:

- Creating a Custom Botnet List on page 233
- Editing a Custom Botnet List Entry on page 234
- Deleting Custom Botnet List Entries on page 235

Creating a Custom Botnet List

IMPORTANT: For the firewall to use the custom Botnet list, you must enable it as described in Configuring Botnet Filtering on page 231.

To create a custom Botnet list:

1 Navigate to the **MANAGE | Security Configuration > Security Services > Botnet Filter**.

2 Click **Settings**.

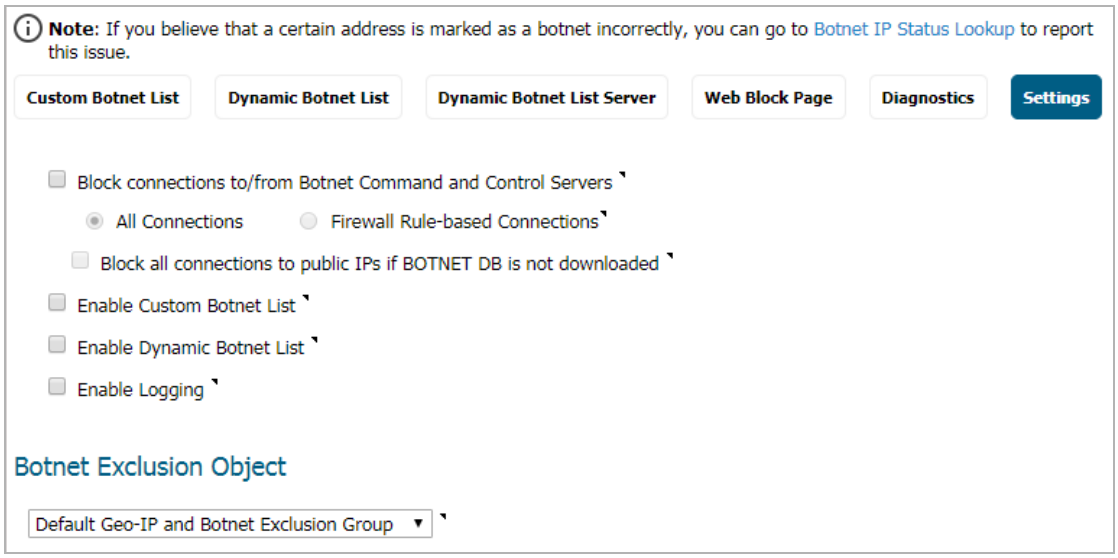

3 Click **Custom Botnet List**.

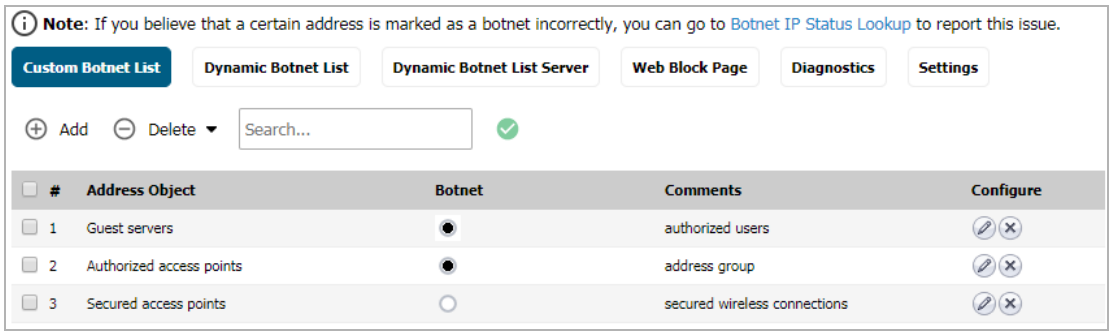

4 Click the **Add** icon. The **Add Custom Botnet List** dialog displays.

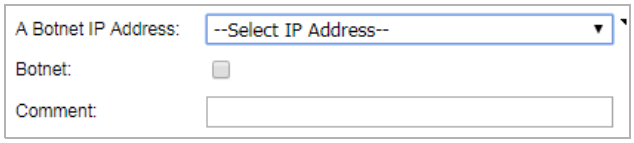

5 Select an IP address object or create a new address object from the **A Botnet IP Address** drop-down menu:

IMPORTANT: An address object cannot overlap any other address objects in the custom country list. Different address objects, however, can have the same country ID.

- **Create new address object...** – the **Add Address Object** dialog displays.

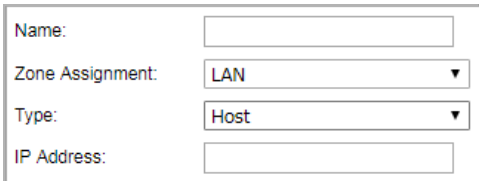

You create a new address object as described in *SonicWall SonicOS 6.5 Policies*, with these restrictions:

- Allowed types are
  - **Host**
  - **Range**
  - **Network**
  - A group of any combination of the first three types

  All other types are disallowed types and cannot be added to the custom Botnet list.

- **Create new address group...** – the **Add Address Object Group** dialog displays.

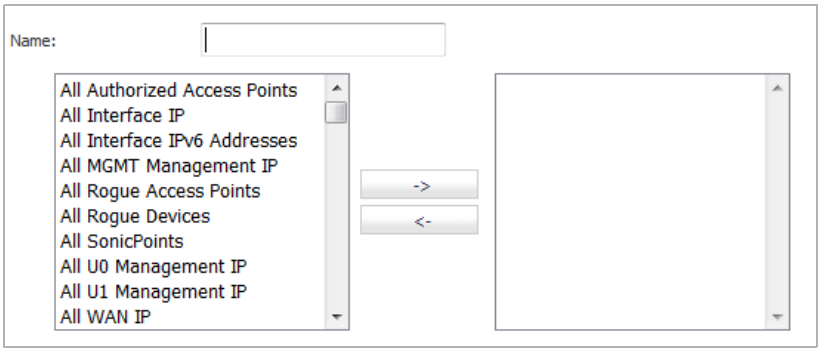

You create a new address object as described in *SonicWall SonicOS 6.5 Policies*

- Already defined address object or address group

6 If this address object is a known Botnet, select a the **Botnet** checkbox.

7 Optionally, add a comment in the **Comment** field.

8 Click **OK**.

Editing a Custom Botnet List Entry

To edit a custom Botnet list entry:

1 In the **Custom Botnet List** table, click the **Edit** icon in the **Configure** column for the entry to be edited. The **Add Custom Botnet List** dialog displays the entry.

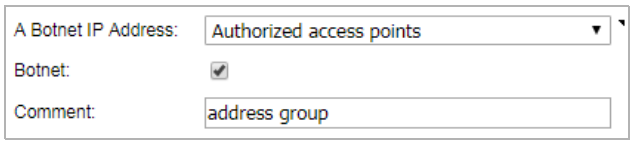

2 Make your changes.

3 Click **OK**. The **Custom Botnet List** table is updated.

Deleting Custom Botnet List Entries

To delete a custom Botnet list entry:

1 Do one of these:

- Click the **Delete** icon in the **Configure** column for the entry.
- Select the checkbox for the entry and then click the **Delete** button.

A confirmation message displays.

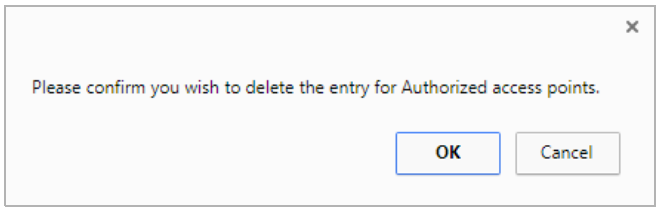

2 Click **OK**.

To delete multiple entries:

1 Select the checkboxes of the entries to be deleted. The **Delete** button becomes available.

2 Click **DELETE**. A confirmation message displays.

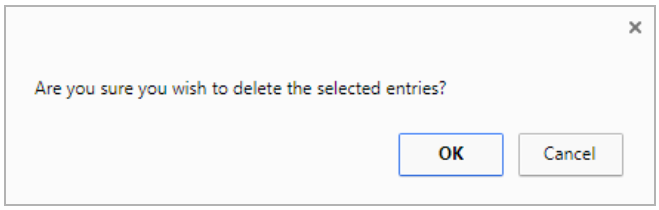

3 Click **OK**.

To delete all entries:

1 Click the checkbox in the table header.

2 Click **Delete**. A confirmation message displays.

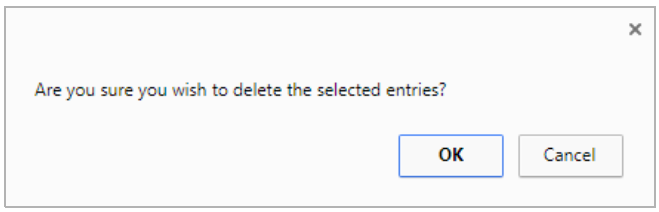

3 Click **OK**.

Configuring Dynamic HTTP Authentication

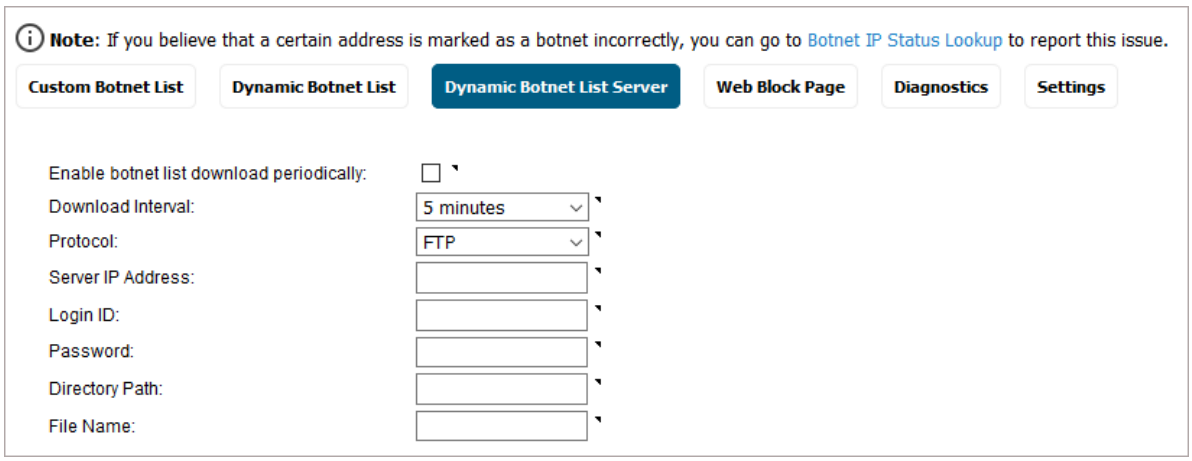

With SonicOS 6.5.2, username and passwords for HTP URLs in the dynamic Botnet configuration are accepted, and the information is transmitted in the HTTP header so the firewall has the required information.

To configure dynamic HTTP authentication:

1 Navigate to **MANAGE | Security Configuration >Security Services > Botnet Filter**.

2 Click **Dynamic Botnet List Server.**

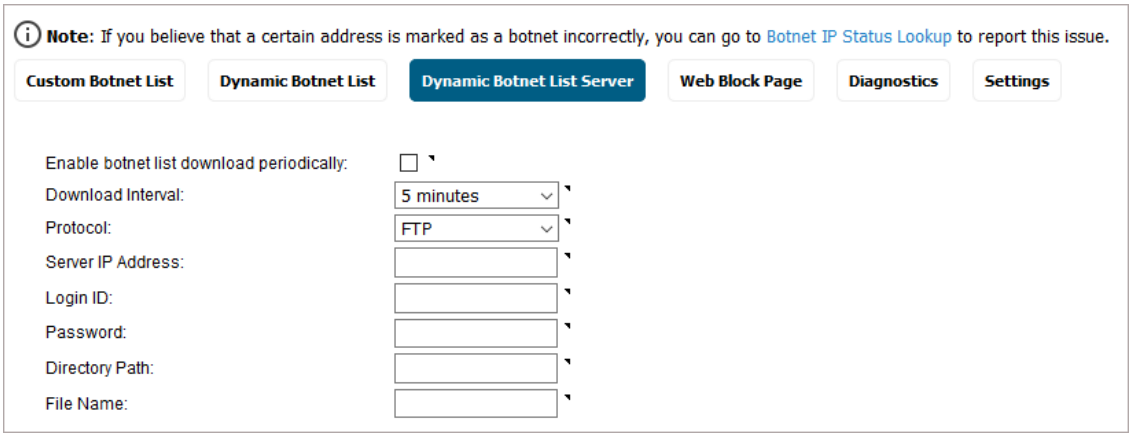

3 Select **Enable botnet list download periodically**. This option is not selected by default.

4 Select the frequency of downloads from **Download Interval**:

- **5 minutes** (default)
- **15 minutes**
- **1 hour**
- **24 hours**

The firewall downloads the Botnet file from the server at the specified interval.

5 Select the protocol in which the firewall has to communicate with the backend server to retrieve the file from **Protocol**:

- **FTP** (default)
- **HTTPS**

6 Enter the IP address of the server to which the Botnet list file will be downloaded in the **Server IP Address** field.

7 Enter the login ID the firewall is to use to connect to the server in the **Login ID** field.

8 Enter the password the firewall is to use to connect to the server in the **Password** field.

9 Enter the directory path the firewall from which the firewall retrieves the Botnet file in the **Directory Path** field. This server directory path is relative to the default root directory.

10 Enter the name of the file on the server to be downloaded in the **File Name** field.

11 Click **ACCEPT**.

Customizing Web Block Page Settings

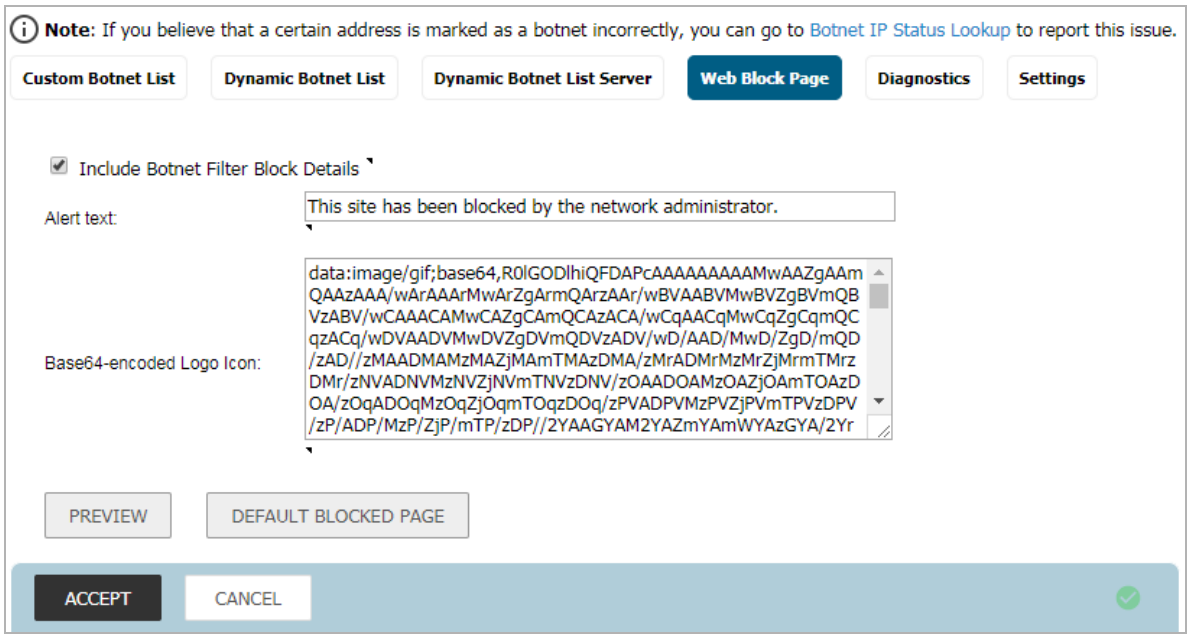
Note: If you believe that a certain address is marked as a botnet incorrectly, you can go to Botnet IP Status Lookup to report this issue.
Custom Botnet List
Dynamic Botnet List
Dynamic Botnet List Server
Web Block Page
Diagnostics
Settings
Include Botnet Filter Block Details
Alert text:
This site has been blocked by the network administrator.
Base64-encoded Logo Icon:

To create a custom message and include a custom logo:

1 Navigate to **MANAGE | Security Configuration > Security Services > Botnet Filter**.

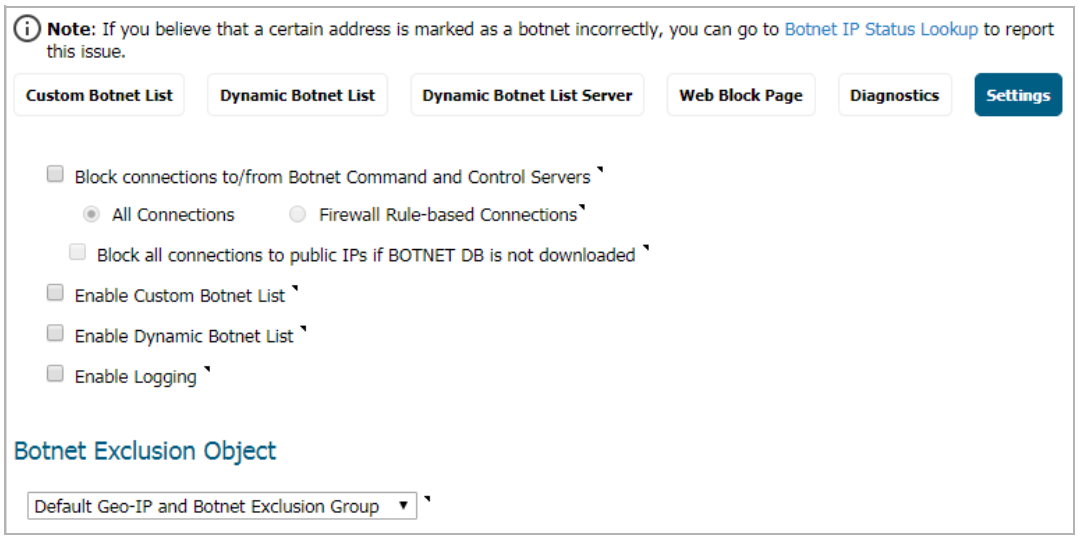

2 Ensure the **Include Botnet Filter Block Details** option is selected. This option is selected by default.

When enabled, this option shows block details such as reason for the block, IP address, and country. When disabled, this option hides all information.

3 Do one of the following:

- To use the default message displayed in the **Alert text** field, `This site has been blocked by the network administrator.`, click the **Default Blocked Page** button and then go to Step 4.
- Specify a custom message to be displayed in the Geo-IP Filter Block page in the **Alert text** field. Your message can be up to 100 characters long.

4 Optionally, in the **Base64-encoded Logo Icon** field, you can specify a Base 64-encoded GIF icon to be displayed as well.

NOTE: Ensure the icon is valid and make the size as small as possible. The recommended size is 400 x 65.

5 To see a preview of your customized message and logo (or the default message and logo), click the **Preview** button. A warning message displays.

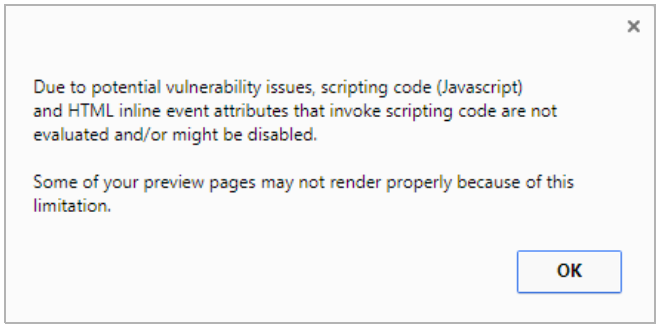

6 Click **OK**. The **Web Site Blocked** message displays.



7 Close the **Web Site Blocked** message.

8 Click **ACCEPT**.

Using Botnet Filter Diagnostics

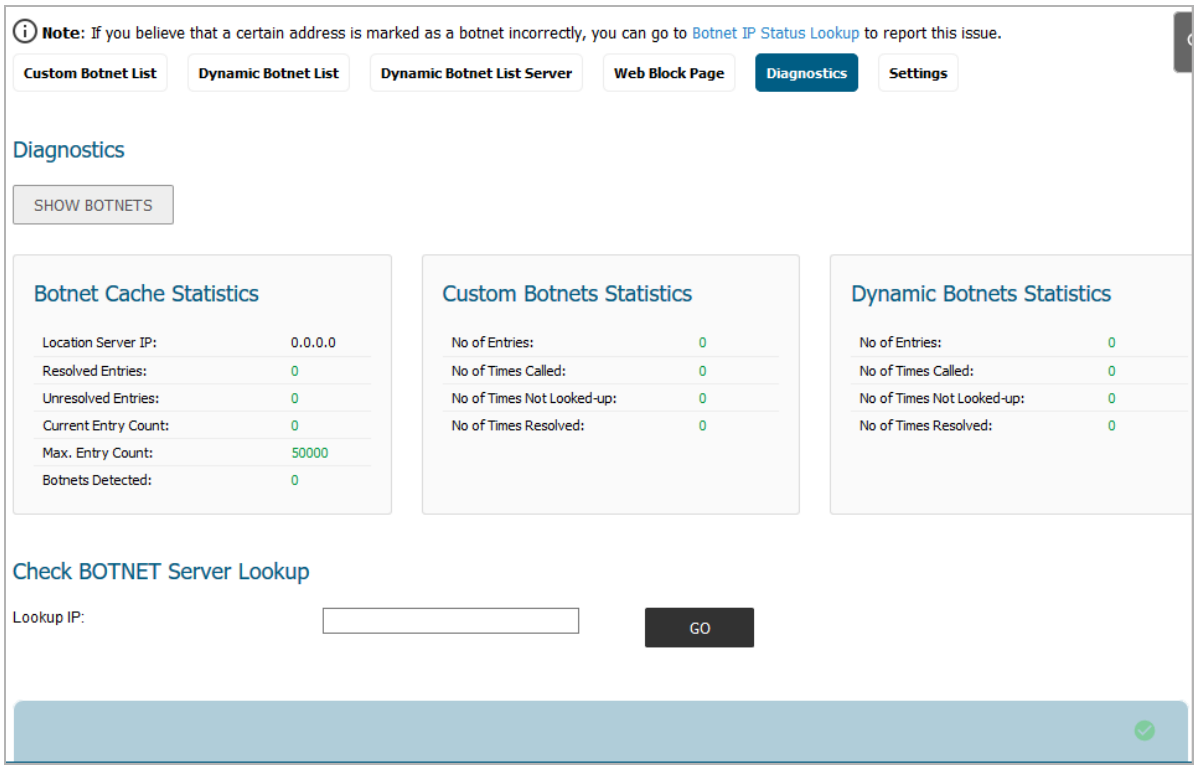

The **MANAGE | Security Configuration > Security Services > Botnet Filter** page has a **Diagnostics** view with several tools:

- Show Resolved Botnet Locations on page 240
- Botnet Cache Statistics on page 240
- Botnets Statistics on page 241
- Check Botnet Server Lookup on page 241
- Incorrectly Marked Address on page 242

Show Resolved Botnet Locations



When you click on **SHOW BOTNETS** in the **Diagnostics** section, a table of resolved IP addresses displays with this information:

- **Index**
- **IP Address** – IP address of the Botnet

Botnet Cache Statistics

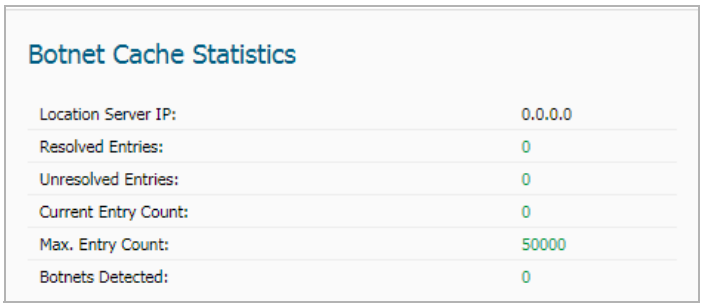

The **Botnet Cache Statistics** table contains this information:

- **Location Server IP**
- **Resolved Entries**
- **Unresolved Entries**
- **Current Entry Count**
- **Max. Entry Count**
- **Botnets Detected**

Botnets Statistics



The **Diagnostics** view displays statistics for both custom and dynamic Botnets. Both the **Custom Botnets Statistics** and **Dynamic Botnet Statistics** tables display the same information about the number of entries in the list and the number of times lookups have occurred for the entries:

- **No of Entries**
- **No of Times Called**
- **No of Times Not Looked-up**
- **No of Times Resolved**

Check Botnet Server Lookup

The Botnet Filter also provides the ability to look up IP addresses to determine:

- Domain name or IP address
- Country of origin and whether the server is classified as a Botnet server

NOTE: The Botnet Server Lookup tool can also be accessed from the **System > Diagnostics** page.

To look up a Botnet server:

1 Navigate to **MANAGE | Security Configuration > Security Services > Botnet Filter**.

2 Click **Diagnostics**.

3 Scroll to the **Check BOTNET Server Lookup** section.



4 Enter the IP address in the **Lookup IP** field,

5 Click **Go**. Details on the IP address are displayed below the **Result** heading.

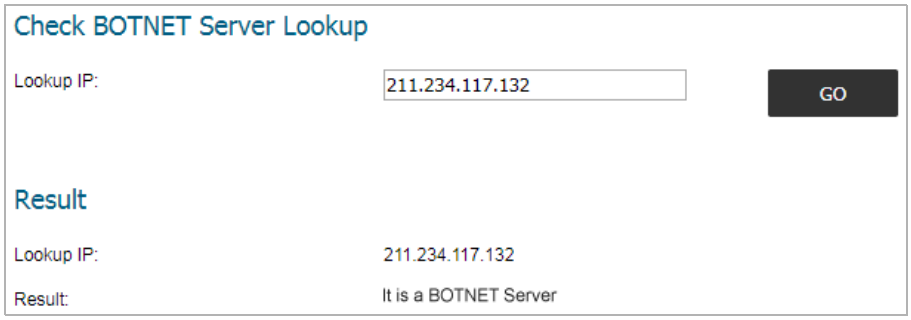

Incorrectly Marked Address

Note: If you believe that a certain address is marked as a botnet incorrectly, you can go to Botnet IP Status Lookup to report this issue.

If you believe that a certain address is marked as a Botnet incorrectly, or if you believe an address should be marked as a Botnet, report this issue at SonicWall Botnet IP Status Lookup by either clicking on the link in the **Note** in the **MANAGE | Security Configuration > Security Services > Botnet Filter** page or going to: SonicWall Botnet IP Status Lookup.

Displaying the Status of the Botnet Feature and Database

To display the status of the Botnet feature and database, click the Status icon. A popup with the status displays.

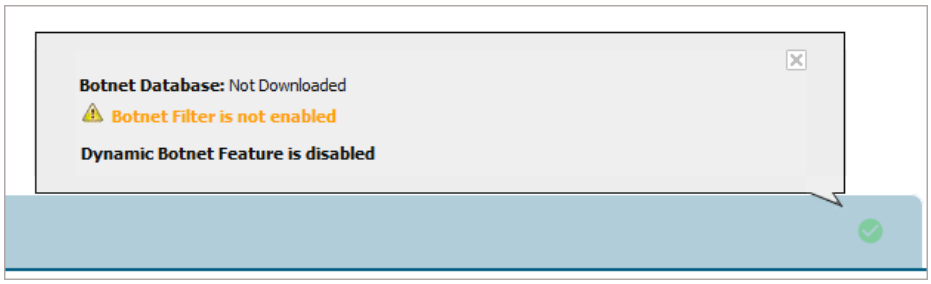

To close the popup, click the **X**.

Part 3

Security Config | Decryption Services

- About DPI-SSL
- Configuring the DPI-SSL/TLS Client
- Configuring DPI-SSL/TLS Server Settings
- Configuring DPI-SSH

20

About DPI-SSL

NOTE: DPI-SSL is a separate, licensed feature that provides inspection of encrypted HTTPS traffic and other SSL-based IPv4 and IPv6 traffic.

- About DPI-SSL on page 244
- Deployment Scenarios on page 247
- Customizing DPI-SSL on page 247
- Connections per Appliance Model on page 248

About DPI-SSL

Topics:

- Supported Features on page 244
- Security Services on page 246

Supported Features

Deep Packet Inspection of Secure Socket Layer (DPI-SSL) extends SonicWall's Deep Packet Inspection technology to the inspection of encrypted HTTPS traffic and other SSL-based traffic. The SSL traffic is decrypted (intercepted) transparently, scanned for threats, and then re-encrypted and, if no threats or vulnerabilities are found, sent along to its destination.

DPI-SSL provides additional security, application control, and data-leakage prevention for analyzing encrypted HTTPS and other SSL-based traffic. DPI-SSL supports:

- Transport Layer Security (TLS) Handshake Protocol 1.2 and earlier versions – Starting with SonicOS 6.2.5.1, the TLS 1.2 communication protocol is supported during SSL inspection/decryption between the firewall and the server in DPI-SSL deployments (previously, TLS 1.2 was only supported between client and firewall). SonicOS also supports TLS 1.2 in other areas as well.
- SHA-256 – All re-signed server certificates are signed with the SHA-256 hash algorithm.
- Perfect Forward Secrecy (PFS) – Perfect Forward Secrecy-based ciphers and other stronger ciphers are prioritized over weak ciphers in the advertised cipher suite. As a result, the client or server is not expected to negotiate a weak cipher unless the client or server does not support a strong cipher.

DPI-SSL also supports application-level Bandwidth Management over SSL tunnels. App Rules HTTP bandwidth management policies also applies to content that is accessed over HTTPS when DPI-SSL is enabled for App Rules.

DPI-SSL for both client and server can be controlled by Access Rules.

Topics:

- Support for Local CRL on page 245
- TLS Certificate Status Request Extension on page 245
- Blocking of SSH X11 Forwarding on page 245
- Support for ECDSA-Related Ciphers on page 246
- DPI-SSL and CFS HTTPS Content Filtering Work Independently on page 246
- Original Port Numbers Retained in Decrypted Packets on page 246

Support for Local CRL

A Certificate Revocation List (CRL) 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. A problem with contacting the CA for this list is that the browser cannot confirm whether it has reached the CA's servers or if an attacker has intercepted the connection to bypass the revocation check.

Local CRL is relative to typical CRL (or online CRL). For typical CRL, the client needs to download the CLR from a CRL distribution point. If the client is unable to download the CRL, then by default, the client trusts the certificate. Contrary to typical CRL, Local CRL maintains a list of revoked certificates locally in import memory for DPI-SSL to verify whether the certificate has been revoked.

For further information about this feature, contact Technical Support.

TLS Certificate Status Request Extension

DPI-SSL now supports the new TLS Certificate Status Request extension (formally known as OCSP stapling). By supporting this extension, the certificate status information is delivered to the DPI-SSL client through an already established channel, thereby reducing overhead and improving performance.

For further information about this feature, see *SonicWall SonicOS 6.5 System Setup* or contact Technical Support.

Blocking of SSH X11 Forwarding

NOTE: X11 Forwarding requires a valid SonicWall DPI-SSH license.

X is a popular window system for Unix workstations. Using X, a user can run remote X applications that open their windows on the user's local display (and vice versa, running local applications on remote displays). If the remote server is outside after a firewall and administrator have blocked remote connections, user can still use SSH tunneling to get the X display on a local machine. A user can thus circumvent the application-based security policies on the firewall, thereby creating security risks. As X protocol sessions between applications and X servers are not encrypted while being transmitted over a network, an X11 protocol connection can be routed through an SSH connection to provide security and stronger authentication. This feature is called X11 forwarding An SSH client requests X forwarding when it connects to an SSH server (assuming X forwarding is enabled in the client). If the server allows X forwarding for this connection, login proceeds normally, but the server takes some special steps behind the scenes. In addition to handling the terminal session, the server sets itself up as a proxy X server running on the remote machine and sets the DISPLAY environment variable in the remote shell to point to the proxy X display. If an X client program is run, it connects to the proxy. The proxy behaves just like a real X server, and in turn instructs the SSH client to behave as a proxy X client, connecting to the X server on the local machine. The SSH client and server then cooperate to pass X protocol information back and forth over the SSH pipe between the two X sessions, and the X client program appears on your screen just as if it had connected directly to your display. DPI-SSH X11 forwarding supports these clients:

- SSH client for Cygwin
- Putty •secureCRT
- SSH on Ubutu
- SSH on centos

DPI-SSH X11 Forwarding supports the SSH servers on:

- Fedora
- Ubuntu

SSH X11 Forwarding supports both route mode and wire mode. For:

- Wire mode, SSH X11 Forwarding is only supported in the secure (active DPI of inline traffic) mode.
- Route mode, here is no limitation.

The maximum number of connections supported for SSH X11 Forwarding is same as for DPI-SSH: 1000.DPI-SSH.

Support for ECDSA-Related Ciphers

DPI-SSL Client supports ECDSA (Elliptic Curve Digital Signature Algorithm) ciphers:

- TLS_ECDHE_ECDSA_WIATH_AES_128_GCM_SHA256
- TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256

DPI-SSL and CFS HTTPS Content Filtering Work Independently

DPI-SSL and CFS HTTPS content filtering can be enabled at the same time and function as follows:

- If DPI-SSL Client Inspection is disabled, Content Filter Service filters HTTPS connections.
- If DPI-SSL Client Inspection is enabled, but the Content Filter option is not selected, Content Filter Service filters HTTPS connections.
- If DPI-SSL Client Inspection is enabled and the Content Filter option is selected, CFS does not filter HTTPS connections.

Original Port Numbers Retained in Decrypted Packets

For encrypted connections DPI-SSL/DPI-SSH connections, the decrypted packet shows the destination port as 80 (in the case of HTTPS). When the decrypted packets are observed in packet capture/Wireshark, they now retain the original port numbers. The port number change applies only to the packet capture and not to the actual packet or connection cache.

Security Services

The following security services and features can use DPI-SSL:

Gateway Anti-Virus	Content Filtering
Gateway Anti-Spyware	Application Firewall
Intrusion Prevention	

Deployment Scenarios

DPI-SSL has two main deployment scenarios:

- **Client DPI-SSL**: Used to inspect HTTPS traffic when clients on the appliance's LAN access content located on the WAN. Exclusions to DPI-SSL can be made on a common-name or category basis.
- **Server DPI-SSL**: Used to inspect HTTPS traffic when remote clients connect over the WAN to access content located on the appliance's LAN.

Proxy Deployment

DPI-SSL supports proxy deployment, where all client browsers are configured to redirect to a proxy server, but an appliance sits between the client browsers and the proxy server. All DPI-SSL features are supported in this scenario, including supporting domain exclusions when the domain is part of a virtual hosting server, or in some cloud deployments, wherein the same server IP can be used by multiple domains.

Additionally, typical data center server farms are fronted with a load balancer and/or reverse SSL Proxy to offload SSL processing on the servers. For a load balancer fronting the servers and doing decryption, the appliance usually only sees the IP of the load balancer, and the load balancer decrypts the content and determines the specific server to assign this connection to. DPI-SSL now has a global policy option to disable an IP-based exclusion cache. The exclusions continues to work even if the IP-based exclusion cache is off.

Customizing DPI-SSL

IMPORTANT: Add the NetExtender SSL VPN gateway to the DPI SSL IP-address exclusion list. As NetExtender traffic is PPP-encapsulated, having SSL VPN decrypt such traffic does not produce meaningful results.

In general, the policy of DPI-SSL is to secure any and all traffic that flows through the appliance. This may or may not meet your security needs, so DPI-SSL allows you to customize what is processed.

DPI-SSL comes with a list (database) of built-in (default) domains excluded from DPI processing. You can add to this list at any time, remove any entries you've added, and/or toggle built-in entries between exclusion from and inclusion in DPI processing. DPI-SSL also allows you to exclude or include domains by common name or category (for example, banking or health care).

Excluded sites, whether by common name or category, however, can become a security risk that can be exploited in the future by exploit kits that circumvent the appliance and are downloaded to client machines or by a man-in-the-middle hijacker presenting a fake server site/certificate to an unsuspecting client. To prevent such risks, DPI-SSL allows excluded sites to be authenticated before exclusion.

As the percentage of HTTPS connections increase in your network and new https sites appear, it is improbable for even the latest SonicOS version to contain a complete list of built-in/default exclusions. Some HTTPS connections fail when DPI-SSL interception occurs due to the inherent implementation of a new client app or the server implementation, and these sites may need to be excluded on the appliance to provide a seamless user experience. SonicOS keeps a log of these failed connections that you can troubleshoot and use to add any trusted entries to the exclusion list.

In addition to excluding/including sites, DPI-SSL provides both global authentication policy and a granular exception policy to the global one. For example, with a global policy to authenticate connection, some connections may be blocked that are in essence safe, such as new trusted CA certificates or a a self-signed server certificate of a private (or local-to-enterprise deployment) secure cloud solution. The granular option allows you to exclude individual domains from the global authentication policy.

You can configure exclusions for a domain that is part of a list of domains supported by the same server (certificate). That is, some server certificates contain multiple domain names, but you want to exclude just one of these domains without having to exclude all of the domains served by a single server certificate. For example, you can exclude `youtube.com` without having to exclude any other domain, such as `google.com`, even though `*.google.com` is the common name of the server certificate that has `youtube.com` listed as an alternate domain under Subject Alternate-Name extension.

Connections per Appliance Model

To learn about the hardware model and its maximum concurrent connections to perform the Client DPI-SSL inspections, refer to the following platform datasheets:

- SonicWall Network Security appliance (NS*a*) series
- SonicWall TZ Series
- SonicWall SuperMassive Series

Refer to the SonicWall resources page for more information about our Product Series. Search for high-end, mid-range, entry level, and virtual firewall details, such as Maximum connections (DPI SSL), from the **By Product Series** drop-down menu.

NOTE: For NS*a* Series; SuperMassive 9200, 9400, and 9600; and NSA 3600 Series (and higher) firewalls with more that 250,000 DPI settings and dynamic connection sizing configured, the firewall can increase the DPI-SSL connection count dynamically. For more information, see Dynamic Connection Sizing on page 18.

21

Configuring the DPI-SSL/TLS Client

- Decryption Services > DPI-SSL/TLS Client on page 249
- Viewing DPI-SSL Status on page 250
- Configuring the DPI-SSL/TLS Client on page 250

Decryption Services > DPI-SSL/TLS Client

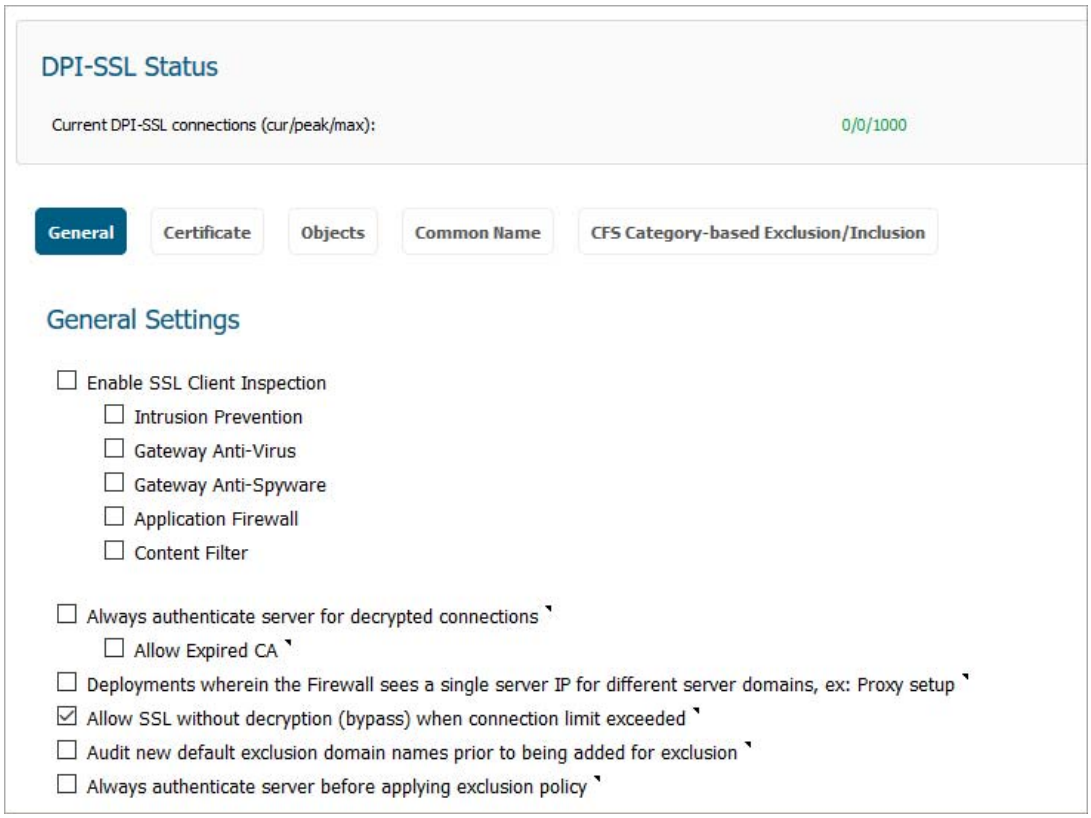

TIP: For information about DPI-SSL, see About DPI-SSL on page 244.

Viewing DPI-SSL Status



The **DPI-SSL Status** section displays the current DPI-SSL connections, peak connections, and maximum connections.

Configuring the DPI-SSL/TLS Client

The DPI-SSL/TLS Client deployment scenario typically is used to inspect HTTPS traffic when clients on the LAN browse content located on the WAN. In this scenario, the firewall typically does not own the certificates and private keys for the content it is inspecting. After performing DPI-SSL inspection, the appliance re-writes the certificate sent by the remote server and signs this newly generated certificate with the certificate specified in the Client DPI-SSL configuration. By default, this is the firewall certificate authority (CA) certificate, but a different certificate can be specified. Users should be instructed to add the certificate to their browser's trusted list to avoid certificate trust errors.

Topics:

- Configuring General Settings on page 250
- Selecting the Re-Signing Certificate Authority on page 253
- Configuring Exclusions and Inclusions on page 254
- Excluding/Including by Common Name on page 256
- Client DPI-SSL Examples on page 264

Configuring General Settings

Topics:

- Enabling SSL Client Inspection on page 250
- Enabling DPI-SSL Client on a Zone on page 253
- Enabling DPI-SSL Server on a Zone on page 253

Enabling SSL Client Inspection

To enable SSL Client inspection:

1 Navigate to **MANAGE | Security Configuration > Decryption Services > DPI-SSL/TLS Client**.

2 Click **General**.

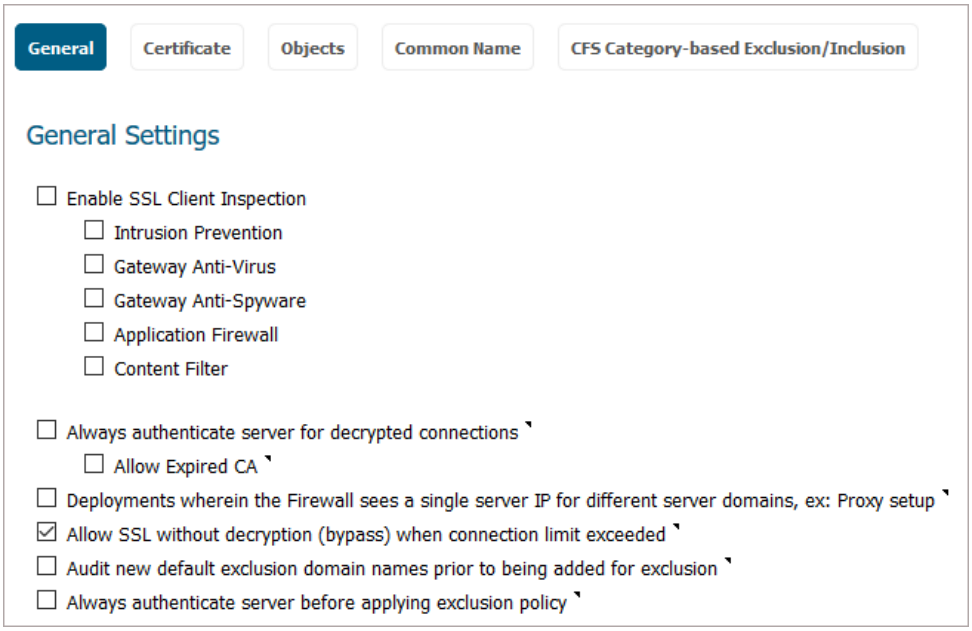

3 Select **Enable SSL Client Inspection**. This option is not selected by default..

4 Select one or more services with which to perform inspection; none are selected by default:

- **Intrusion Prevention**
- **Gateway Anti-Virus**
- **Gateway Anti-Spyware**
- **Application Firewall**
- **Content Filter**

5 To authenticate servers for decrypted/intercepted connections, select **Always authenticate server for decrypted connections**. When enabled, DPI-SSL blocks connections:

- To sites with untrusted certificates.
- If the domain name in the Client Hello cannot be validated against the Server Certificate for the connection.

This option is not selected by default. When this option is selected, **Allow Expired CA** becomes available.

IMPORTANT: Only enable this option if you need a high level of security. Blocked connections show up in the connection failures list, as described in Showing Connection Failures on page 260.

TIP: If you enable this option, use the **Skip CFS Category-based Exclusion** option (see Excluding/Including Common Names on page 257) to exclude a particular domain or domains from this global authenticate option. This is useful to override any server authentication-related failures of trusted sites.

6 To allow expired or intermediate CAs, select **Allow Expired CS**. This option is not selected by default. If it is not selected, connections are blocked if the domain name in the Client Hello cannot be validated against the server certificate for the connections.

7 To disable use of the server IP address-based dynamic cache for exclusion, select **Deployments wherein the Firewall sees a single server IP for different server domains, ex: Proxy setup**. This option is not selected by default.

This option is useful for proxy deployments, where all client browsers redirect to a proxy server, including if appliance is between the client browsers and the proxy server. All DPI-SSL features are supported, including domain exclusions when the domain is part of a virtual hosting server, as part of a server farm fronted with a load balancer, or in some cloud deployments, wherein the same server IP can be used by multiple domains.

In such deployments, all server IPs as seen by the appliance are the proxy server's IP. It is, therefore, imperative that in proxy deployments, IP-based exclusion cache is disabled. Enabling this option does not affect SonicOS's capability to perform exclusions.

8 By default, new connections over the DPI-SSL connection limit are bypassed. To allow new connections to bypass decryption instead of being dropped when the connection limit is exceeded, select the **Allow SSL without decryption (bypass) when connection limit exceeded** checkbox. This option is selected by default.

To ensure new connections over the DPI-SSL connection limit are dropped, deselect/disable this checkbox.

9 To audit new, built-in exclusion domain names before they are added for exclusion, select the **Audit new built-in exclusion domain names prior to being added for exclusion** checkbox. By default, this checkbox is not enabled.

When this option is enabled, whenever changes to the built-in exclusion list occur, for example, an upgrade to a new firmware image or other system-related actions, a notification pop-up dialog displays over the **Decryption Services > DPI-SSL/TLS Client** page with the changes. You can inspect/audit the new changes and accept or reject any, some, or all of the new changes to the built-in exclusion list. At this point, the run-time exclusion list is updated to reflect the new changes.

If this option is disabled, SonicOS accepts all new changes to the built-in exclusion list and adds them automatically.

10 To always authenticate a server before applying a common-name or category exclusion policy, select the **Always authenticate server before applying exclusion policy** checkbox. This option is not selected by default. When enabled, DPI-SSL blocks excluded connections:

- To sites with untrusted certificates.
- If the domain name in the Client Hello cannot be validated against the Server Certificate for the connection.

This is a useful feature to authenticate the server connection before applying exclusion policies. Enabling this option ensures that the appliance does not blindly apply exclusion on connections and thereby create a security hole for exclusion sites or sites belonging to excluded categories. This is especially relevant when banking sites, as a category, are excluded.

By validating both the server certificate and the domain name in the Client Hello before applying an exclusion policy, SonicOS can reject untrusted sites and potentially block a type of zero-day attack from taking place. The SonicOS implementation takes the "trust-but-verify" approach to ensure that a domain name that matches the exclusion policy criteria is validated first, thus preventing an unsuspecting client from phishing or URL-redirect-related attacks.

IMPORTANT: If you are excluding alternate domains in the Subject-Alternate-Name extension, it is recommended that you enable this option.

TIP: If you enable this option, use the **Skip CFS Category-based Exclusion** option (see Excluding/Including Common Names on page 257) to exclude a particular domain or domains from this global authenticate option. This is useful to override any server authentication-related failures of trusted sites.

11 Click **ACCEPT**.

Enabling DPI-SSL Client on a Zone

To enable DPI-SSL Client on a zone:

12 Navigate to **MANAGE | System Setup > Network > Zones**.

> **TIP:** For information about configuring zones, see *SonicWall SonicOS 6.5 System Setup*.

13 Click the **Edit** icon for the zone to be configured. The **Edit Zone** dialog displays.

14 Select **Enable SSL Client Inspection**. This option is not selected by default.

15 Finish configuring the zone.

16 Click **OK**.

17 Repeat Step 13 through Step 16 for each zone on which to enable DPI-SSL client inspection

Enabling DPI-SSL Server on a Zone

To enable DPI-SSL Server on a zone:

1 Navigate to **MANAGE | Security Configuration > Decryption Services > DPI-SSL/TLS Client**.

> **TIP:** For information about configuring DPI-SSL servers, see Configuring DPI-SSL/TLS Server Settings on page 267.

2 Select **Enable SSL Server Inspection**. This option is not selected by default.

3 Select one or more types of inspection.

4 Click **ACCEPT**.

5 Navigate to **MANAGE | System Setup > Network > Zones**.

> **TIP:** For information about configuring zones, see *SonicWall SonicOS 6.5 System Setup*.

6 Click the **Edit** icon for the zone to be configured. The **Edit Zone** dialog displays.

7 Select **Enable SSL Server Inspection**. This option is not selected by default.

8 Finish configuring the zone.

9 Click **OK**.

10 Repeat Step 6 through Step 8 for each zone on which to enable DPI-SSL server inspection

Selecting the Re-Signing Certificate Authority

The re-signing certificate replaces the original certificate signing authority only if that authority certificate is trusted by the firewall. If the authority is not trusted, then the certificate is self-signed. To avoid certificate errors, choose a certificate that is trusted by devices protected by DPI-SSL.

> **NOTE:** For information about requesting/creating a DPI SSL Certificate Authority (CA) certificate, see the Knowledge Base article, *How to request/create DPI-SSL Certificate Authority (CA) certificates for the purpose of DPI-SSL certificate resigning* (SW14090).

To select a re-signing certificate

1 Navigate to the **MANAGE | Security Configuration > Decryption Services > DPI-SSL/TLS Client** page.

2 Click **Certificate**.

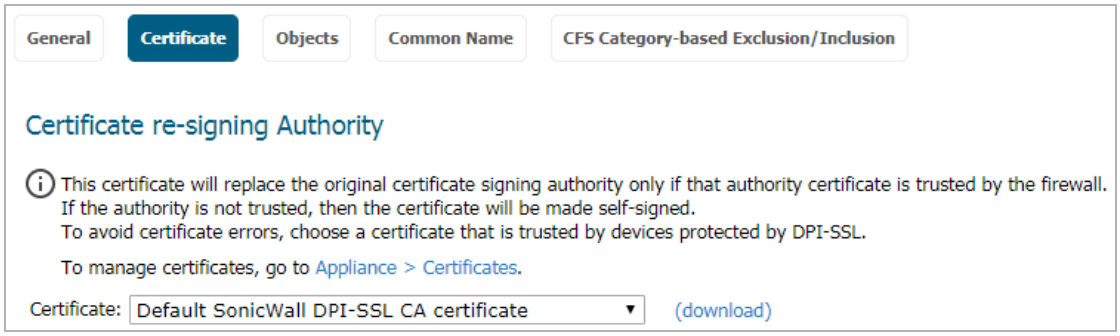

3 Select the certificate to use from the **Certificate** drop-down menu. By default, DPI-SSL uses the **Default SonicWall DPI-SSL CA certificate** to re-sign traffic that has been inspected.

NOTE: If the certificate you want is not listed, you can import it from the **MANAGE | System Setup > Appliance > Certificates** page. See *SonicWall SonicOS 6.5 System Setup*.

For PKCS-12-formatted certificates, see *SonicWall SonicOS 6.5 System Setup*.

4 To download the selected certificate to the firewall, click the **(download)** link. The **Opening *filename*** dialog appears.

TIP: To view available certificates, click on the **(Manage Certificates)** link to display the **MANAGE | System Setup > Appliance > Certificates** page

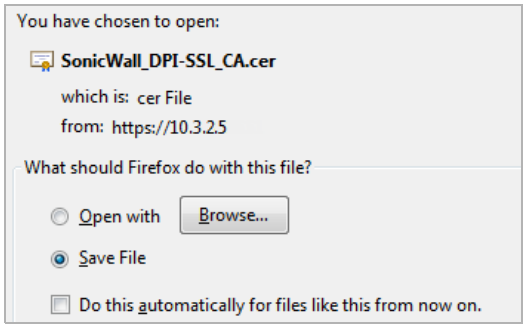

a Ensure the **Save File** radio button is selected.

b Click **OK**.

The file is downloaded.

5 Click **ACCEPT**.

Adding Trust to the Browser

For a re-signing certificate authority to successfully re-sign certificates, browsers have to trust the certificate authority. Such trust can be established by having the re-signing certificate imported into the browser's trusted CA list. Follow your browser's instructions for importing re-signing certificates.

Configuring Exclusions and Inclusions

By default, when DPI-SSL is enabled, it applies to all traffic on the appliance. You can customize to which traffic DPI-SSL inspection applies:

- **Exclusion/Inclusion** lists exclude/include specified objects and groups

- **Common Name** exclusions excludes specified host names
- **CFS Category-based Exclusion/Inclusion** excludes or includes specified categories based on CFS categories

This customization allows individual exclusion/inclusion of alternate names for a domain that is part of a list of domains supported by the same server (certificate). In deployments that process a large amount of traffic, to reduce the CPU impact of DPI-SSL and to prevent the appliance from reaching the maximum number of concurrent DPI-SSL inspected connections, it can be useful to exclude trusted sources.

NOTE: If DPI-SSL is enabled on the firewall when using Google Drive, Apple iTunes, or any other application with pinned certificates, the application may fail to connect to the server. To allow the application to connect, exclude the associated domains from DPI-SSL; for example, to allow Google Drive to work, exclude:

- `.google.com`
- `.googleapis.com`
- `.gstatic.com`

As Google uses one certificate for all its applications, excluding these domains allows Google applications to bypass DPI-SSL.

Alternatively, exclude the client machines from DPI-SSL.

Topics:

- Excluding/Including Objects/Groups on page 255
- Excluding/Including by Common Name on page 256
- Specifying CFS Category-based Exclusions/Inclusions on page 262
- Content Filtering on page 264
- App Rules on page 266

Excluding/Including Objects/Groups

To customize DPI-SSL client inspection:

1. Navigate to the **MANAGE | Security Configuration > Decryption Services > DPI-SSL/TLS Client** page.
2. Click **Objects**.

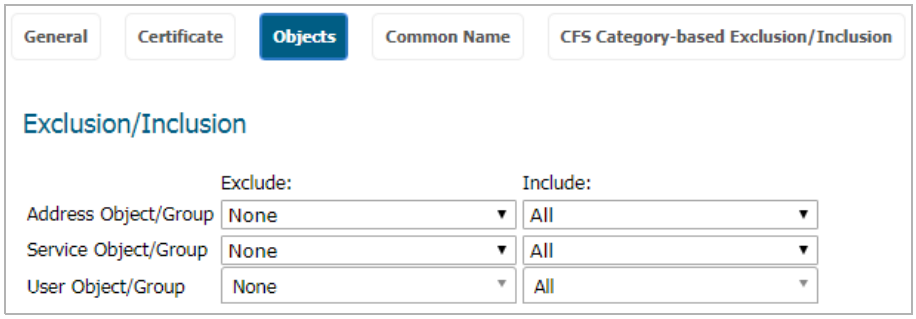

3. From the **Address Object/Group Exclude** and **Include** drop-down menus, select an address object or group to exclude or include from DPI-SSL inspection. By default, **Exclude** is set to **None** and **Include** is set to **All**.

TIP: The **Include** drop-down menu can be used to fine tune the specified exclusion list. For example, by selecting the **Remote-office-California** address object in the **Exclude** drop-down menu and the **Remote-office-Oakland** address object in the **Include** drop-down menu.

4 From the **Service Object/Group Exclude** and **Include** drop-down menus, select an address object or group to exclude or include from DPI-SSL inspection. By default, **Exclude** is set to **None** and **Include** is set to **All**.

5 From the **User Object/Group Exclude** and **Include** drop-down menus, select an address object or group to exclude or include from DPI-SSL inspection. By default, **Exclude** is set to **None** and **Include** is set to **All**.

6 Click **ACCEPT**.

Excluding/Including by Common Name

You can add trusted domain names to the exclusion list. Adding trusted domains to the Built-in exclusion database reduces the CPU effect of DPI-SSL and prevents he appliance from reaching the maximum number of concurrent DPI-SSL inspected connections.

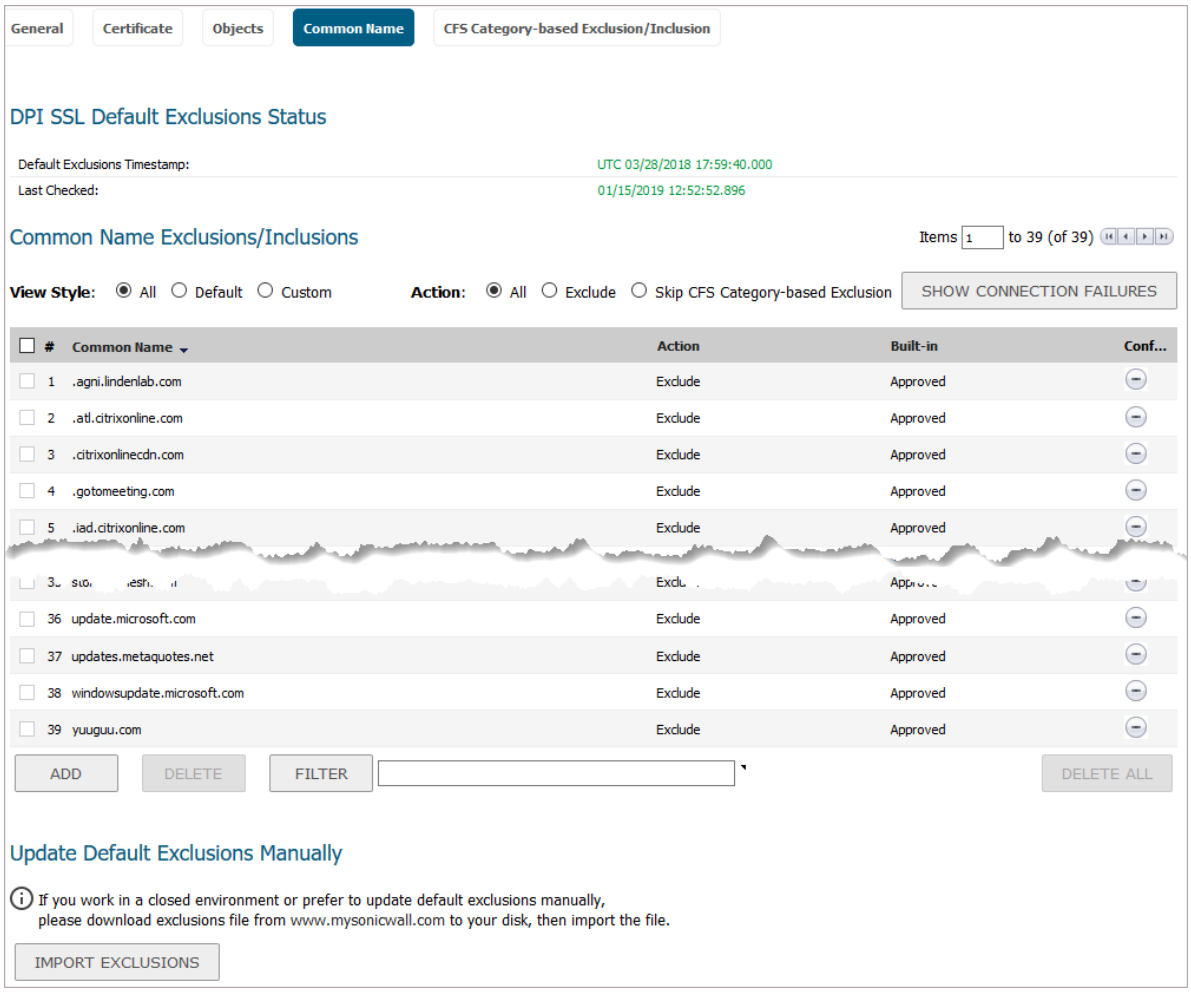

Topics:

- Viewing Status of DPI SSL Default Exclusions on page 257
- Excluding/Including Common Names on page 257
- Deleting Custom Common Names on page 260
- Showing Connection Failures on page 260
- Updating Default Exclusions Manually on page 261

Viewing Status of DPI SSL Default Exclusions

The firewall periodically checks for updates to the DPI SSL default exclusions database on MySonicWall and displays the latest status of the database in the **DPI SSL Default Exclusions Status** section. You can update the database on the firewall manually, as described in Updating Default Exclusions Manually on page 261.

To view the status of default exclusions:

1 Navigate to **MANAGE | Security Configuration > Decryption Services > DPI-SSL/TLS Client**.

2 Scroll to **DPI SSL Default Exclusions Status**.



Default Exclusions Timestamp	Date and time the default exclusions database was updated.
Last Checked	Date and time the firewall checked the default exclusions database.

Excluding/Including Common Names

To exclude/include entities by common name:

1 Navigate to the **MANAGE | Security Configuration > Decryption Services > DPI-SSL/TLS Client** page.

2 Click **Common Name**.

3 Scroll to **Common Name: Exclusions/Inclusions**.

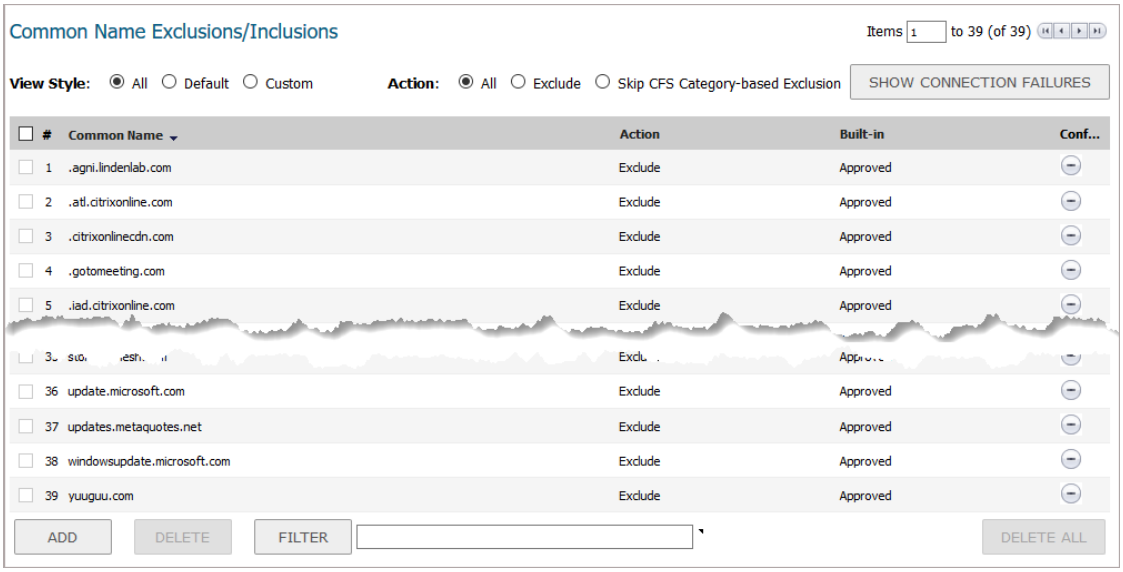

4 You can control the display of the common names by selecting the following options:

- **View Style** options:
  - **All** (default) – Displays all common names.
  - **Built-in** – Displays only non-custom common names.
  - **Custom** – Displays only common names you've added.

- **Action** options:
  - **All** (default) – Displays both excluded and CFS Category-exclusion overrides.
  - **Exclude** – Displays only excluded common names.
  - **Skip CFS Category-based Exclusion** – Displays only custom common names that have the override CFS category-based exclusion option selected.

NOTE: Use the **Skip CFS Category-based Exclusion** option to exclude a particular domain from the global inclusion options, **Always authenticate server for decrypted connections** and **Always authenticate server before applying exclusion policy**.

5 By default, all Built-in common names are approved. You can reject the approval of a Built-in common name by:

a Clicking on the **Reject** icon in the **Configure** column for the common name. A confirmation message displays.

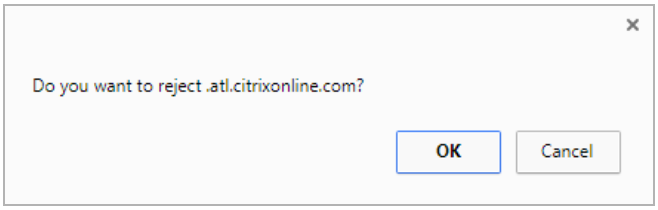

b Click **OK**.

The **Reject** icon becomes an **Accept** icon, and **Approved** in the **Built-in** column become **Rejected**.

TIP: Built-in common names cannot be modified or deleted, but you can reject or accept them.

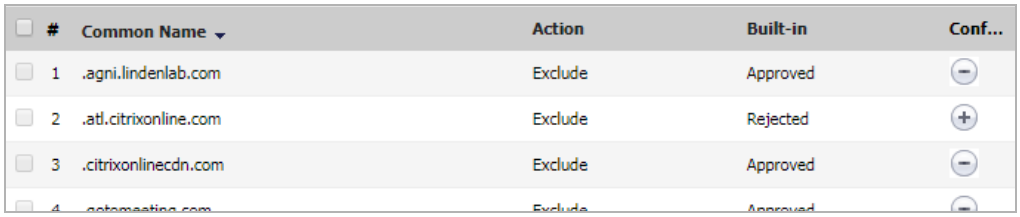

#	Common Name	Action	Built-in	Conf...
1	.agni.lindenlab.com	Exclude	Approved	
2	.atl.citrixonline.com	Exclude	Rejected	
3	.citrixonlinecdn.com	Exclude	Approved	

To accept a rejected Built-in common name:

a Click its **Accept** icon. A confirmation message displays.

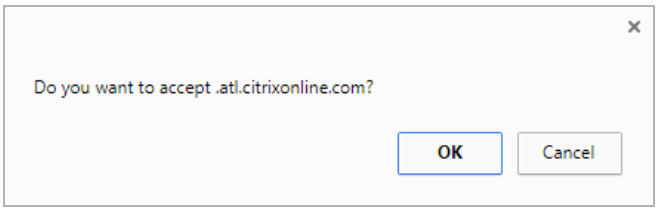

b Click **OK**.

6 To add a custom common name, click the **ADD** button below the **Common Name Exclusions/Inclusions** table. The **Add Common Names** dialog displays.

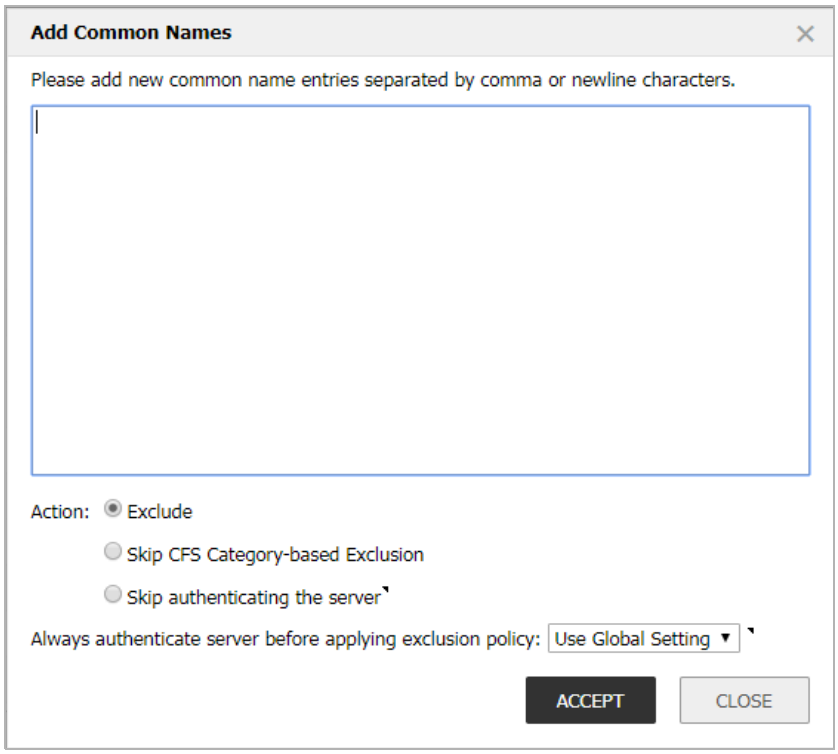

a Add one or more common names in the field. Separate multiple entries with commas or newline characters.

b Specify the type of **Action**:

- **Exclude** (default)
- **Override CFS Category-based Exclusion**
- **Skip authenticating the server** to opt out of authenticating the server for this domain if doing so results in the connection being blocked. Enable this option only if the server is a trusted domain.

c DPI-SSL dynamically determines if a connection should be intercepted (included) or excluded, based on policy or configuration. When DPI-SSL extracts the domain name for the connection, exclusion information is readily available for subsequent connections to the same server/domain.

To disable use of dynamic exclusion cache (both server IP and common-name based), select the **Always authenticate server before applying exclusion policy** checkbox. This option is not selected by default.

d Click **ACCEPT**.

The **Common Name Exclusions/Inclusions** table is updated, with **Custom** in the **Built-in** column. If the **Always authenticate server before applying exclusion policy** option has been selected an **Information** icon displays next to **Custom** in the **Built-in** column.

#	Common Name	Action	Built-in	Confi...
1	sonicwall.com	Exclude	Custom	
2	support.sonicwall.com	Exclude	Custom	

Mouse over the **Information** icon to see which custom attributes were selected. If a common name was added through the **Connection Failure List**, the Information icon indicates the type of failure:

- **Skip CFS category exclusion**
- **Skip Server authentication**
- **Failed to authenticate server**
- **Failed Client handshake**
- **Failed Server handshake**

To delete the entry, click the **Delete** icon in the **Configure** column.

7 You can search for common names by specifying a filter.

a In the **Filter** field, enter a name by specifying the name in this syntax: *name*:*mycommonname*.

b Click the **FILTER** button.

8 Click **ACCEPT**.

Deleting Custom Common Names

To delete custom common names:

1 Do one of the following:

- Clicking a custom common name's **Delete** icon in the **Configure** column.
- Selecting the name in the **Exclusions**, and then clicking the **DELETE** button.
- Clicking **DELETE ALL** to delete all custom common names. A confirmation message displays. Click **OK**.

2 Click **ACCEPT**.

Showing Connection Failures

SonicOS keeps a list of recent DPI-SSL client-related connection failures. This is a powerful feature that:

- Lists DPI-SSL failed connections.
- Allows you to audit the failed connections.
- Provide a mechanism to automatically exclude some failing domains.

The dialog displays the run-time connection failures. The connection failures could be any of the following reasons:

- Failure to handshake with the Client
- Failure to handshake with the Server
- Failed to validate the domain name in the Client Hello
- Failure to authenticate the server (the server certificate issuer is not trusted)

The failure list is only available at run-time. The number logged for each failure is limited to ensure a single failure type does not overrun the entire buffer.

To use the connection failure list:

1 Click the **SHOW CONNECTION FAILURES** button. The **Connection Failure List** dialog displays.

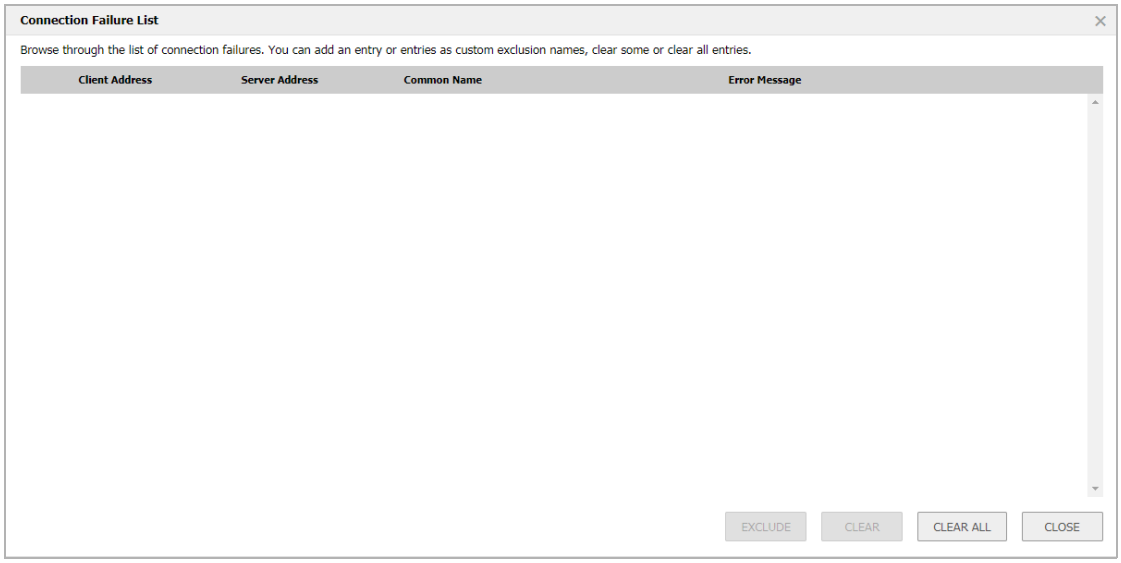

Each entry in this lists displays the:

- **Client Address**
- **Server Address**
- **Common Name** – The common name of the failed connection's domain. You can edit this entry inline before adding it to the automatic exclusion list.
- **Error Message** – Provides contextual information associated with the connection that enables you to make appropriate choices about excluding this connection.

2 To add an entry to the exclusion list:

   a Select the entry.

   b Make any edits to the entry.

   c Click the **EXCLUDE** button.

3 To delete an entry:

   a Select it.

   b Click the **CLEAR** button.

4 To delete all entries, click the **CLEAR ALL** button.

5 When you have finished, click the **CLOSE** button.

Updating Default Exclusions Manually

If your environment is closed or you prefer to update default exclusions manually, you can download the default exclusions database from www.mysonicwall.com and then import them.

To update default exclusions manually:

1 Import the default exclusions database from www.mysonicwall.com.

2 Navigate to the **MANAGE | Security Configuration > Decryption Services > DPI-SSL/TLS Client** page.

3 Scroll to the **Update Default Exclusions Manually** section.

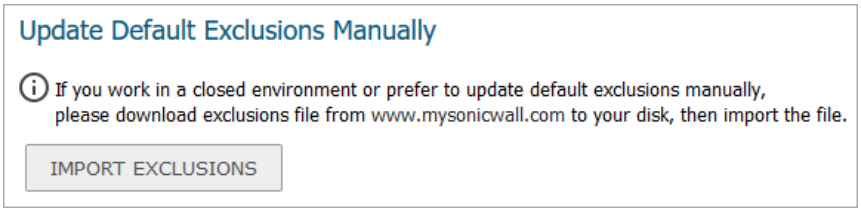

4 Click **IMPORT EXCLUSIONS**. The **Import Default Exclusions** dialog displays.

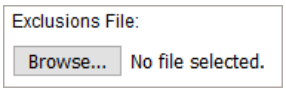

5 Click **Browse**. The **File Upload** dialog displays.

6 Open the downloaded default exclusions database file.

The **Common Name Exclusions/Inclusions** table and the status of the default database used by the firewall in the **DPI SSL Default Exclusions Status** section are updated.

Specifying CFS Category-based Exclusions/Inclusions

You can exclude/include entities by content filter categories.

To specify CFS category-based exclusions/inclusions:

1 Navigate to the **MANAGE | Security Configuration > Decryption Services > DPI-SSL/TLS Client** page.

2 Click **CFS Category-based Exclusions/Inclusions**.

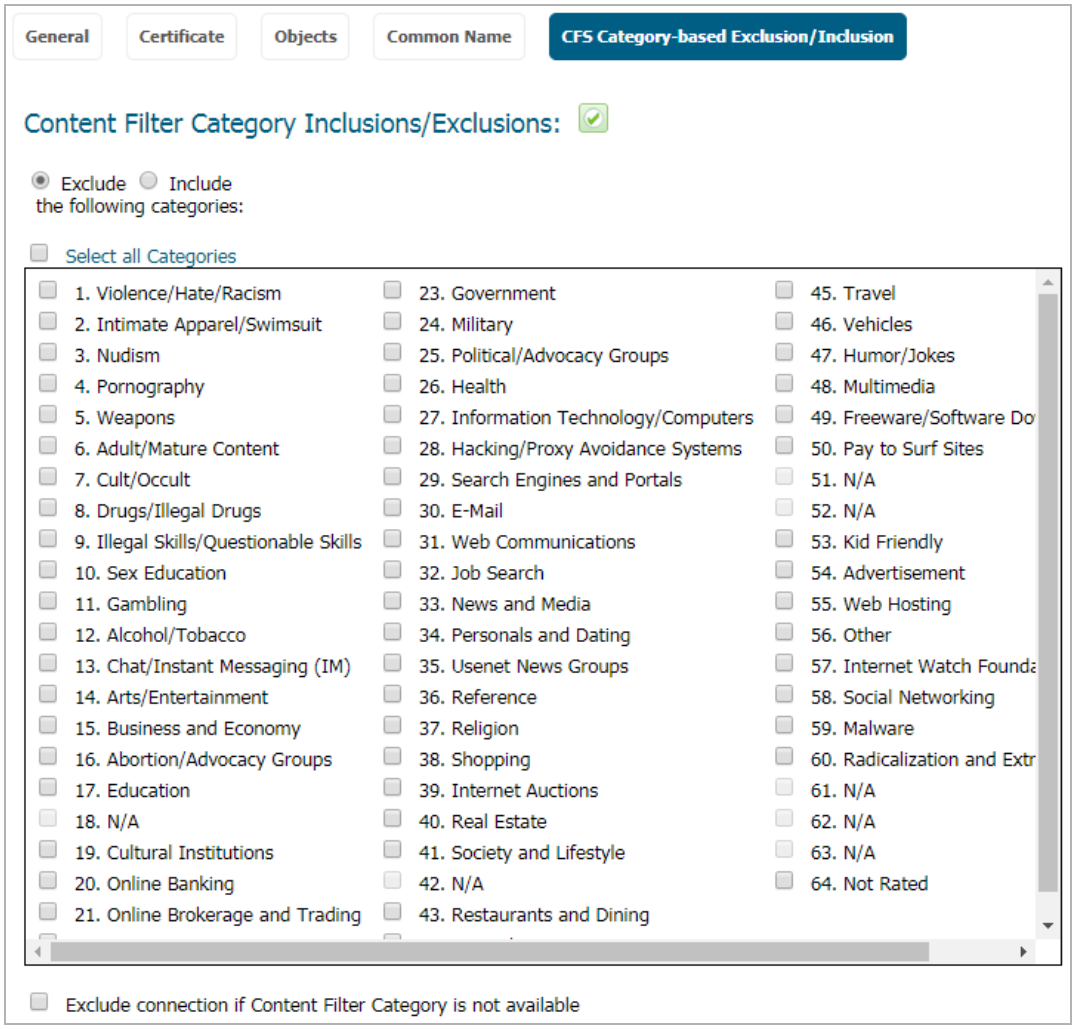

The status of the list is shown by an icon at the top of the view. A green icon indicates Content Filtering is licensed, a red icon that it is not. Mousing over the icon displays a popup with the status.

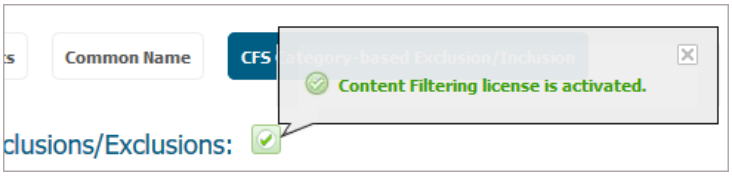

3 Choose whether you want to include or exclude the selected categories by clicking either:

- **Exclude** (default)
- **Include**

By default, all categories are unselected.

4 Select the categories to be included/excluded. To select all categories, click the **Select all Categories** checkbox.

5 Optionally, repeat Step 3 and Step 4 to create the opposite list.

6 Optionally, to exclude a connection if the content filter category information for a domain is not available to DPI-SSL, select the **Exclude connection if Content Filter Category is not available** checkbox. This option is not selected by default.

In most cases, category information for a HTTPS domain is available locally in the firewall cache. When the category information is not locally available, DPI-SSL obtains the category information from the cloud without blocking the client or server communication. In rare cases, the category information is not available for DPI-SSL to make a decision. By default, such sites are inspected in DPI-SSL.

7 Click **ACCEPT**.

Client DPI-SSL Examples

Topics:

- Content Filtering on page 264
- App Rules on page 266

Content Filtering

To perform SonicWall Content Filtering on HTTPS and SSL-based traffic using DPI-SSL:

1 Navigate to **MANAGE | Security Configuration > Security Services > Content Filter**.

2 Ensure **SonicWall CFS** is selected for the **Content Filter Type** from the drop-down menu.

3 Scroll to the **Global Settings** section.

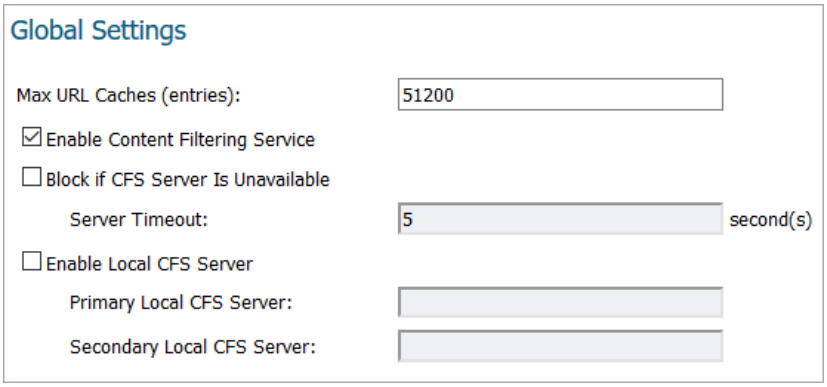

4 Select **Enable Content Filter Service**.

5 Click **ACCEPT**.

6 Navigate to the **MANAGE | Security Configuration > Decryption Services > DPI-SSL/TLS Client** page.

7 Click **General**.

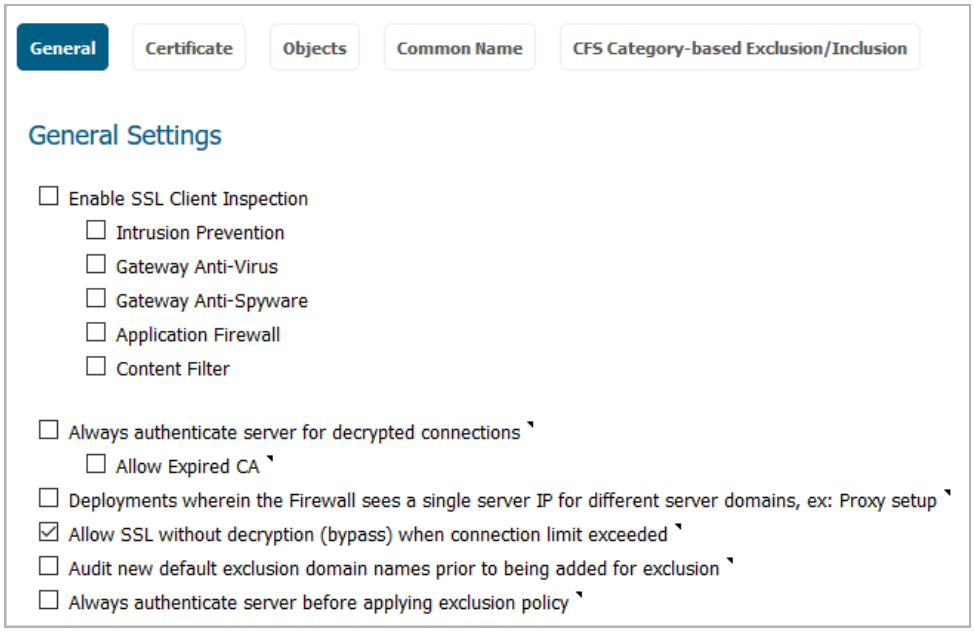

8 Select the **Enable SSL Inspection** checkbox.

9 Select the **Content Filter** checkbox.

10 Click **ACCEPT**.

11 Navigate to a blocked site using the HTTPS protocol to verify that it is properly blocked.

NOTE: For content filtering over DPI-SSL, the first time HTTPS access is blocked results in a blank page being displayed. If the page is refreshed, the user sees the firewall block page.

App Rules

To filter by application firewall rules, you need to enable them on both the **MANAGE | Security Configuration > Decryption Services > DPI-SSL/TLS Client** page and the **MANAGE | Policies > Rules > App Control** page.

1 Navigate to the **MANAGE | Security Configuration > Decryption Services > DPI-SSL/TLS Client** page.

2 Click **General**.

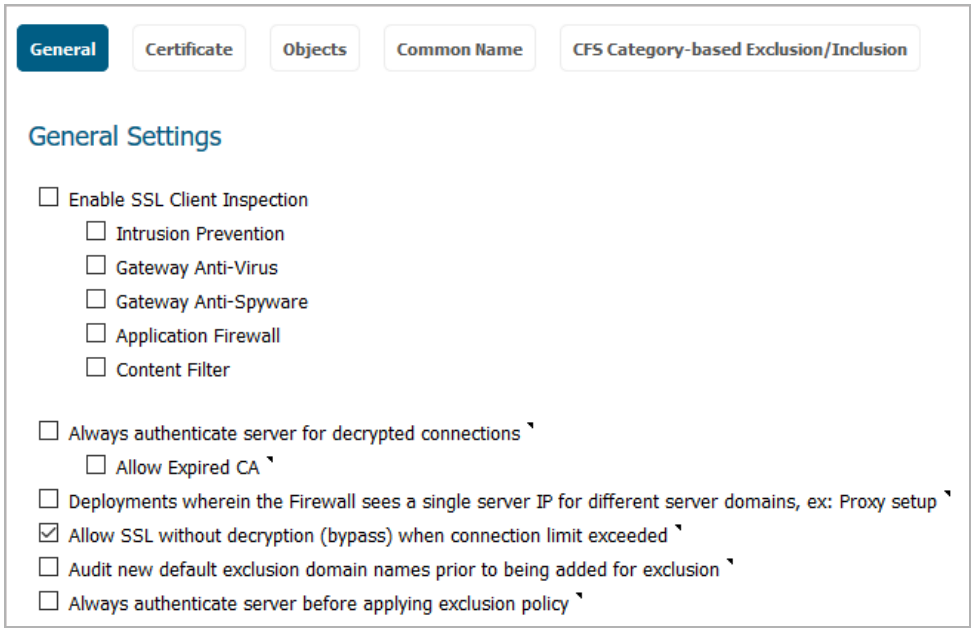

3 Select the **Enable SSL Client Inspection** checkbox.

4 Select the **Application Firewall** checkbox.

5 Click **ACCEPT**.

6 Navigate to **MANAGE | Policies > Rules > App Control** page.

7 Scroll to the **App Rules Global Settings** section.

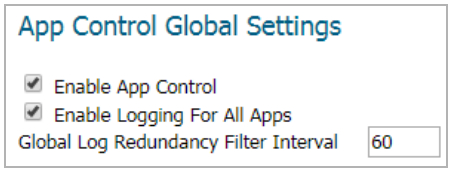

8 Select **Enable App Control**. This option is not selected by default.

9 Configure an HTTP Client policy to block Microsoft Internet Explorer browser with **block page** as an action for the policy. For how to configure an App Rule, see *SonicWall SonicOS 6.5 Policies*.

10 Click **ACCEPT**.

11 Access any website using the HTTPS protocol with Internet Explorer to verify it is blocked.

22

Configuring DPI-SSL/TLS Server Settings

- Decryption Services > DPI-SSL/TLS Server on page 267
- Configuring DPI-SSL/TLS Server Settings on page 268

Decryption Services > DPI-SSL/TLS Server

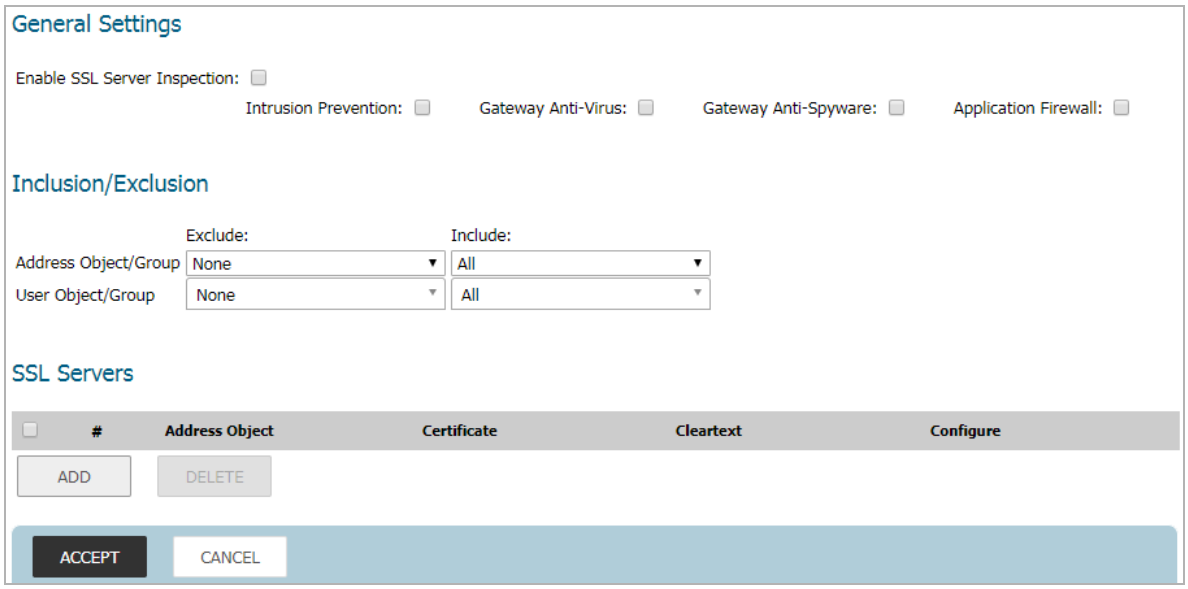

NOTE: For information about DPI SSL, see About DPI-SSL on page 244.

The Server DPI-SSL deployment scenario is typically used to inspect HTTPS traffic when remote clients connect over the WAN to access content located on the firewall's LAN. Server DPI-SSL allows you to configure pairings of an address object and certificate. When the appliance detects SSL connections to the address object, it presents the paired certificate and negotiates SSL with the connecting client.

Afterward, if the pairing defines the server to be cleartext, then a standard TCP connection is made to the server on the original (post NAT remapping) port. If the pairing is not defined to be cleartext, then an SSL connection to the server is negotiated. This allows for end-to-end encryption of the connection.

NOTE: In this deployment scenario, the owner of the firewall owns the certificates and private keys of the origin content servers. You would have to import the server's original certificate onto the appliance and create an appropriate server IP address to server certificate mappings in the Server DPI-SSL UI.

Configuring DPI-SSL/TLS Server Settings

Topics:

- Configuring General DPI-SSL/TLS Server Settings on page 268
- Configuring Exclusions and Inclusions on page 268
- Configuring Server-to-Certificate Pairings on page 269

Configuring General DPI-SSL/TLS Server Settings

To enable Server DPI-SSL inspection:

1 Navigate to the **MANAGE | Security Configuration** > **Decryption Services > DPI-SSL/TLS Server** page.

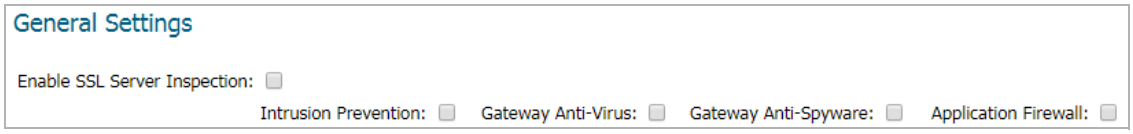

2 Scroll to the **General Settings** section.

3 Select **Enable SSL Server Inspection**.

4 Select one or more of the services with which to perform inspection:

- **Intrusion Prevent**
- **Gateway Anti-Virus**
- **Gateway Anti-Spyware**
- **Application Firewall**

5 Click **ACCEPT**.

6 Scroll down to the **SSL Servers** section to configure the server or servers to which DPI-SSL inspection is applied. See Configuring Server-to-Certificate Pairings on page 269.

Configuring Exclusions and Inclusions

By default, the DPI-SSL applies to all traffic on the appliance when it is enabled. You can configure inclusion/exclusion lists to customize to which traffic DPI-SSL inspection applies. The **Inclusion/Exclusion** lists provide the ability to specify certain objects or groups. In deployments that process a large amount of traffic, to reduce the CPU impact of DPI-SSL and to prevent the appliance from reaching the maximum number of concurrent DPI-SSL inspected connections, it can be useful to exclude trusted sources.

To customize DPI-SSL server inspection:

1 Navigate to the **MANAGE | Security Configuration > Decryption Services > DPI-SSL/TLS Server** page.

2 Scroll to the **Inclusion/Exclusion** section.



3 From **Address Object/Group Exclude**, select an address object or group to exclude from DPI-SSL inspection. By default, **Exclude** is set to **None**.

4 From **Address Object/Group Include**, select an address object or group to include in DPI-SSL inspection. By default, **Include** is set to **All**.

> **TIP:** **Include** can be used to fine tune the specified exclusion list. For example, by selecting the **Remote-office-California** address object from **Exclude** and the **Remote-office-Oakland** address object from **Include**.

5 From **User Object/Group Exclude**, select an address object or group to exclude from DPI-SSL inspection. By default, **Exclude** is set to **None**.

6 From **User Object/Group Include**, select an address object or group to include in DPI-SSL inspection. By default, **Include** is set to **All**.

7 Click **ACCEPT**.

Configuring Server-to-Certificate Pairings

Server DPI-SSL inspection requires that you specify which certificate is used to sign traffic for each server that has DPI-SSL inspection performed on its traffic.

To configure a server-to-certificate pairing:

1 Navigate to the **MANAGE | Security Configuration > Decryption Services > DPI-SSL/TLS Server** page.

2 Scroll to the **SSL Servers** section.

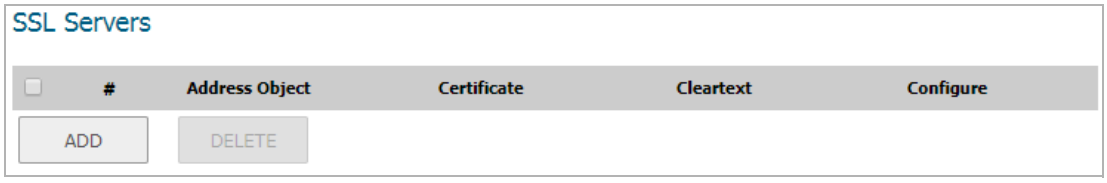

3 Click **ADD**. The **Server DPI-SSL - SSL Server Setting** dialog displays.

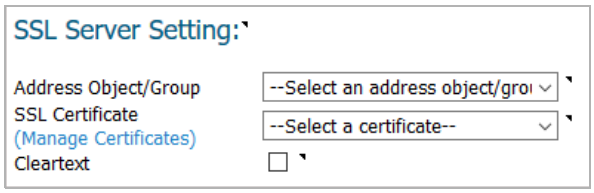

4 From **Address Object/Group**, select the address object or group for the server or servers to which you want to apply DPI-SSL inspection.

5 From **SSL Certificate**, select the certificate to be used to sign the traffic for the server. For more information on:

- Importing a new certificate to the appliance, see Selecting the Re-Signing Certificate Authority on page 253.
- Creating a Linux certificate, see *SonicWall SonicOS 6.5 System Setup*.

> **TIP:** Clicking the `(Manage Certificates)` link displays the **MANAGE | System Setup > Appliance > Certificates** page.

6 Select **Cleartext** to enable SSL offloading. When adding server-to-certificate pairs, the **Cleartext** option provides a method of sending unencrypted data onto a server. This option is not selected by default..

IMPORTANT: For such a configuration to work properly, a NAT policy needs to be created for this server on the **MANAGE | Policies > Rules > NAT Policies** page to map traffic destined for the offload server from an SSL port to a non-SSL port. Traffic must be sent over a port other than 443. For example, for HTTPS traffic used with SSL offloading, an inbound NAT policy remapping traffic from port 443 to port 80 needs to be created for things to work properly.

7 Click **ADD**.

23

Configuring DPI-SSH

- About DPI-SSH
- Activating Your DPI-SSH License
- Configuring DPI-SSH

About DPI-SSH

IMPORTANT: Gateway Anti-Spyware service does not work for DPI-SSH because TCP streams for Anti-Spyware are not supported. If the option is checked, the system takes no action.

Deep Packet Inspection (DPI) technology allows a packet filtering-firewall to classify passing traffic based on signatures of the Layer 3 and Layer 4 contents of the packet. DPI also provides information that describes the contents of the packet's payload (the Layer 7 application data). DPI is an existing SonicOS feature that examines the data and the header of a packet as it passes through the SonicWall firewall, searching for protocol non-compliance, viruses, spam, intrusions, or defined criteria to decide whether the packet may pass or if it needs to be routed to a different destination for action or other tracking.

SSH (Secure Shell) is a cryptographic network protocol for secure data communication, remote command-line login, remote command execution, and other secure network services between two networked computers. SSH connects, via a secure channel over an insecure network—a server and a client running SSH server and SSH client programs, respectively. The protocol distinguishes between two different versions, referred to as SSH-1 and SSH-2. SonicWall only supports SSH-2; SSH-1 sessions are not intercepted and inspected.

IMPORTANT: SSH clients with different version numbers cannot be used at the same time.f

To effectively inspect an encrypted message, such as SSH, the payload must be decrypted first. DPI-SSH works as a man-in-the-middle (MITM) or a packet proxy. Any preset end-to-end communication is broken, and pre-shared keys cannot be used.

DPI-SSH divides the one SSH tunnel into two tunnels as it decrypts the packets coming from both tunnels and performs the inspection. If the packet passes the DPI check, DPI-SSH sends the re-encrypted packet to the tunnels. If the packet fails the check, it's routed to another destination, based on the policies, or submitted for collecting statistical information, and DPI-SSH resets the connection.

Topics:

- Supported Clients/Servers and Connections on page 272
- Supported Key Exchange Algorithms on page 272
- Caveats on page 272

Supported Clients/Servers and Connections

SSH is not a shell, but a secure channel that provides different services over this channel (tunnel), including shell, file transfer, or X11 forwarding.

DPI-SSH supports both route mode and Wire Mode. For Wire Mode, DPI-SSH is only supported in the secure (active DPI of inline traffic) mode. For route mode, there is no limitation.

SSH supports different client and server implementations, as listed in Supported clients/servers.

Supported clients/servers

DPI-SSH Client Supported	DPI-SSH Servers Supported
SSH client for Cygwin	SSH server on Fedorz
Putty	SSH server on Ubuntu
secureCRT	
SSH on Ubuntu	
SSH n centos	
SFTP client on Cygwin	
SCP on Cygwin	
Winscp	

DPI-SSH supports up to 250 connections.

Supported Key Exchange Algorithms

DPI-SSH supports these key exchange algorithms:

- Diffie-Hellman-group1-sha1
- Diffie-Hellman-group14-sha1
- ecdh-sha2-nistp256

DPI-SSH supports DSA keys on the client side and RSA keys on the server side.

Caveats

If there is already an SSH server key stored in the local machine, it must be deleted. For example, if you already SSH to a server, and the server DSS key is saved, the SSH session fails if the DSS key is not deleted from the local file.

The `ssh-keygen` utility cannot be used to bypass the password.

Putty uses GSSAPI. This option is for SSH2 only, which provides stronger encrypted authentication. It stores a local token or secret in the local client and server for the first time communication. It exchanges messages and operations before DPI-SSH starts, however, so DPI-SSH has no knowledge about what was exchanged before, including he GSSAPI token. DPI-SSH fails with the GSSAPI option enabled.

On the client side, either the SSH 2.x or 1.x client can be used if DPI-SSH is enabled. Clients with different version numbers, however, cannot be used at the same time.

Gateway Anti-Spyware and Application Firewall inspections are not supported even if these options are selected in the **MANAGE | Security Configuration > Decryption Services > DPI-SSH** page.

Activating Your DPI-SSH License



DPI-SSH is fully licensed by default, but you need to activate your license. When you first select **MANAGE | Security Configuration** > **Decryption Services > DPI-SSH**, you receive the message: `Upgrade Required`.

If the upgrade isn't required, skip to Configuring DPI-SSH on page 273. For information about activating your license, see the *Quick Start Guide* for your appliance.

Configuring DPI-SSH

IMPORTANT: Gateway Anti-Spyware service doesn't work for DPI-SSH because TCP streams for Anti-Spyware are not supported. If the checkbox is checked, the system takes no action.

You configure DPI-SSH on the **MANAGE | Security Configuration > Decryption Services > DPI-SSH** page.

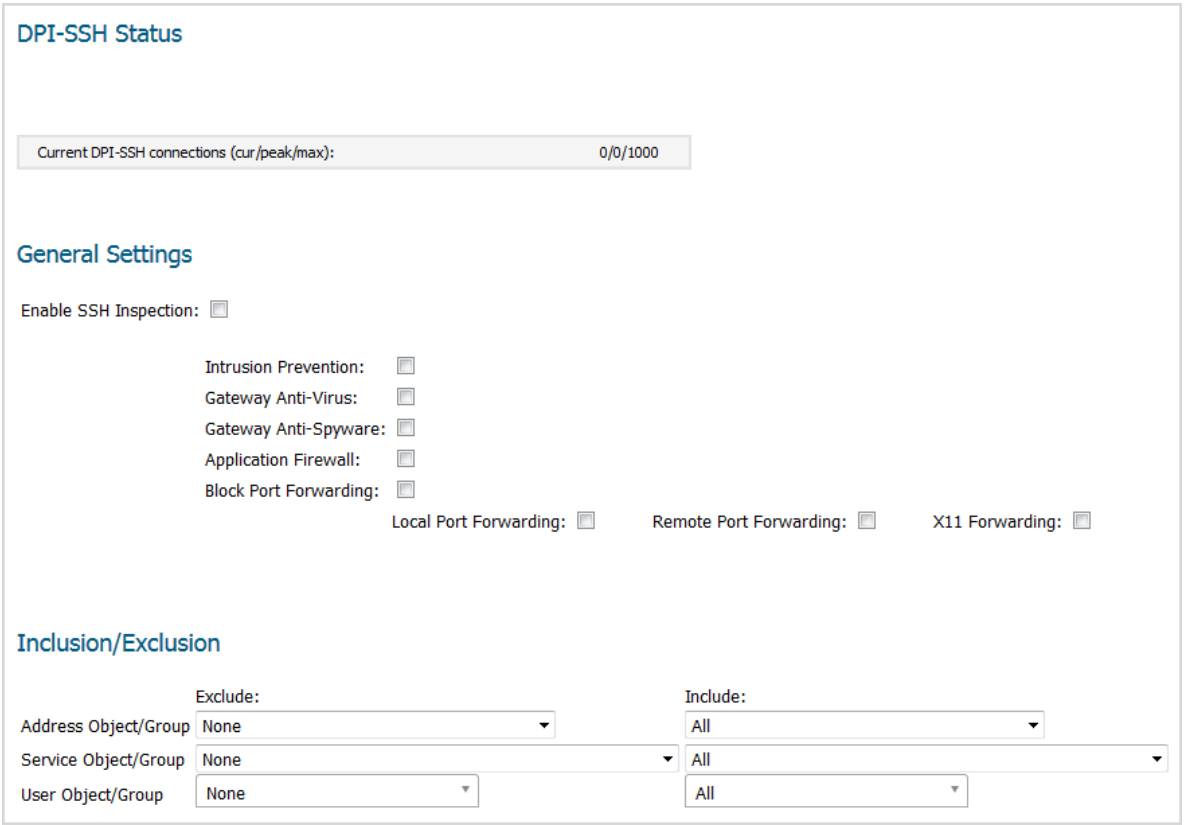

Topics:

- Viewing Connection Status on page 274
- Configuring Client DPI-SSH Inspection on page 274
- Customizing Client DPI-SSH Inspection on page 276

Viewing Connection Status

To view the status of DPI-SSH connections:

1 Navigate to **MANAGE | Security Configuration > Decryption Services > DPI-SSH**.

2 Scroll to **DPI-SSH Status**.

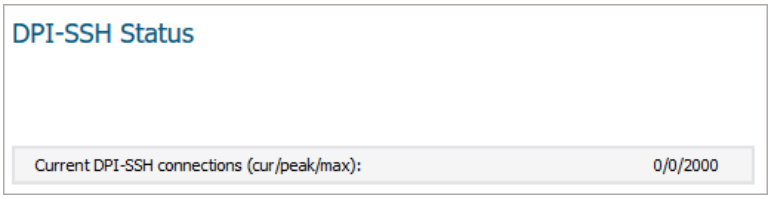

The status displays the number of:

- Current DPI-SSH connections
- Peak DPI-SSH connections
- Maximum number of DPI-SSH connections

Configuring Client DPI-SSH Inspection

You configure Client DPI-SSH inspection in the **General Settings** section of **Decryption Services > DPI-SSH**.

To enable Client DPI-SSH inspection:

1 In the **General Settings** section, select the **Enable SSH Inspection** option. This option is not selected by default.

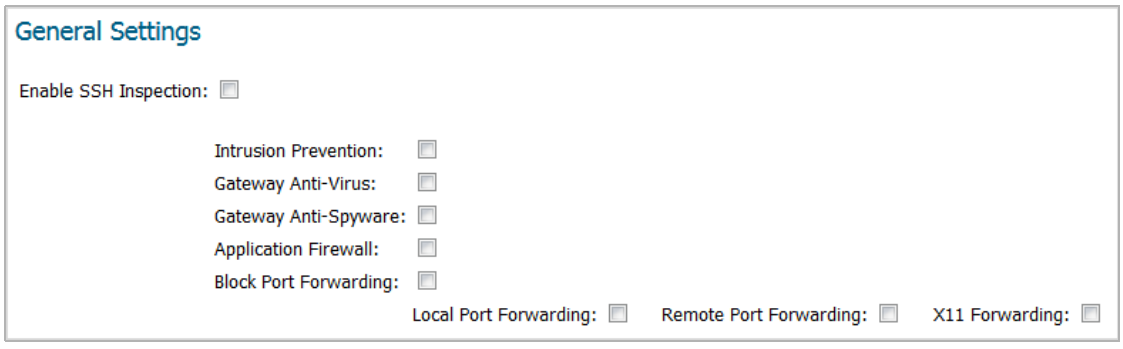

2 Select one or more types of service inspections; none are selected by default:

- **Intrusion Prevention**
- **Gateway Anti-Virus**

- **Gateway Anti-Spyware**

  IMPORTANT: Gateway Anti-Spyware service doesn't work for DPI-SSH because TCP streams for Anti-Spyware are not supported. If the option is checked, the system takes no action.

- **Application Firewall**
- **Block Port Forwarding**: for more information about these options, see DPI-SSH Blocking of Port Forwarding on page 275:
  - **Local Port Forwarding**
  - **Remote Port Forwarding**
  - **X11 Forwarding**

3 Click **ACCEPT**.

DPI-SSH Blocking of Port Forwarding

SSH makes it possible to tunnel other applications through SSH by using port forwarding. Port forwarding allows local or remote computers (for example, computers on the internet) to connect to a specific computer or service within a private LAN. Port forwarding translates the address and/or port number of a packet to a new destination address and forwards it to that destination according the routing rules. Since these packets have new destination and port numbers, they can bypass the firewall security policies.

To prevent circumvention of the application-based security policies on the SonicWall network security appliance, SonicOS supports blocking SSH port forwarding for both Local and Remote port forwarding.

- *Local port forwarding* allows a computer on the local network to connect to another server, which might be an external server.
- *Dynamic port forwarding* allows you to configure one local port for tunneling data to all remote destinations. This can be considered as a special case of *Local port forwarding*.
- *Remote port forwarding* allows a remote host to connect to an internal server.

SSH port forwarding supports the following servers:

- SSH server on Fedora
- SSH server on Ubuntu

SSH port forwarding supports both:

- Route mode
- Wire mode – only supported in Secure Mode

SSH port forwarding supports a maximum of 1000 connections, matching the maximum supported by DPI-SSH.

DPI-SSH must be enabled for blocking of SSH port forwarding to work. If any local or remote port forwarding requests are made when the blocking feature is enabled, SonicOS blocks those requests and resets the connection.

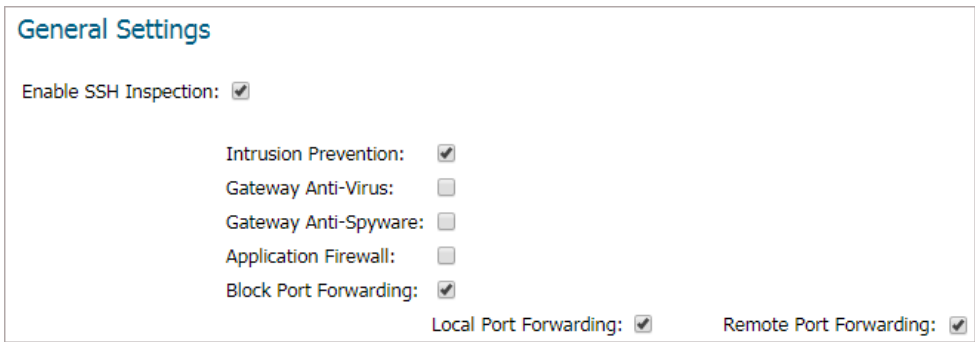

To enable blocking of SSH port forwarding:

1. Navigate to the **MANAGE | Security Configuration > Decryption Services > DPI-SSH** page.
2. In the **General Settings** section, select **Block Port Forwarding**.
3. Select either or both **Local Port Forwarding** and **Remote Port Forwarding** to block that type of port forwarding.
4. Click **ACCEPT**.

DPI-SSH port forwarding supports the following clients:

- SSH client for Cygwin
- Putty
- SecureCRT
- SSH on Ubuntu
- SSH on CentOS

Customizing Client DPI-SSH Inspection



By default, when DPI-SSH is enabled, it applies to all traffic on the firewall. You can customize to which traffic DPI-SSH inspection applies in the **Inclusion/Exclusion** section.

To customize DPI-SSH client inspection:

1. Go to the **Inclusion/Exclusion** section of the **Decryption Services > DPI-SSH** page.
2. From the **Address Object/Group Exclude** and **Include** drop-down menus, select an address object or group to exclude or include from DPI-SSH inspection. By default, **Exclude** is set to **None** and **Include** is set to **All**.

3 From the **Service Object/Group Exclude** and **Include** drop-down menus, select an address object or group to exclude or include from DPI-SSH inspection. By default, **Exclude** is set to **None** and **Include** is set to **All**.

4 From the **User Object/Group Exclude** and **Include** drop-down menus, select an address object or group to exclude or include from DPI-SSH inspection. By default, **Exclude** is set to **None** and **Include** is set to **All**.

5 Click **ACCEPT**.

Part 4

Security Config | Anti-Spam

IMPORTANT: Anti-Spam is not supported on the SuperMassive series or the NS*a* 9250-9650 firewalls.

- About Anti-Spam
- Enabling and Activating Anti-Spam
- Configuring Anti-Spam Logging
- Configuring the RBL Filter
- Specifying Relay Domains
- Configuring Junk Box Settings
- Managing the Junk Summary
- Configuring the Junk Box View
- Configuring User-Visible Settings
- Configuring Corporate Allowed and Blocked Lists
- Managing Users
- Configuring the LDAP Server
- Downloading Anti-Spam Desktop Buttons

24

About Anti-Spam

IMPORTANT: Anti-Spam is not supported on the SuperMassive series or the NS*a* 9250-9650 firewalls.

NOTE: Anti-Spam is a separate, licensed feature that provides a quick, efficient, and effective way to add anti-spam, anti-phishing, and anti-virus capabilities to your existing firewall.

- About Anti-Spam on page 279
- How Does the Anti-Spam Service Work? on page 280
- Purchasing an Anti-Spam License on page 285

About Anti-Spam

Topics:

- What is Anti-Spam? on page 279
- Benefits on page 280

What is Anti-Spam?

The Anti-Spam feature provides a quick, efficient, and effective way to add anti-spam, anti-phishing, and anti-virus capabilities to your existing firewall.

In a typical Anti-Spam configuration, you choose to add Anti-Spam capabilities by selecting it in the SonicOS interface and licensing it. The firewall then uses the same advanced spam-filtering technology as the SonicWall Email Security products to reduce the amount of junk email delivered to users.

There are two primary ways inbound messages are analyzed by the Anti-Spam feature:

- Advanced IP Reputation Management
- Cloud-based Advanced Content Management

IP Address Reputation uses the GRID Network to identify the IP addresses of known spammers, and reject any mail from those senders without even allowing a connection. GRID Network Sender IP Reputation Management checks the IP address of incoming connecting requests against a series of lists and statistics to ensure that the connection has a probability of delivering valuable email. The lists are compiled using the collaborative intelligence of the SonicWall GRID Network. Known spammers are prevented from connecting to the firewall, and their junk email payloads never consume system resources on the targeted systems.

Email that does not come from known spammers is analyzed based on “GRIDprints” generated by SonicWall’s research laboratories and are based on data from millions of business endpoints, hundreds of millions of messages, and billions of reputation votes from the users of the GRID Network. Our Grid Network uses data from multiple SonicWall solutions to create a collaborative intelligence network that defends against the

worldwide threat landscape. GRIDprints uniquely identify messages without exposing data contained in the email message.

The Anti-Spam service determines that an email fits *only one* of the following threats: Spam, Likely Spam, Phishing, Likely Phishing, Virus, or Likely Virus. It uses the following precedence order when evaluating threats in email messages:

- Phishing
- Likely Phishing
- Virus
- Likely Virus
- Spam
- Likely Spam

For example, if a message is both a virus and spam, the message is categorized as a virus as virus is higher in precedence than spam.

If the Anti-Spam service determines that the message is *not* any of the above threats, it is judged as good email and is delivered to the destination server.

Benefits

Adding anti-spam protection to your firewall increases the efficiency of your system as a whole by filtering and rejecting junk messages before users see them in their inboxes.

- Reduced amount of bandwidth and resources consumed by junk email in your network
- Reduced number of incoming messages sent to the mail server
- Reduced threat to the organization, because users cannot accidentally infect their computers by clicking on virus spam
- Better protection for users from phishing attacks

How Does the Anti-Spam Service Work?

This section describes the Anti-Spam feature, including the SonicWall GRID Network, and how it interacts with SonicOS as a whole. The two points of significant connection with SonicOS are Address and Service Objects. You use the address and service objects to configure the Anti-Spam feature to function smoothly with SonicOS. For example, use the Anti-Spam Service Object to configure NAT policies to archive inbound email as well as sending it through a filter.

The Comprehensive Anti-Spam Service analyzes messages' headers and contents and uses collaborative GRID printing to block spam email.

Topics:

- GRID Network on page 280
- Address and Service Objects on page 282

GRID Network

The GRID Connection Management with Sender IP Reputation feature is used by SonicWall Email Security and by the Anti-Spam service in SonicOS. GRID Network Sender IP Reputation is the reputation a particular IP address has with members of the SonicWall GRID Network. When this feature is enabled, email is not accepted from IP addresses with a bad reputation. When SonicOS does not accept a connection from a known bad IP address, mail from that IP address never reaches the email server.

GRID Network Sender IP Reputation checks the IP address of incoming connection requests against a series of lists and statistics to ensure that the connection has a probability of delivering valuable email. The lists are compiled using the collaborative intelligence of the SonicWall GRID Network. Known spammers are prevented from connecting to the firewall, and their junk email payloads never consume system resources on the targeted systems.

Topics:

- Benefits
- GRID Connection Management with Sender IP Reputation and Connection Management Precedence Order

Benefits

- As much as 80 percent of junk email is blocked at the connection level, before the email is ever accepted into your network. Fewer resources are required to maintain your level of spam protection.
- Your bandwidth is not wasted on receiving junk email on your servers, only to analyze and delete it.
- A global network watches for spammers and helps legitimate users restore their IP reputations if needed.

GRID Connection Management with Sender IP Reputation and Connection Management Precedence Order

When a request is sent to your first-touch firewall, the Anti-Spam service evaluates the 'reputation' of the requester. The reputation is compiled from white lists of known-good senders, block lists of known spammers, and denial-of-service thresholds.

If IP Reputation is enabled, the source IP address is checked in the order shown in Evaluation order:

Evaluation order

Evaluation	Description
Allow-list	If an IP address is on this list, it is allowed to pass messages through Connection Management. The messages are analyzed by your firewall as usual.
Block-list	This IP address is banned from connecting to the firewall.
Reputation-list	If the IP address is not in the previous lists, the firewall checks with the GRID Network to see if this IP address has a bad reputation.
Defer-list	Connections from this IP address are deferred. A set interval must pass before the connection is allowed.
DoS	If the IP address is not on the previous lists, the firewall checks to see if the IP address has crossed the Denial of Service threshold. If it has, the appliance uses the existing DoS settings to take action.

Only if the IP address passes all of these tests does the firewall allow that server to make a connection and transfer mail. If the IP address does not pass the tests, there is a message from SonicOS to the requesting server indicating that there is no SMTP server. The connection request is not accepted.

Address and Service Objects

The Anti-Spam feature of SonicOS supports Address and Service Objects to manage a customer's email server(s). These objects are used by the Anti-Spam Service for its NAT and Access Rule policies. Automatically-created rules are not editable and will be deleted if the Anti-Spam Service is disabled.

When enabled, the Anti-Spam service creates NAT policies and Access Rules to control and redirect email traffic. The policies and rules are visible in the **MANAGE | Policies > Rules > NAT Policies** page, but are not editable. These automatically-created policies are only available when the Anti-Spam service is enabled. For further information about these rules and policies, see *SonicWall SonicOS 6.5 Policies*.

When the Anti-Spam service is licensed and activated, the **MANAGE | Security Configuration > Anti-spam > Base Setup** page shows a single option to enable Anti-Spam. Selecting the option invokes the **Destination Mail Server Policy Wizard** if there is no existing custom access rule and NAT policy for an already-deployed scenario. When you set up generated policies, the Anti-Spam service must know where the emails are routed behind the firewall. Specifically it needs the destination mail server IP address and its zone assignment. The **Destination Mail Server Policy Wizard** is launched if this data cannot be found.

You need the following information for the wizard:

- **Destination Mail Server Public IP Address** – The IP address to which external MTAs (message transfer agents) connect by SMTP.
- **Destination Mail Server Private IP Address** – The internal IP address of the Exchange or SMTP server (behind the firewall).
- **Zone Assignment** – The zone to which the Exchange server is assigned.
- **Inbound Email Port** – The TCP service port number to which emails will be sent, also known as the inbound SMTP port.

If this information is needed, this message displays:

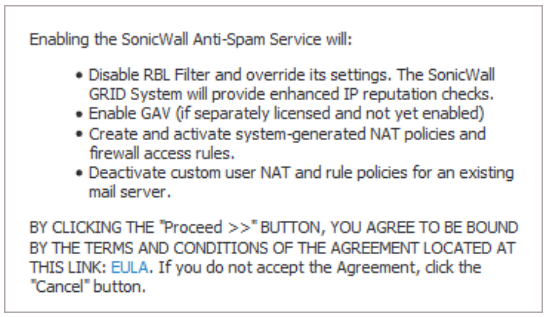

Clicking PROCEED walks you through the wizard's requests

Policies and Address Objects created by the wizard are editable and persist even if the Anti-Spam service is disabled.

Topics:

- Objects Created When the Anti-Spam Service Is Enabled
- Objects Created by the Wizard
- Policy and Object Changes

Objects Created When the Anti-Spam Service Is Enabled

This section provides an example of the type of rules and objects generated automatically as Firewall Access Rules, NAT Policies and Service Objects. These objects are not editable and will be removed if the Anti-Spam service is disabled.

The **MANAGE | Policies > Rules > Access Rules** page shows the generated rules used for Anti-Spam.

#	From	To	Priority	Source	Destination	Service	Action	Users Incl.	Users Excl.
183	VPN	WLAN	9	Any	All Interface IPv6 Addresses	HTTPS Management	Allow	All	None
184	WAN	WAN	4	Any	Public Mail Server Address Group	SMTP (Anti-Spam Inbound Port)	Allow	All	None
185	WAN	DMZ	1	Any	Any	Any	Deny	All	None
186	WAN	LAN	1	Any	User Mail Server Public IP	SMTP (Anti-Spam Inbound Port)	Allow	All	None
187	WAN	LAN	2	Any	Default Active WAN IP	SonicWALL Anti-Spam Service	Allow	All	None
188	WAN	LAN	3	Any	Public Mail Server Address Group	SMTP (Anti-Spam Inbound Port)	Allow	All	None
189	WAN	LAN	4	Any	Any	Any	Deny	All	None

The rows outlined in red are the access rules generated when Anti-Spam is activated. The row outlined in green is the default rule that Anti-Spam creates if there are no existing mail server policies.

You could also create the following access rules:

- WAN to WAN rule for incoming email (SMTP) from any source to all the WAN IP addresses
- WAN to LAN rule for processed email from Email Security Service to all the WAN IP address using the Anti-Spam service port (default:10025)

The Anti-Spam Service Object is created in the **Policies | Objects > Service Objects** page.

#	Name	Protocol	Port Start	Port End	Class	Comments	Configure
156	SonicWALL Anti-Spam Service	TCP	10025	10025	Default		

This Service Object is referenced by the generated NAT policies.

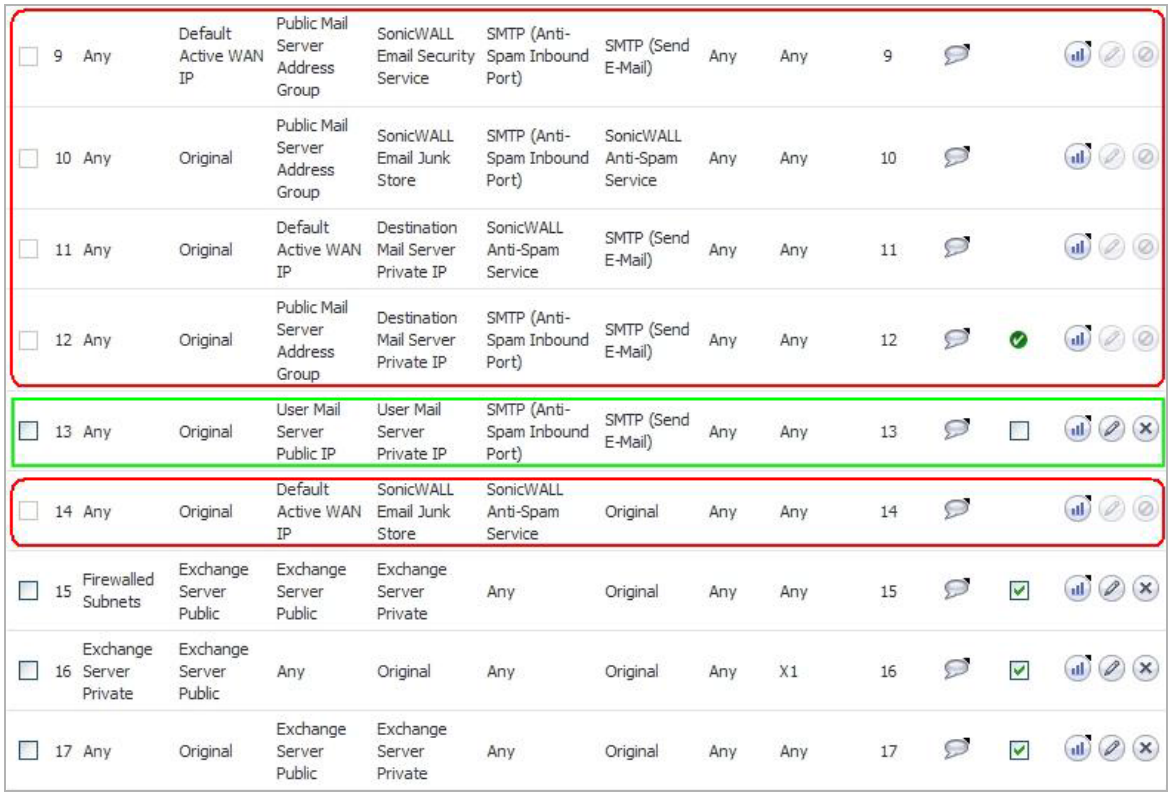

9	Any	Default Active WAN IP	Public Mail Server Address Group	SonicWALL Email Security Service	SMTP (Anti-Spam Inbound Port)	SMTP (Send E-Mail)	Any	Any	9			
10	Any	Original	Public Mail Server Address Group	SonicWALL Email Junk Store	SMTP (Anti-Spam Inbound Port)	SonicWALL Anti-Spam Service	Any	Any	10			
11	Any	Original	Default Active WAN IP	Destination Mail Server Private IP	SonicWALL Anti-Spam Service	SMTP (Send E-Mail)	Any	Any	11			
12	Any	Original	Public Mail Server Address Group	Destination Mail Server Private IP	SMTP (Anti-Spam Inbound Port)	SMTP (Send E-Mail)	Any	Any	12			
13	Any	Original	User Mail Server Public IP	User Mail Server Private IP	SMTP (Anti-Spam Inbound Port)	SMTP (Send E-Mail)	Any	Any	13			
14	Any	Original	Default Active WAN IP	SonicWALL Email Junk Store	SonicWALL Anti-Spam Service	Original	Any	Any	14			
15	Firewalled Subnets	Exchange Server Public	Exchange Server Public	Exchange Server Private	Any	Original	Any	Any	15			
16	Exchange Server Private	Exchange Server Public	Any	Original	Any	Original	Any	X1	16			
17	Any	Original	Exchange Server Public	Exchange Server Private	Any	Original	Any	Any	17			

The rows outlined in red are the policies generated when Anti-Spam is activated. The row outlined in green is the default policy that Anti-Spam creates if there are no existing mail server policies.

Objects Created by the Wizard

Objects created from an administrator's interaction with the wizard can be edited and stay in the system even if the Anti-Spam service is disabled.

The following considerations apply to the auto-generation of policies:

- A system Address Group Object called the **Public Mail Server Address Group** is created as a default for the original destination for generated policies. This group contains the Address Object, **Destination Mail Server Public IP**, which takes the IP address value provided during the wizard.
- If a SonicWall device already has existing policies for SMTP, the following procedures occur:
  - If the existing policy's original destination is a host-type Address Object, then the generated policies use the **Public Mail Server Address Group** object as their original destination.
  - If the existing policy's original destination is a non-host-type Address Object, the generated policies use this non-host type Address Object as their original destination.
  - If there is more than one public IP address for SMTP, you can manually add Address Objects to the **Public Mail Server Address Group**.

Policy and Object Changes

In the `diag.html` page, the **Reset GRID Name Cache** button can be used to clear all the entries in the GRID name cache.

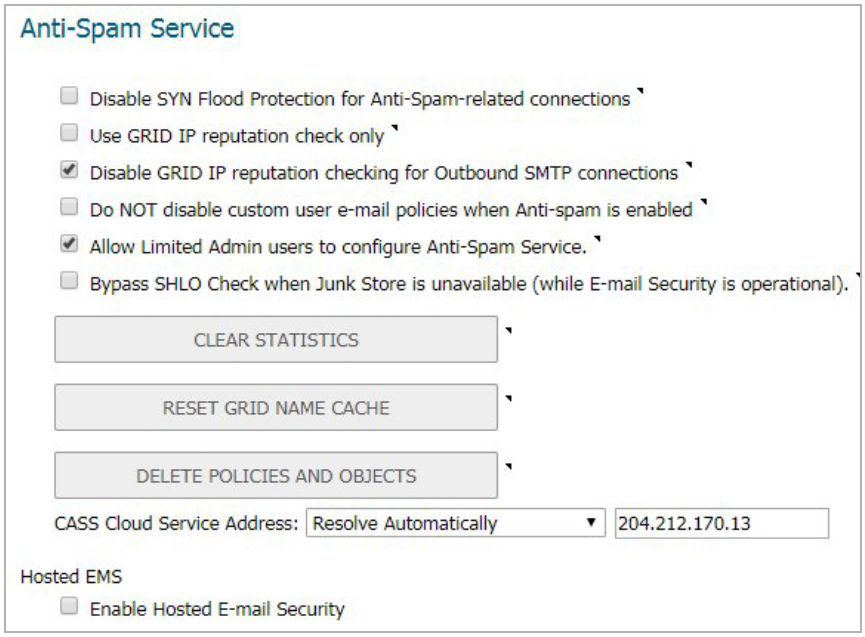

The **Delete Policies and Objects** button can be used to remove Anti-Spam Address and Service Objects and policies that are not deleted when the service is turned off. When this button is clicked, SonicOS attempts to remove all the automatically generated objects and policies. This operation is only allowed when the Anti-Spam service is off.

The other **diag.html** page options relating to Anti-Spam are:

- **Disable SYN Flood Protection for Anti-Spam related connections** – SYN Flood protection by default is turned on for SMTP (25) and Anti-Spam service (10025) ports. This disables the protection.
- **Use GRID IP reputation check only** – When selected, this overrides the probing result and simulates the Anti-Spam service being unavailable (admin down). When an email is sent, it still goes through both the SYN FLOOD check and GRID IP check, but other email scanning is not performed.

Purchasing an Anti-Spam License

The following deployment prerequisites are required to use the Anti-Spam feature:

- A licensed SonicWall network security appliance
- Anti-Spam License for the appliance
- One of the following Microsoft Windows Servers:
    - Windows Server 2012 R2 (64-bit)
    - Windows Server 2012 (64-bit)
    - Windows SBS 2008 R2 Server (64-bit)
    - SBS 2008 (64-bit)

Purchasing an Anti-Spam license for the firewall can be done directly through mySonicWall.com or through your reseller.

NOTE: Your SonicWall network security appliance must be registered with mySonicWall.com before use.

To purchase an Anti-Spam license:

1. Open a Web browser on the computer you use to manage your SonicWall appliance.
2. Enter **http://www.mySonicWall.com** in the **location** or **address** field.
3. Enter your mySonicWall.com account **user name** and **password** in the appropriate fields.
4. Click the **submit** button.
5. Navigate to **My Products** in the left-hand navigation bar.

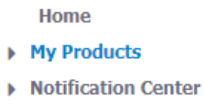

6. Select the appliance to which you wish to add Anti-Spam capability.
7. Register for an Anti-Spam license.
8. Login to your appliance's web management interface.
9. Navigate to the **MANAGE | Updates > Licenses** page from the navigation bar.mySonicWall.com.

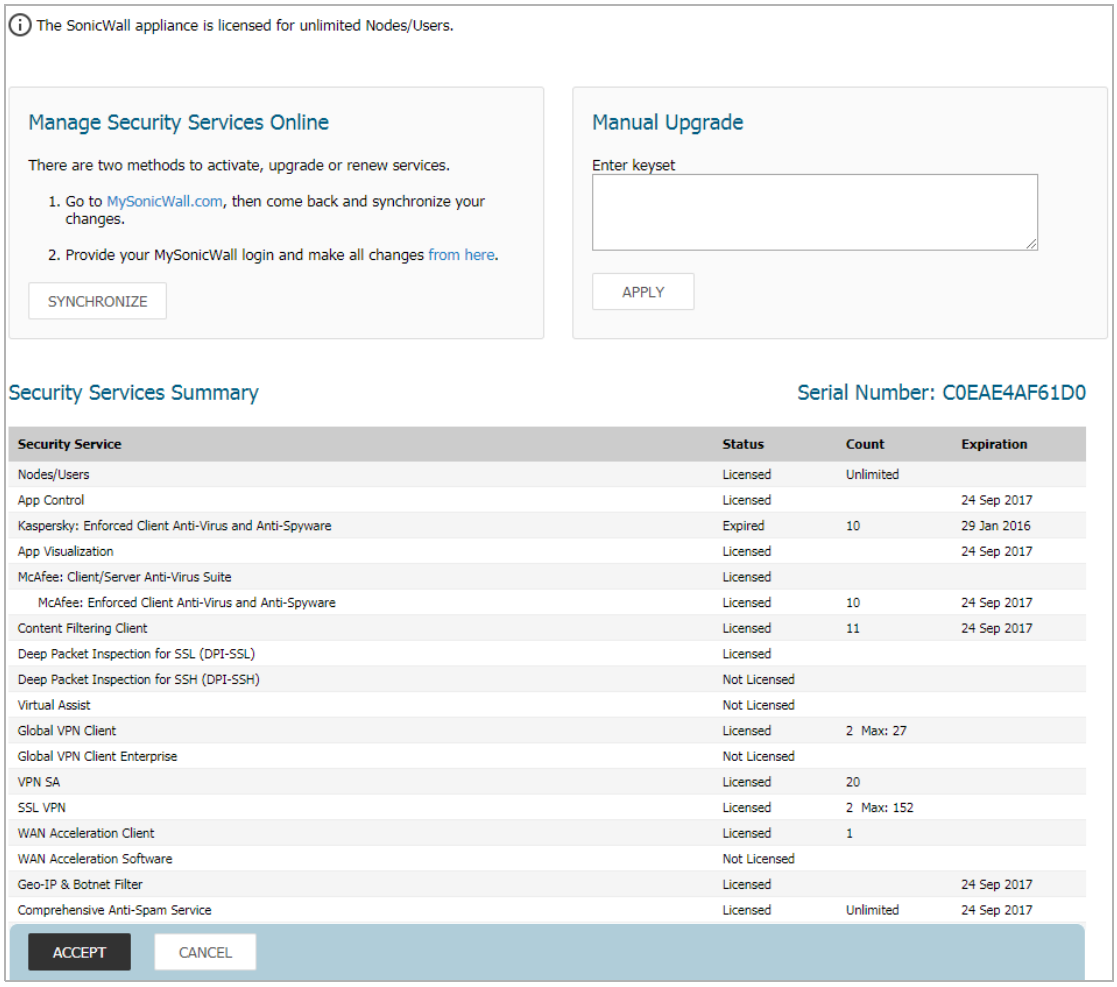

10 In the **Manage Security Services Online** section, click the link to activate or renew your license. Alternately, enter your key or keyset in the **Manual Upgrade** section.

11 Enter your `mySonicWall.com` login information.

25

Enabling and Activating Anti-Spam

NOTE: Anti-Spam is not supported on the SuperMassive series or the NS*a* 9250-9650 firewalls.

TIP: For information about the Anti-Spam feature and how to license it, see About Anti-Spam on page 279.

- Anti-Spam > Base Setup on page 289
- Activating Anti-Spam on page 289
- Installing the Junk Store on page 291
- Configuring Email Threat Categories on page 292
- Configuring Access Lists on page 293
- Configuring Advanced Options on page 295

Anti-Spam > Base Setup

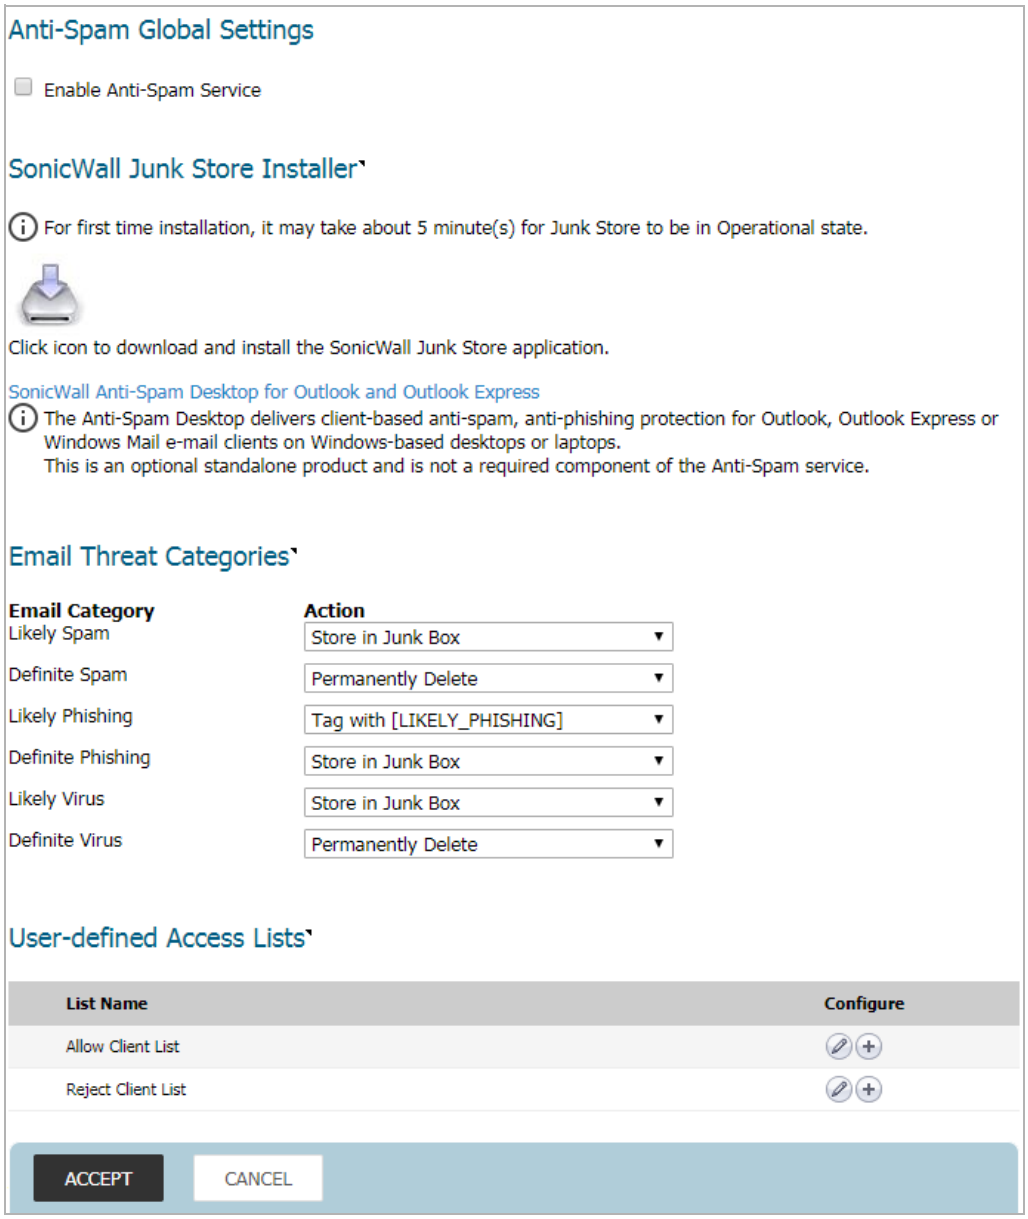

The **MANAGE | Security Configuration > Anti-Spam > Base Setup** page allows you to activate the Anti-Spam feature, configure email threat categories, modify access lists, and set advanced options.

TIP: For information about the Anti-Spam feature and how to license it, see About Anti-Spam on page 279.

Activating Anti-Spam

After you have registered Anti-Spam, activate it to start your appliance-level protection from spam, phishing, and virus messages.

To activate Anti-Spam:

1 Navigate to **MANAGE | Security Configuration > Anti-Spam > Base Setup**.

2 Scroll to the **Anti-Spam Global Settings** section.

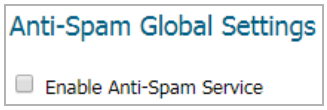

3 Click **Enable Anti-Spam Service** to activate the Anti-Spam feature. A message displays describing the effects of enabling the Anti-Spam Service and requesting agreement to proceed.

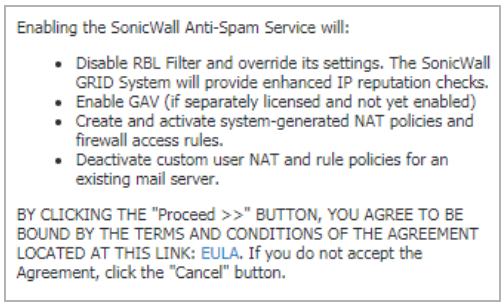

4 To proceed, click the **PROCEED>>** button. Another message about the mail server to be used displays.

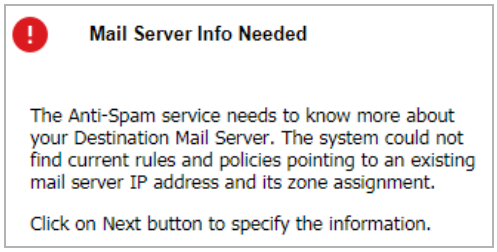

5 Click the **NEXT>>** button. A dialog requesting information about the server displays. The dialog's settings are populated with information taken from the system.

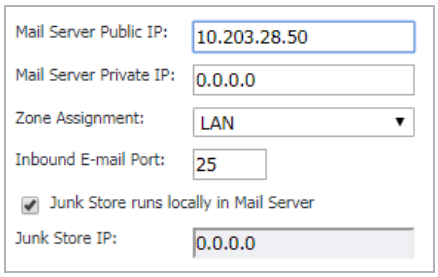

6 Optionally, change the information.

7 Click **NEXT>>**. A message displays explaining what is created during the installation.

8 Click **CONFIRM**.

When the Anti-Spam application is installed, you can:

- Download and install the Junk Box; see Installing the Junk Store on page 291
- Configure the email threat categories; see Configuring Email Threat Categories on page 292.

Installing the Junk Store

Anti-Spam can create a Junk Store on your Microsoft Exchange Server. The Junk Store quarantines messages for end-user analysis and provides statistics. Log in to your Exchange system, then open a browser to log in to the management interface, and install the Junk Store.

NOTE: While SonicWall supports non-Exchange SMTP servers, such as Sendmail and Lotus Domino, it is not required to install the Junk Store on one of these servers. Similar to the SonicWall Email Security product, the CASS 2.0 feature allows you to install the Junk Store on a stand-alone server.

To fully utilize the newest functionality available with CASS 2.0, SonicWall recommends installing Junk Store on a stand-alone server.

To install the Junk Store:

1 Log in to your Exchange system.

2 Open a web browser.

IMPORTANT: To download and install the SonicWall Junk Store application, you need the following on the system where you will install the Junk Store application:

- Internet Explorer 6 or above
- Microsoft Exchange Server
- Email Downloader ActiveX component for IE

3 Log in to the SonicOS interface.

4 Navigate to the **MANAGE | Security Configuration > Anti-Spam > Base Setup** page.

5 Go to the **SonicWall Junk Store Installer** section.

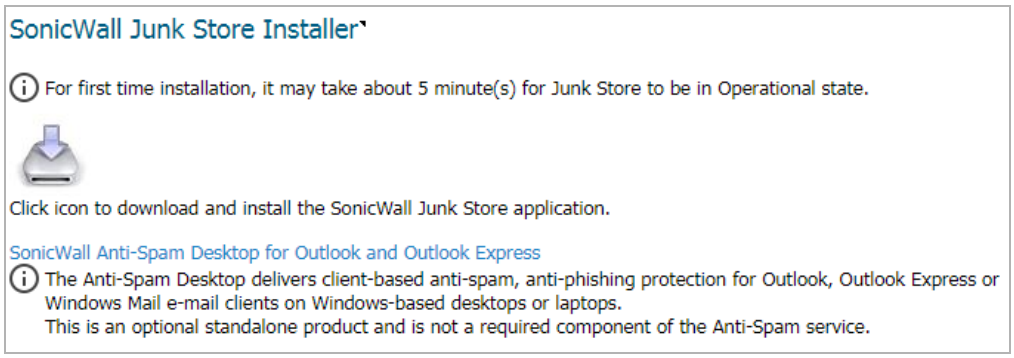

6 Click the **Junk Store Installer** icon to install the junk store on your Windows server.

NOTE: The first time the Junk Store application is installed, it takes about 5 - 15 minutes for the Junk Store to be operational.

7 If your browser warns you that the Web site is trying to load the SonicWall Email Security add-on:

a Click in the Information Bar.

b Select **Install ActiveX Control** in the pop-up menu. The Security Warning Screen displays.

8 Click **Install** to install the ActiveX Control.

9 On the **MANAGE | Security Configuration > Anti-Spam > Base Setup** page, click the **Junk Store Installer** icon again. A progress bar is displayed on the page.

10 The installer launches when it is fully downloaded.

NOTE: Migrating data to the Junk Store may take a long time to complete.

11 Navigate to the **MONITOR | Current Status > Anti-Spam Status** page and verify that the SonicWall Junk Store is **Operational**.

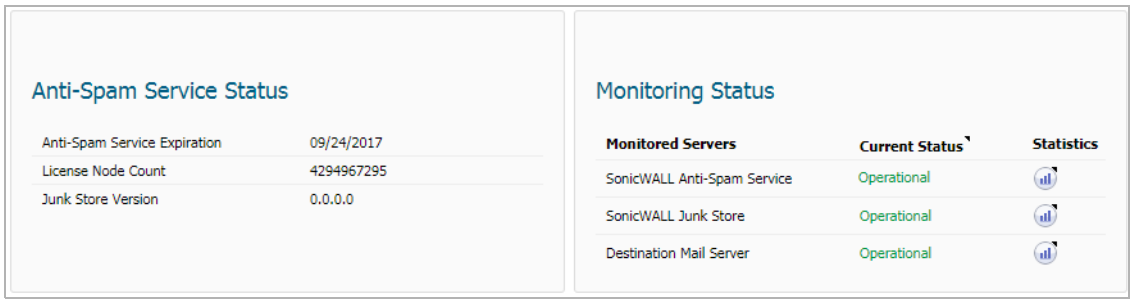

Configuring Email Threat Categories

After activating Anti-Spam, set your preferences. After these are configured, your email is filtered and sorted according to your configuration.

To set default settings for users' messages:

1 On the **MANAGE | Security Configuration > Anti-Spam > Base Setup** page, scroll to the **Email Threat Categories** section.

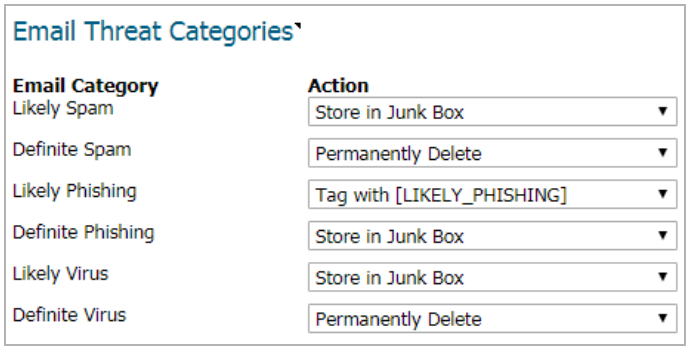

2 Choose default settings for messages that contain or may contain spam, phishing, and virus issues; see Email Threat Category Settings: Options for options available in the drop-down menus:

- **Likely Spam** (default: **Store in Junk Box**)
- **Definite Spam** (default: **Permanently Delete**)
- **Likely Phishing** (default: **Tag with [LIKELY_PHISHING]**)
- **Definite Phishing** (default: **Store in Junk Box**)
- **Likely Virus** (default: **Store in Junk Box**)
- **Definite Virus** (default: **Permanently Delete**)

Email Threat Category Settings: Options

Category	Action
Filtering off	Anti-Spam does not scan and filter any email for this threat category, so all the email messages are delivered to the recipients.
Tag With [*TAG***]**	The email is tagged with a term in the subject line: • **[LIKELY_SPAM]** • **[SPAM]** • **[LIKELY_PHISHING]** • **[PHISHING]** • **[LIKELY_VIRUS]** • **[VIRUS]** Selecting this option allows the user to have control of the email and can junk it if it is unwanted.
Store in Junk Box	The email message is stored in the Junk Box. It can be unjunked by users and administrators with appropriate permissions.
Permanently Delete	The email message is permanently deleted. CAUTION: **If you select this option, your organization risks losing wanted email.**

TIP: If you are using more than one domain, choose the Multiple Domains option and contact SonicWall or your SonicWall reseller for more information.

3 Click **ACCEPT**.

Configuring Access Lists

The two lists in the **User-defined Access Lists** section allow you to manage static allow and reject lists by designating which clients are allowed or denied connection to deliver email.

NOTE: Entry settings in these lists take precedence over GRID IP reputation check results.

Topics:

- Configuring the Access Lists on page 294
- Adding a Host to the Access Lists on page 294

Configuring the Access Lists

To configure the lists:

1 On the **MANAGE | Security Configuration > Anti-Spam > Base Setup** page, scroll to the **User-defined Access Lists** section.

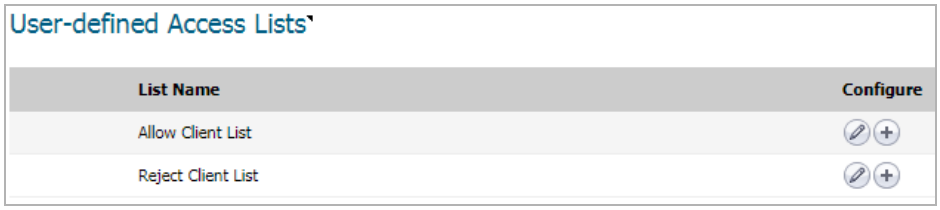

2 Click the **Edit** icon for the list, **Allow Client List** or **Reject Client List**, you want to configure. The **Allow/Reject Client List** dialog displays.

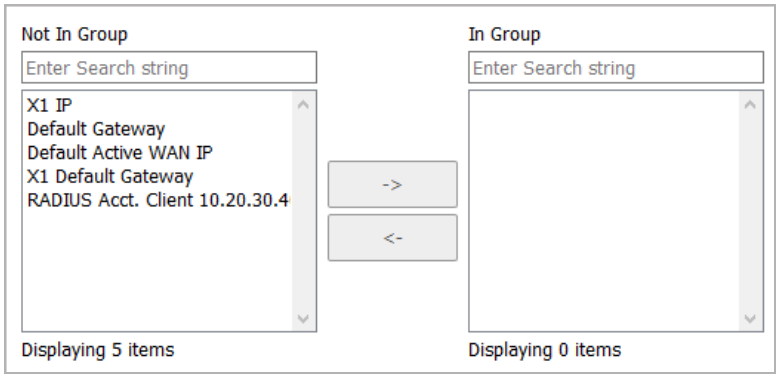

3 Select items from the **Not In Group** column you want to add to the **In Group** column.

4 Click the **Right Arrow** button.

To remove items from the **In Group** column:

a Select the item(s) from the **In Group** column.

b Click the **Left Arrow** button.

5 When finished, click the **OK** button.

Adding a Host to the Access Lists

To add a host to the lists:

1 Scroll to the **User-defined Access Lists** section.

2 Click the **Add Host** icon. The **Add Host to Allow/Reject List** dialog displays.

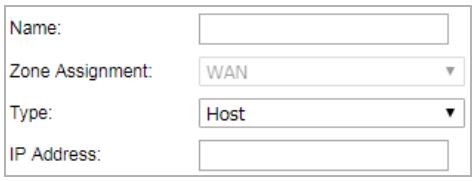

3 Enter a name for the host in the **Name** field.

4 Select the type of host from the **Type** drop-down menu. The following setting(s) change, depending on the host type selected.

5 If you selected:

- **Host** (default) – enter the IP address in the **IP Address** field.
- **Range** – enter the starting and ending IP addresses in the **Starting IP Address** and **Ending IP Address** fields.



- **FQDN** – enter the FQDN hostname in the **FQDN Hostname** field.

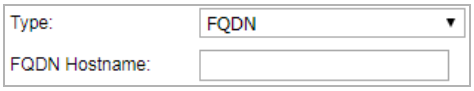

6 Click **OK**.

Configuring Advanced Options

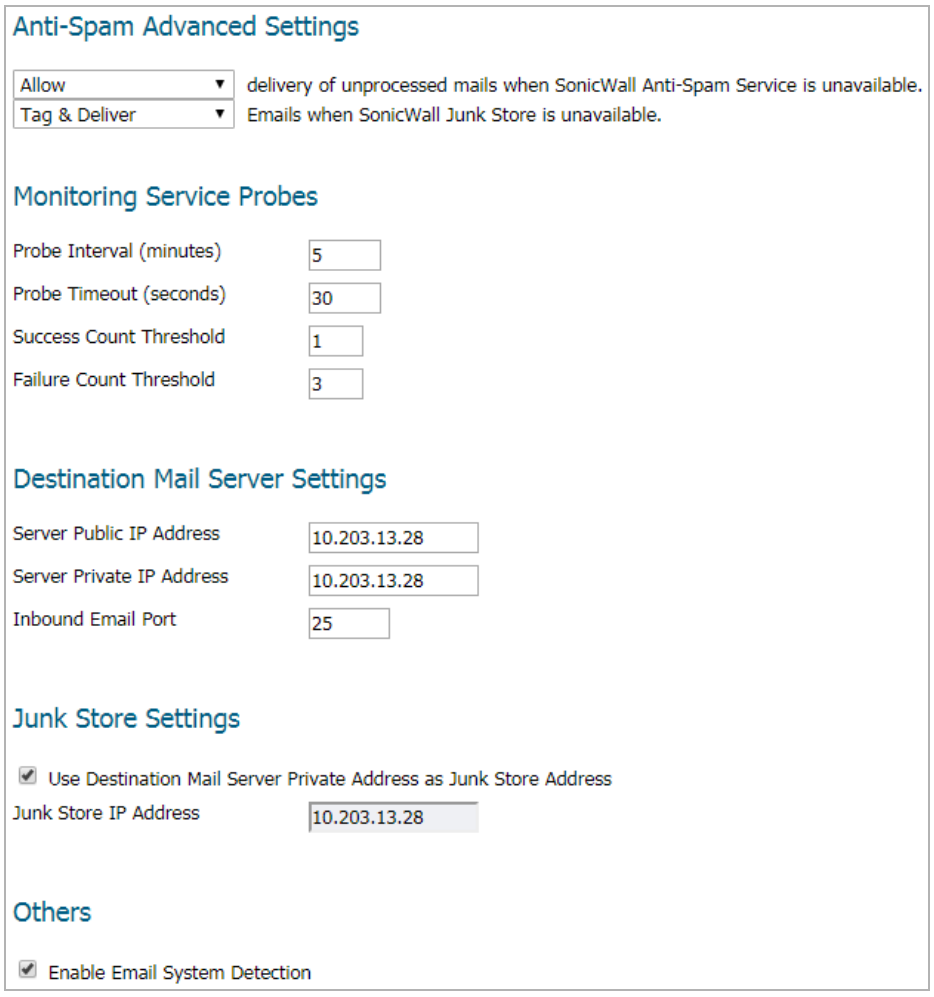

In the **Advanced Options** section, you can set the email options described in Anti-Spam > Base Settings: Advanced Options:.

Anti-Spam > Base Settings: Advanced Options

Setting type	Setting	Description
Anti-Spam Advanced Settings	**Allow/Reject delivery of unprocessed mails when SonicWall Anti-Spam Service is unavailable**	If the Anti-Spam service is not enabled or unavailable for some other reason, you can choose to let all unprocessed emails go through or to reject all unprocessed emails. Spam messages are delivered to users as well as good email. Choose from the drop-down menu: • **Allow** (default) • **Reject**
	Tag and Deliver/Delete Emails when SonicWall Junk Store is unavailable	If Junk Store cannot accept spam messages, you can choose to delete them or deliver them with cautionary subject lines such as `[Phishing] Please renew your account.` Choose from the drop-down menu: • **Tag & Deliver** (default) • **Delete**
Monitoring Service Probes	**Probe Interval (minutes)**	Set the timer frequency, in minutes, for probing Email Security components in the WAN and LAN networks. The minimum time is 1 minute, the maximum is 60 minutes, and the default is **5** minutes.
	Probe Timeout (seconds)	Set the time, in seconds, for the probe to wait for response from the target before flagging as failure. The minimum time is 30 seconds, the maximum is 300 seconds, and the default is **30** seconds.
	Success Count Threshold	Set the number of consecutive successful responses before declaring the entity as operational. The minimum number is 1 response, the maximum is 10 responses, and the default is **1** response.
	Failure Count Threshold	Set the number of consecutive successful responses before declaring the entity as unreachable. The minimum number is 1 response, the maximum is 10 responses, and the default is **3** response.
Destination Mail Server Settings	**Server Public IP Address**	The IP address of the server that is available for external connections. MTAs use this WAN IP address for SMTP connection. This number is populated by the address you specified when activating and installing Anti-Spam and Junk Store. You can change the address.
	Server Private IP Address	The IP address of the server for internal traffic. This is the internal mail server IP address behind the appliance. This number is populated automatically by the address you specified when activating and installing Anti-Spam and Junk Store. You can change the address.

Anti-Spam > Base Settings: Advanced Options

Setting type	Setting	Description
	Inbound Email Port	The TCP service port your appliance has open to receive inbound emails. The minimum is 0, the maximum is 65535, and the default is ***function generated***.
Junk Store Settings	**Use Destination Mail Server Private Address as Junk Store Address**	If the Junk Store is on the destination mail server, select the checkbox. The address is populated automatically by the address you specified when activating and installing Anti-Spam and Junk Store. You can change the address. This checkbox is selected by default, and the Junk Store IP Address field is dimmed. ***To change the address:*** 1 Uncheck the checkbox. The Junk Store IP Address field becomes available. 2 Enter the Junk Store IP address of where the server is located.
Others	**Enable Email Subsystem Detection**	Enables discover of available email system resources in the network. This checkbox is selected by default.

26

Configuring Anti-Spam Logging

NOTE: Anti-Spam is not supported on the SuperMassive series or the NS*a* 9250-9650 firewalls.

- Anti-Spam > Advanced
- Downloading System/Log Files
- Selecting the Amount and Level of Log Information

Anti-Spam > Advanced

The **Anti-Spam > Advanced Settings** page allows you to download log and system configuration files from your server as well as configure the log level.

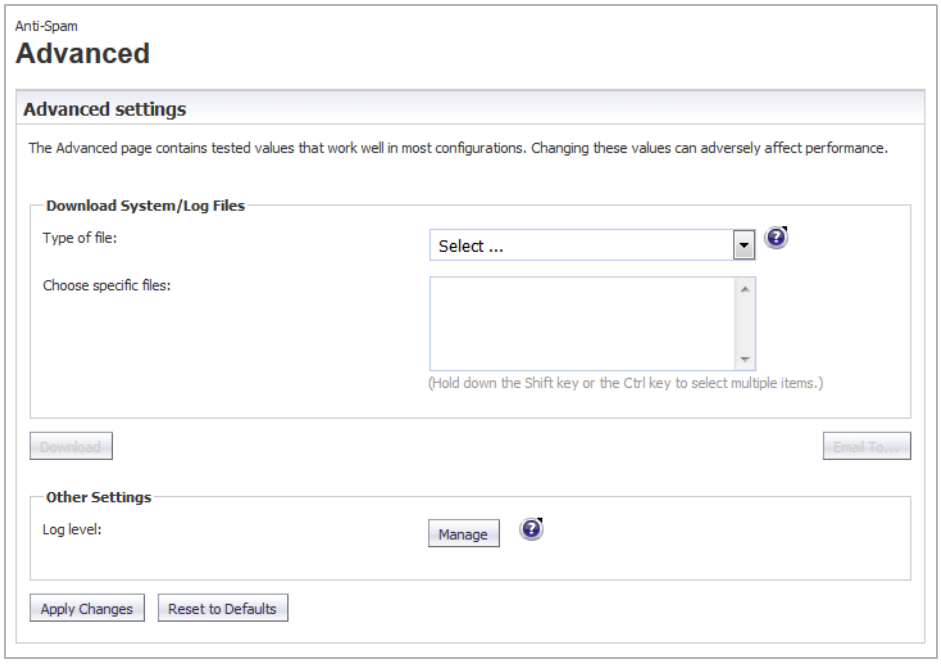

Topics:

- Downloading System/Log Files
- Selecting the Amount and Level of Log Information

Downloading System/Log Files

NOTE: Some log file names, such as those found in the `commonlogs` directory, contain a two-digit number such as `12.log`. The "12" indicates that the log is for the 12th day of the most recent month. Some log file names end with a digit, such as `MlfThumbUpdate_2.log`. The "2" indicates that this is an older log. The current log is `MlfThumbUpdate.log`. The next most recent log is `MlfThumbUpdate_0.log`, followed by `MlfThumbUpdate_1.log`, and so forth.

Most log data is in Greenwich Mean Time (GMT), not in the local time of the server the logs come from. This applies to the names of the log files as well.

To download log or system configuration files from your SonicWall Email Security server:

1 Navigate to the **Download System/Log Files** section of **Anti-Spam > Advanced Settings**.



2 Select the type of file to download from the **Type of file** drop-down menu. The **Choose specific files** list becomes populated with that type of file.

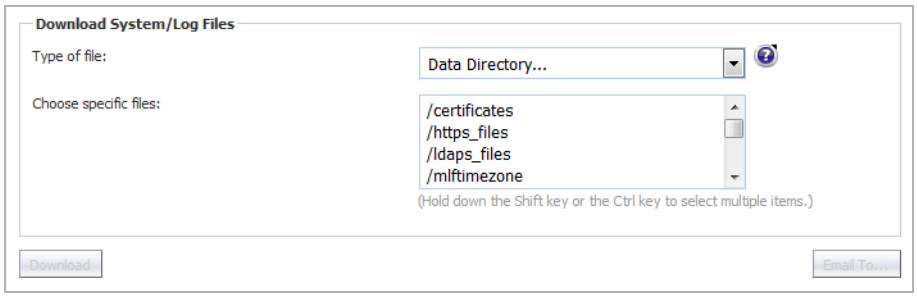

3 From the **Choose specific files** list, select one or more specific items. To select multiple files, hold down the Shift key or Ctrl key while selecting the files. The **Download** Download and **Email To...** Email To... buttons become active.

NOTE: The selected files are combined into a zip file.

4 Click either:

- **Download** button to download the file(s) to your local hard drive.
- **Email To...** button to email the file(s). the **Send To** dialog displays.

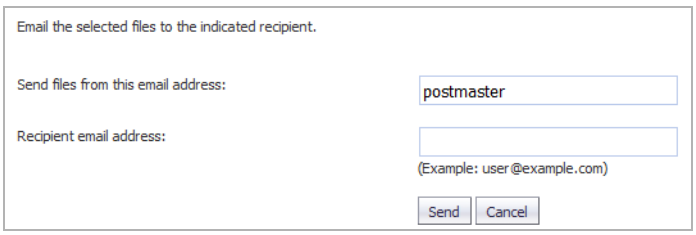

a) Enter the sender's email address in the **Send files from this email address** field. The default is **postmaster**.

b) Enter the recipient's email address in the **Recipient email address** field.

c) Click the **Send** button.

NOTE: Emailing very large files and directories can be problematic depending on the limitations of your email system.

Selecting the Amount and Level of Log Information

You can select the level and amount of system report information to be stored in your logs in the **Other Settings** section.

To configure the level and amount of log information:

1 Navigate to the **Other Settings** section of **Anti-Spam > Advanced Settings**.

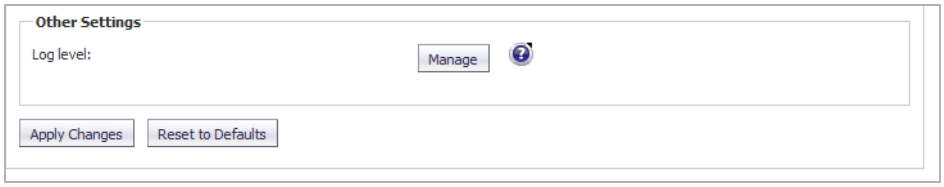

2 Click the **Manage** Manage button. The **Set Log Level** dialog displays.

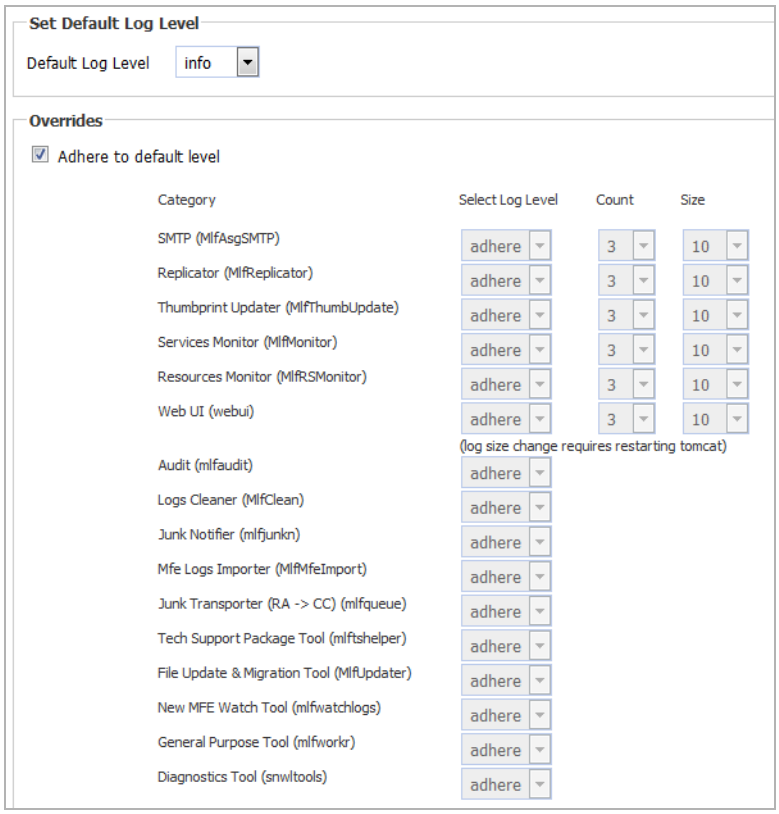

3 Select the default log level from the **Default Log Level** drop-down menu; levels are listed from lowest to highest:

NOTE: The higher the default log level, the more events are recorded. For example, the **info** level also records **trace** and **debug** levels.

- **trace** – lowest level
- **debug**
- **info** – default
- **warn**
- **error**
- **fatal** – highest level

All logs adhere to the default level set here unless specifically overridden.

4 To make changes to the logs in the **Overrides** section, deselect the **Adhere to default level** checkbox. All drop-down menus for all service categories become active.

5 To change the log level for specific services and subservices. from the **Select Log Level** drop-down menu for the service/subservice to be changed, select the desired log level. The levels are the same as for those in Step 3, plus the **adhere** option.

NOTE: The default log level for all service and subservice categories is **adhere**, that is, the log level set by the **Default Log Level** drop-down menu is used.

6 Optionally, select the number of log files to retain. By default, Junk Box keeps 3 log files for these services:

- SMTP
- Replicator
- Thumbprint Updater
- Services Monitor
- Resources Monitor
- Web UI

When a fourth log file is generated, the oldest log file is discarded, the second oldest becomes the oldest, and the third oldest becomes the second oldest.

a You can increase the number of logs kept for a service by selecting a number from the **Count** drop-down menu for that service:

- **3**
- **5**
- **6**
- **7**
- **8**
- **9**
- **10**

A lower number of logs saves disk space, but older data may not be available. A larger number of logs retains more data, but takes more disk space.

7 Optionally, select a size for the service logs (see Step 6) from the **Size** drop-down menus. The default size of each log is **10 Mb**.

You can increase the size of he logs, in 10 MB increments, from 10 Mb (default) to 100 Mb. A smaller log size saves disk space, but larger logs contain more data.

IMPORTANT: Changing the size of a log requires restarting the Tomcat server.

8 Click the **Apply Changes** button to save any changes made.

To return the logging level to default value:

1 Click the **Reset to Defaults** Reset to Defaults button.

27

Configuring the RBL Filter

NOTE: The Anti-Spam service is an advanced superset of the standard SonicOS RBL Filtering. When Anti-Spam is enabled, therefore, RBL Filtering is disabled automatically and a message displays with that information and a link to the **MANAGE | Security Configuration > Anti-Spam > Base Setup** page.

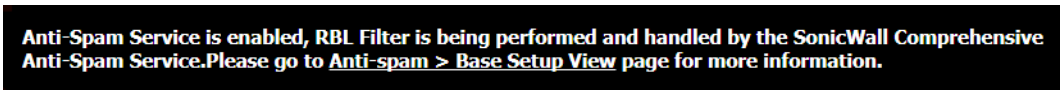

If Anti-Spam is not enabled, you can configure the settings on the **MANAGE | Security Configuration > Anti-Spam > Real-Time Black List Settings** page. All Anti-Spam and Junk Box pages, are unavailable, however.

NOTE: Anti-Spam is not supported on the SuperMassive series or the NS*a* 9250-9650 firewalls.

- Anti-spam > Real-Time Blacklist Filter on page 303
- About RBL Lists on page 303
- Enabling the RBL Filter on page 304
- Managing RBL Services on page 305
- User-Defined SMTP Server Lists on page 308
- Testing the Real-time Black List on page 310

Anti-spam > Real-Time Blacklist Filter

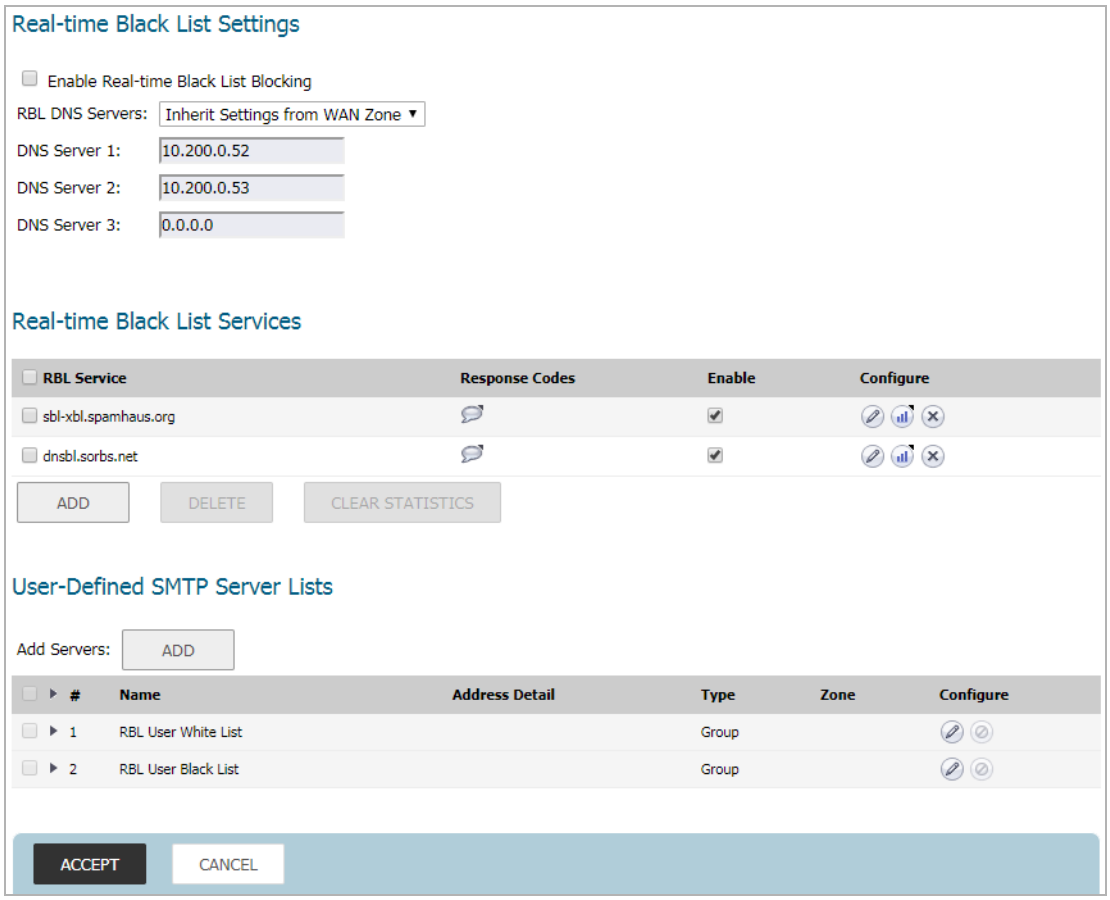

About RBL Lists

SMTP Real-Time Black List (RBL) is a mechanism for publishing the IP addresses of SMTP servers from which or through which spammers operate. There are a number of organizations that compile this information both for free: http://www.spamhaus.org, and for profit: https://ers.trendmicro.com/.

NOTE: SMTP RBL is an aggressive, spam-filtering technique that can be prone to false-positives because it is based on lists compiled from reported spam activity. The SonicOS implementation of SMTP RBL filtering provides a number of fine-tuning mechanisms to help ensure filtering accuracy.

RBL list providers publish their lists using DNS. Blacklisted IP addresses appear in the database of the list provider's DNS domain using inverted IP notation of the SMTP server in question as a prefix to the domain name. A response code from `127.0.0.2` to `127.0.0.11` indicates some type of undesirability:

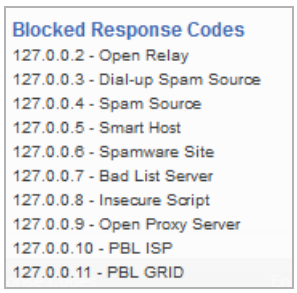

For example, if an SMTP server with IP address `1.2.3.4` has been blacklisted by RBL list provider `sbl-xbl.spamhaus.org`, then a DNS query to `4.3.2.1.sbl-xbl.spamhaus.org` provides a `127.0.0.4` response, indicating that the server is a known source of spam, and the connection is dropped.

NOTE: Most spam today is known to be sent from hijacked or zombie machines running a thin SMTP server implementation.Unlike legitimate SMTP servers, these zombie machines rarely attempt to retry failed delivery attempts. After the delivery attempt is blocked by RBL filter, no subsequent delivery attempts for that same piece of spam is made.

SonicOS Response to a Blacklist Query

The DNS responses are collected and cached. If any of the queries result in a blacklisted response, the server is filtered. Responses are cached using TTL values, and non-blacklisted responses are assigned a cache TTL of 2 hours. If the cache fills up, then cache entries are discarded in a FIFO (first-in-first-out) fashion.

The IP address check uses the cache to determine if a connection should be dropped. Initially, IP addresses are not in the cache, and a DNS request must be made. In this case, the IP address is assumed innocent until proven guilty, and the check results in the allowing of the connection. A DNS request is made and results are cached in a separate task. When subsequent packets from this IP address are checked, if the IP address is blacklisted, the connection is dropped.

Enabling the RBL Filter

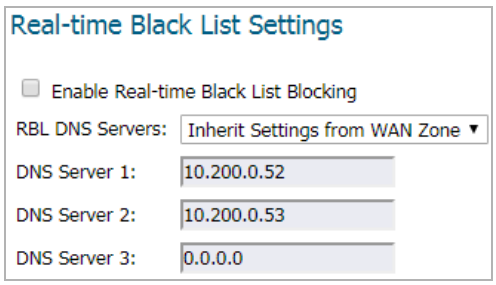

When Real-time Black List blocking is enabled, inbound connections from hosts on the WAN, or outbound connections to hosts on the WAN, are checked against each enabled RBL service with a DNS request to the DNS servers configured under RBL DNS Servers.

To enable the Real-time Black List filter:

1. Navigate to the **Security Configuration | Anti-Spam > Real-Time Blacklist Filter** page.
2. Select the **Enable Real-time Black List Blocking** checkbox.

3 Select the DNS Servers from the RBL DNS Servers drop-down menu:

- **Inherit Settings from WAN Zone** (default) — The DNS server(s) IP address(es) are displayed, but dimmed in the **DNS Server 1/2/3** fields.
- **Specify DNS Servers Manually** — The **DNS Server 1/2/3** fields become available.

    a) Enter one or more DNS server IP addresses in the **DNS Server 1/2/3** fields.

4 Click **ACCEPT**.

Managing RBL Services

You can add additional RBL services in the **Real-time Black List Services** section.

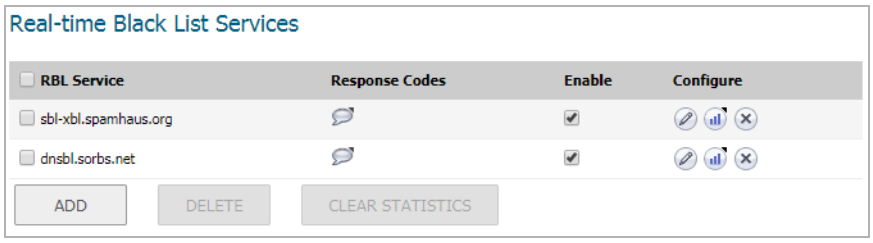

The **Real-time Black List Services** section displays information about and actions for the available RBL services:

- **RBL Service** – The name of the RBL service. Two are provided by SonicWall, but you can add others:
    - sbl-xbl.spamhaus.org – Spamhaus Project, which provides real-time anti-spam protection for Internet networks
    - dnsbl.sorbs.net – SORBS (Spam and Open Relay Blocking System), which provides access to its DNS-based Black List (DNSBL) database
- **Response Codes** – Mouse over the **Comment** icon to display a list of response codes. For information about response codes, see About RBL Lists.
- **Enable** – Select the checkbox to enable the RBL service. The checkboxes for the two provided services are selected by default.

    To disable an RBL service, unselect its checkbox. This does not delete the entry from the table, so you can enable the service in the future.
- **Configure** – Displays icons for various actions:
    - **Edit** icon – Displays the **Edit RBL Domain** dialog. See Editing an RBL Service.
    - **Statistics** icon – Displays information about connections blocked:

        

        To clear these statistics, click the Clear STatistics button.
    - **Delete** icon – Deletes the RBL service entry. See Deleting an RBL Service.

Topics:

- Clearing Statistics on page 306
- Adding an RBL Service on page 306
- Editing an RBL Service on page 307

- Deleting an RBL Service on page 308

Clearing Statistics

You can clear statistics kept for the Black List services.

To clear statistics:

1 Select a service by clicking its checkbox. To clear the statistics of all services, select the checkbox in the header next to **RBL Service**. The **CLEAR STATISTICS** button becomes active.

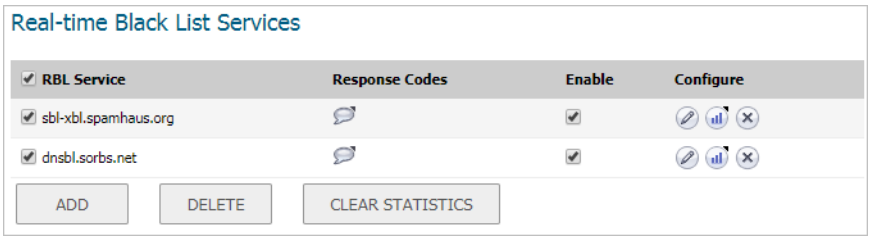

2 Click the **CLEAR STATISTICS** button.

Adding an RBL Service

To add an RBL service:

1 On the **Security Configuration | Anti-Spam > Real-Time Blacklist Filter** page, scroll to the **Real-Time Black List Services** section.

2 Click the **ADD** button. The **RBL Domain Settings** dialog displays.

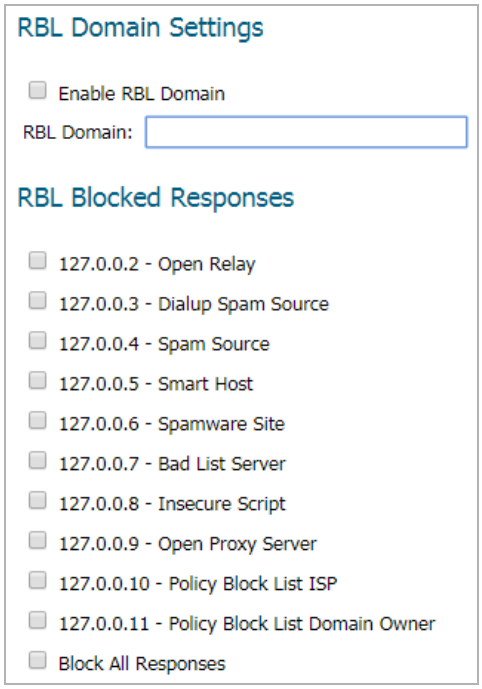

3 Specify the domain name of the RBL service to be queried in the **RBL Domain** field.

4 Enable the service for use by selecting the **Enable RBL Domain** checkbox.

5 Specify the expected response codes by selecting their checkboxes. Most RBL services list the responses they provide on their Web site, although selecting **Block All Responses** is generally acceptable.

> **TIP:** Selecting the **Block All Responses** checkbox selects the checkboxes for all the blocked responses. Deselecting the **Block All Responses** checkbox deselects the checkboxes of all the blocked responses.

6 Click **OK**. The RBL service is added to the **Real-Time Black List Services** table.

Editing an RBL Service

To edit an RBL Service:

1 On the **Security Configuration | Anti-Spam > Real-Time Blacklist Filter** page, scroll to the **Real-Time Black List Services** section.

2 Click the **Edit** icon associated with the RBL Service you want to change. The **Add RBL Domain** dialog displays.



3 Optionally, edit the domain name of the RBL service to be queried in the **RBL Domain** field.

> **TIP:** You can enable or disable an RBL service by selecting/deselecting its **Enable** checkbox in the **Real-time Black List Services** table.

4 Optionally, enable or disable the service for use by selecting/deselecting the **Enable RBL Domain** checkbox.

5 Optionally, select or deselect the expected response codes by selecting their checkboxes.

> **TIP:** Selecting the **Block All Responses** checkbox selects the checkboxes for all the blocked responses. Deselecting the **Block All Responses** checkbox deselects the checkboxes of all the blocked responses.

6 Click **OK**.

Deleting an RBL Service

To delete one RBL service:

1 Click the **Delete** icon for the service in the **Real-time Black List Services** table. A warning message displays:

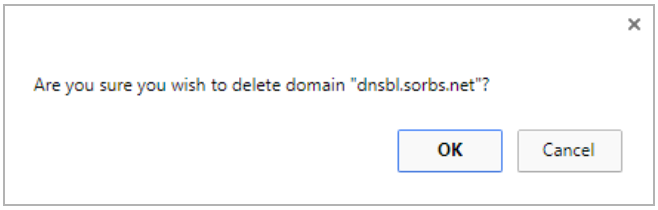

2 Click **OK**. The entry is deleted from the **Real-Time Black List Services** table.

To delete one or more RBL services:

1 Select the checkbox of one or more services in the **Real-time Black List Services** table. The **Delete** button becomes active.

2 Click the **DELETE** button. A warning message displays:

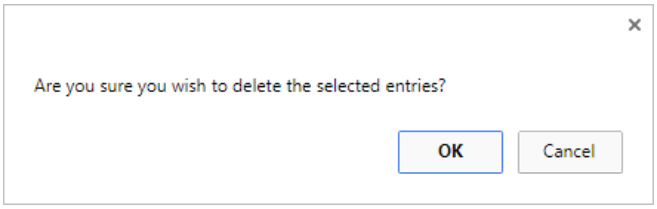

3 Click **OK**. The entry is deleted from the **Real-Time Black List Services** table.

User-Defined SMTP Server Lists

NOTE: You can modify, but not delete, the **RBL User White List** or the **RBL User Black List**.

The **User Defined SMTP Server Lists** section allows for Address Objects to be used to construct a white-list (explicit allow: **RBL User White List**) or black-list (explicit deny: **RBL User Black List**) of SMTP servers. Entries in these lists bypass the RBL querying procedure.

To ensure that you always receive SMTP connections from a partner site's SMTP server:

1 On the **Security Configuration | Anti-Spam > Real-Time Blacklist Filter** page, scroll to the **User-Defined SMTP Server Lists** section.

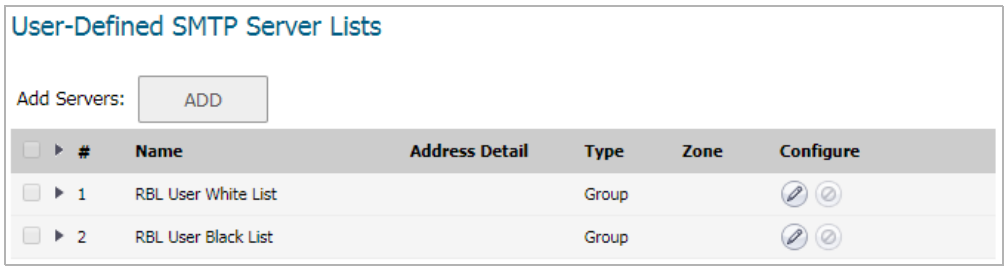

2 Create an Address Object for the server you want to add:

a Click the **ADD** button. The **Add Address Object** dialog displays.

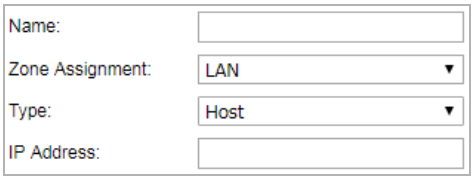

b Enter a friendly name for the server in the **Name** field.

c From the **Zone Assignment** drop-down menu, select the server's zone.

d From the Type drop-down menu, select the type of host from the **Type** drop-down menu. The following setting(s) change, depending on the host type selected.

e If you selected:

- **Host** (default) – Enter the IP address in the **IP Address** field.
- **Range** – Enter the starting and ending IP addresses in the **Starting IP Address** and **Ending IP Address** fields.



- **Network** – Enter the:

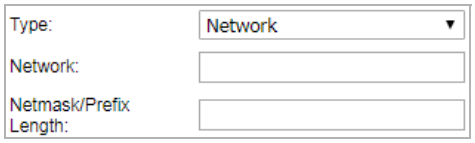

  - Network in the **Network** field.
  - Netmask in the **Netmask** field.
- **MAC:**

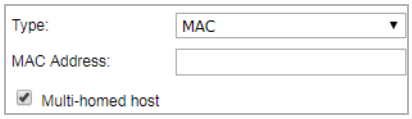

  - Enter the MAC address in the MAC Address field.
  - If the host is a multi-homed hose, select the **Multi-homed host** checkbox. Otherwise, deselect the checkbox. This checkbox is selected by default.
- **FQDN** – Enter the FQDN hostname in the **FQDN Hostname** field.



f Click **OK**.

3 Click the **Edit** icon in the **Configure** column of the **RBL User White List**. The **Edit Address Object Group** dialog displays.

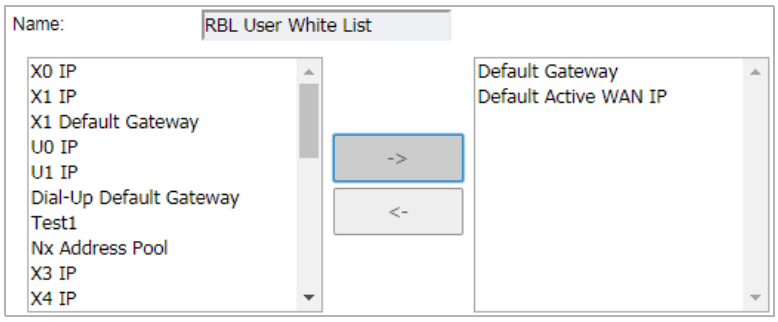

4 Select the address objects to be added from the left column. Multiple address objects can be selected at one time.

5 Click the **Right Arrow** button.

To delete an address object from the group, select the address object and click the **Left Arrow** button.

6 Click **OK**. The table is updated, and that server is always allowed to make SMTP exchanges.

Testing the Real-time Black List

The **INVESTIGATE | Tools | System Diagnostics page** also provides a **Real-time Black List Lookup** feature in the **Diagnostic Tools** section that allows for SMTP IP addresses (or RBL services or DNS servers) to be specifically tested. For information about this feature, see *SonicWall SonicOS 6.5 Investigate*.

For a list of known spam sources to use in testing, refer to: http://www.spamhaus.org/sbl/latest/.

28

Specifying Relay Domains

NOTE: Anti-Spam is not supported on the SuperMassive series or the NS*a* 9250-9650 firewalls.

- Anti-Spam > Relay Domains on page 311
- About Open Relay on page 312
- Listing Allowed Relay Domains on page 312

Anti-Spam > Relay Domains

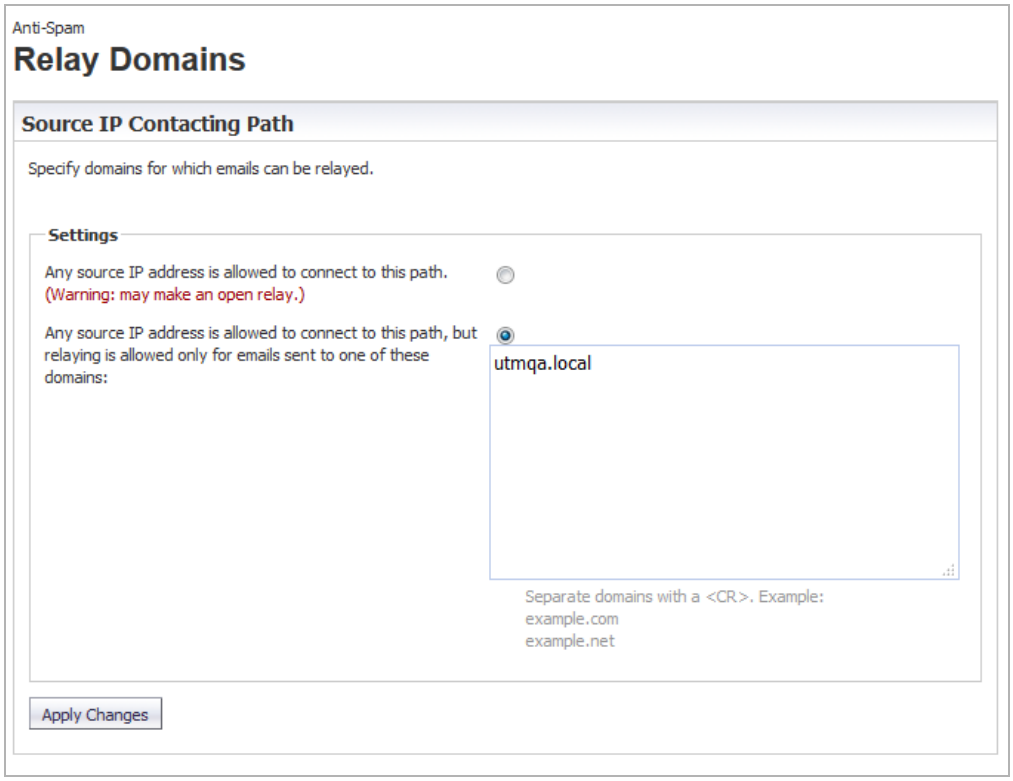

The **MANAGE | Security Configuration > Anti-Spam > Relay Domains** page allows you to list domains authorized for relaying email by CASS. Restricting domains that can relay emails avoids open-relay issues.

About Open Relay

An open relay is a SMTP server configured in such a way that it allows a third party to relay (send/receive email messages) that are neither from nor for local users. Such servers, therefore, are usually targets for spammers.

When CASS is configured as an open relay, the mail is relayed even if the mail is not destined to the recipient domain. When CASS is not configured as an open relay, it relays the emails that have one of the listed recipient domains; for domains not listed, the mails are rejected. Listing allowed relay domains avoid unnecessary relaying of emails even when mails are not destined to the user.

Listing Allowed Relay Domains

You can list all domains used for relay.

To list an authorized relay domain:

1 Navigate **MANAGE | Security Configuration > Anti-Spam > Relay Domains**.

2 Scroll to the **Settings** section

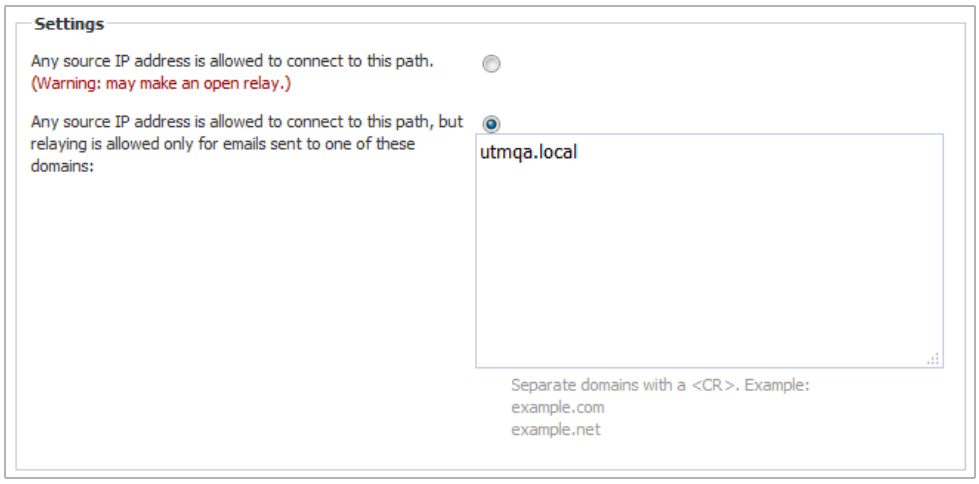

3 Select whether to restrict relay domains:

- **Any source IP address is allowed to connect to this path** – Allows any domain to relay messages. Go to Step 5.

CAUTION: Selecting this option may make a CASS open relay. Even if the mail is not destined to the recipient's domain, the mail is relayed, which could result in spamming

- **Any source IP address is allowed to connect to this path, but relaying is allowed only for emails sent to one of these domains** – Allows only listed domains to relay messages.

4 Enter the domain(s) allowed to relay messages in the field. Separate domains with a carriage return (<CR>).

5 Click **Apply Changes**.

29

Configuring Junk Box Settings

NOTE: Anti-Spam is not supported on the SuperMassive series or the NS*a* 9250-9650 firewalls.

- Anti-Spam > Junk Box Settings

Anti-Spam > Junk Box Settings

The **Anti-Spam > Junkbox Settings** page allows you to set the:

- Length of time that messages are stored in the Junk Box before being deleted.
- Number of Junk Box messages to be displayed per page.
- Action performed when a user unjunks a message.

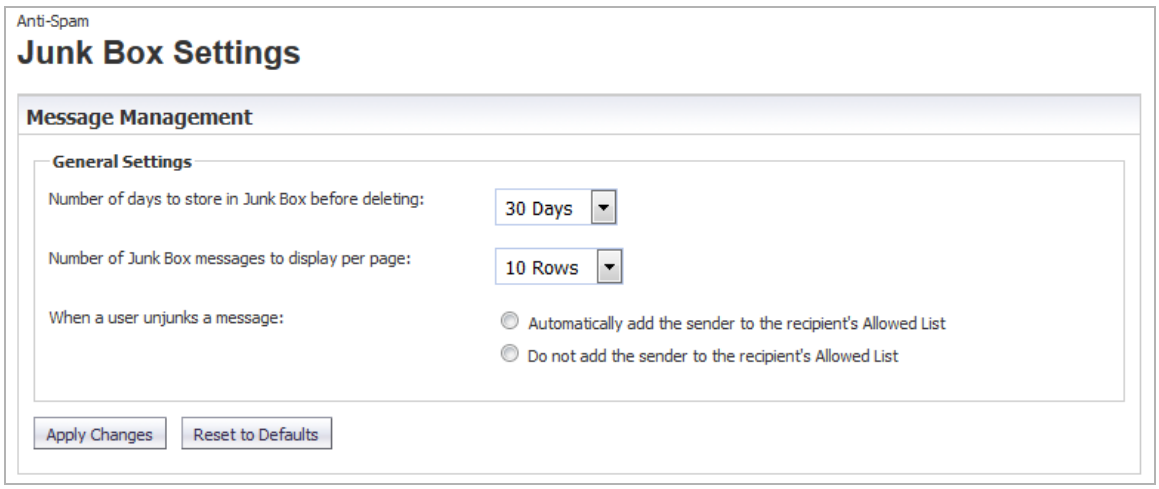

To perform message management:

1. In the **Message Management** section, select the number of days to retain junk mails before deleting them from the **Number of days to store in Junk Box before deleting** drop-down menu. The minimum is 1 Day, the maximum is 180 Days, and the default is **15 Days**.
2. Select the number of rows of messages to display in the **Messages Found** section on the **Inbound** view of the **INVESTIGATE | Logs | Anti-Spam Junkbox** page from the **Number of Junk Box messages to display per page** drop-down menu. The minimum is 10 Rows, the maximum is 400 Rows, and the default is **400 Rows**.
3. Select whether an unjunked sender is added to the recipient's Allowed List from **When a user unjunks a message**; neither option is selected by default:
    - **Automatically add the sender to the recipient's Allowed List**

- **Do not add the sender to the recipient's Allowed List**

4 Click **Apply Changes**.

To revert to default settings:

1 Click **Reset to Defaults**.

30

Managing the Junk Summary

NOTE: Anti-Spam is not supported on the SuperMassive series or the NS*a* 9250-9650 firewalls.

- Anti-Spam > Junkbox Summary on page 315
- Managing the Junk Summary on page 317
- Reverting to Defaults on page 319

Anti-Spam > Junkbox Summary

The Junk Store sends an email message to users listing all the messages placed in their Junk Summary. The **Anti-Spam > Junk Box Summary** page allows you to set up the Junk Summary for users.

To configure the types of messages that are logged, there is a link to the **Anti-Spam > Advanced Settings** page.

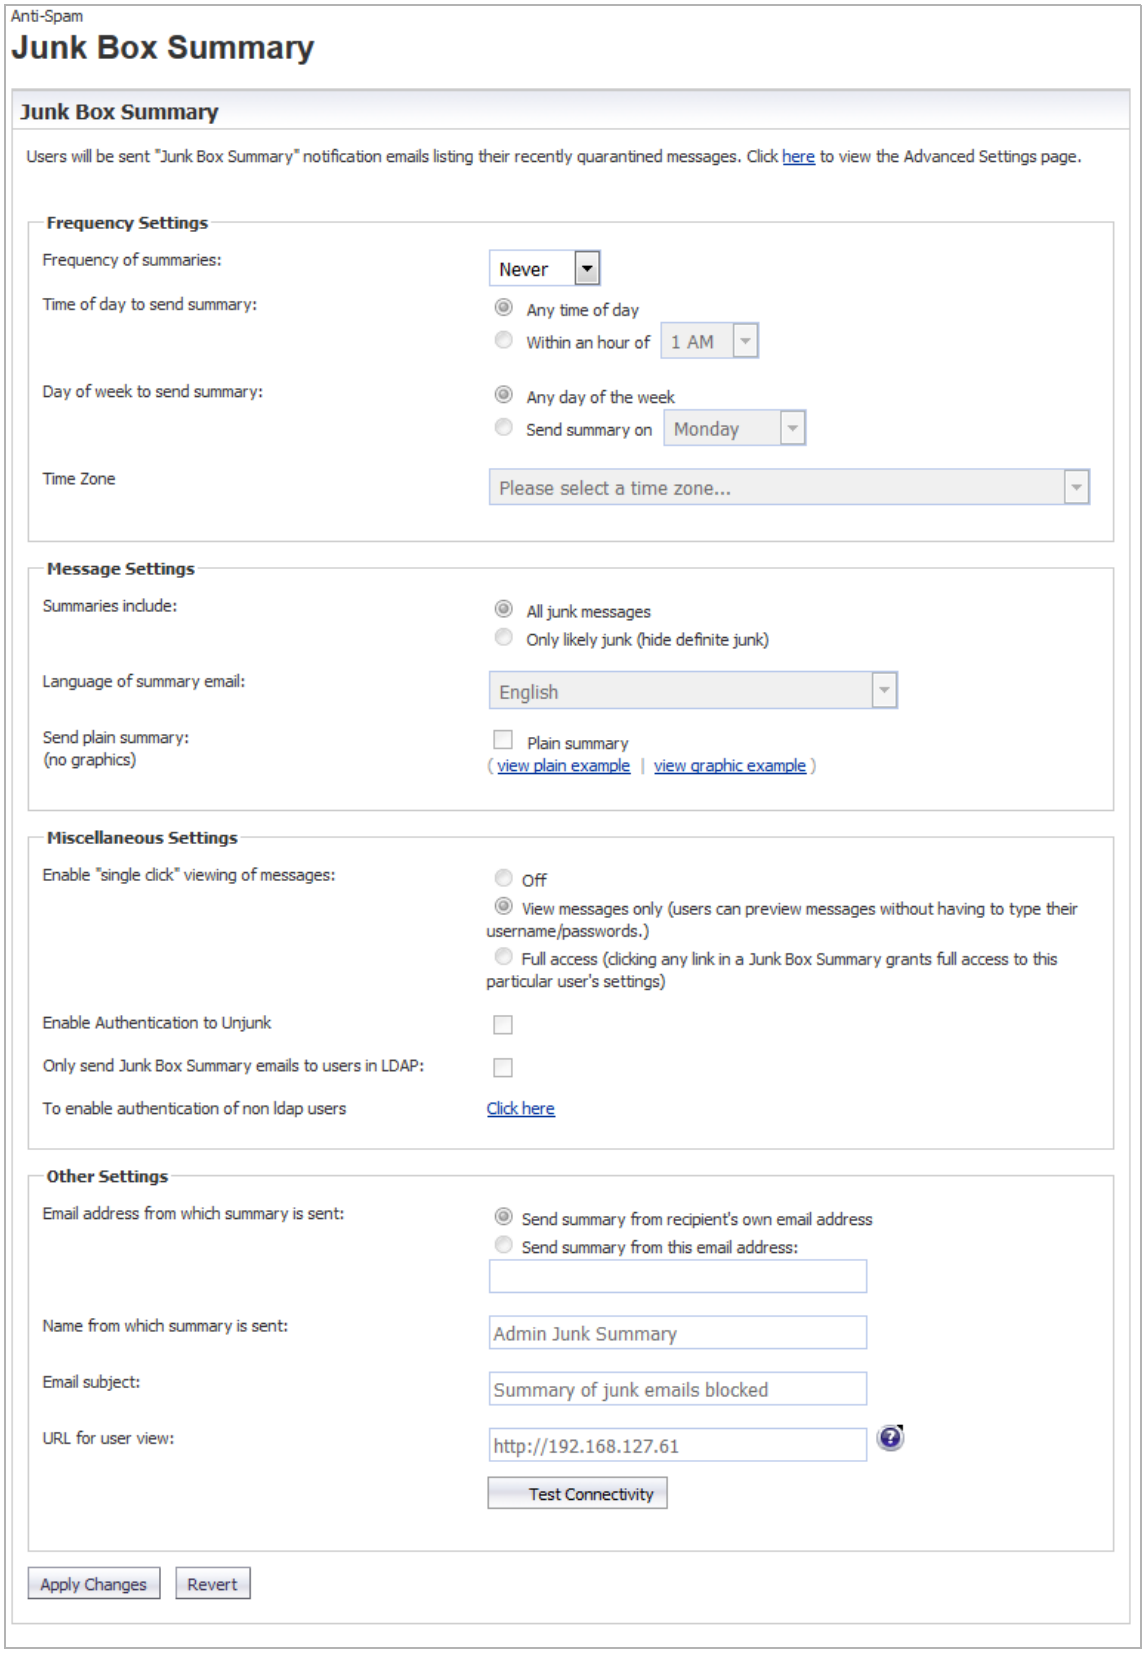

The **Anti-Spam > Junkbox Summary Settings** page allows you to set these options:

- **Frequency Settings** – Set the frequency and time Junk Box summaries are sent to you.
- **Message Settings** – Configure what is included in the summary, the language, and whether the summary contains graphics.
- **Miscellaneous Settings** – Set options such as single-click viewing of messages and authentication.
- **Other Settings** – Set options such as sender of summary, email subject, and URL for users.

Managing the Junk Summary

To manage the junk summary:

1 In the **Frequency Settings** section of the **Junkbox Summary Settings** page, select how often summaries are sent to you from the **Frequency of Summaries** drop-down menu.

Minimum frequency is **14 Days**, maximum is **1 Hour**, the default is **1 Day**. To prevent summaries from being sent to you, select **Never**.

2 Select from the **Time of day to send summary** options to customize the time your users receive email notifications.

NOTE: Individual users can override this setting.

- **Any time of day** (default)
- **Within an hour of** – select a time of day from the drop-down menu; the default is **12 AM**

3 If you selected **7 Days** or **14 Days** from the **Frequency of summaries** drop-down menu, the **Day of week to send summary** options become available. To customize the date your users receive email notifications select either:

NOTE: Individual users can override this setting.

- **Any day of the week** (default)
- **Send summary on** – select a day of the week from the drop-down menu; the default is **Monday**

4 Optionally, from the **Time Zone** drop-down menu, select the Greenwich Mean Time (GMT) to be used in determining the frequency.

5 In the **Message Settings** section, select what to include in the message summary from the **Summaries include** options:

- **All Junk Messages** (default)
- **Likely Junk Only (hide definite junk)**

6 Optionally, select a language for the emails from the **Language of summary emails** drop-down menu.

7 For **Send plain summary (no graphics)**, select whether the summary does not contain graphics by clicking the **Plain summary** checkbox. By default, graphics are included in the summary.

a To see an example for either version, click the appropriate link:

- **view plain example**

Junk Box Summary for: biz@example.com

In the past 24 hours, your organization has received 8040 Junk emails and 1122 Good emails.

Junk Emails Blocked: 24
The emails listed below have been placed in your personal Junk Box since your last Junk Box Summary and will be deleted after 90 days. To receive any of these messages, click Unjunk. The message will be delivered to your inbox.

Junk Box Summary

[Unjunk] [View] johnn@180solutions.com Re: 180 Advertising
[Unjunk] [View] dmcswzzain@hotmail.com -*- YES, Earn a Doctors income wi...
[Unjunk] [View] support@ebay.com Win Free Stuff
[Unjunk] [View] spammer@corp.net Take Some Viagra, its Cheap
▪
▪
▪
[Unjunk] [View] warning@alertsPC.com *!Alert. Read this. Click on buttons or BOOM
[Unjunk] [View] 31331@haxor.i.ua 133t H@x0r eZ xP10ts
[Unjunk] [View] ez@speller.com Learn to read words like a Pro
[Unjunk] [View] biggy@fat-guru.com Secret strategies of staying unemployed and fat
[Unjunk] [View] opportunity@yesyoucan.com Crop dusting jobs for Arab Americans

---

Junk blocking by SonicWALL, Inc.

- **view graphic example**

SONICWALL **Junk Box Summary**
for biz@example.com

Junk Emails Blocked: 8

The emails listed below have been placed in your personal Junk Box since your last Junk Box Summary and will be deleted after 90 days.
To receive any of these messages, click Unjunk. The message will be delivered to your inbox.

Email sent to: biz@example.com Visit Junk Box

	From	Subject	Threat
Unjunk | View	support@ebay.com	Official notice to biz@mailfrontier.com from Ebay Inc.	Phishing
Unjunk | View	dmcswzzain@hotmail.com	-*- YES, Earn a Doctors income wi...	Spam
Unjunk | View	spammer@corp.net	Win Free Stuff	Spam
Unjunk | View	jlef@mb12.com	Take Some Viagra, its Cheap	Spam
Unjunk | View	sally@getitup.com	Enlarge another body part	Spam
Unjunk | View	edd@aled.net	Nigerian Prince wants your PIN number	Spam
Unjunk | View	aber@ls.ua	Morgage rates that are really just ok	Spam
Unjunk | View	savenow@yahts.com	95% off of our Yahts	Spam

Anti-Spam Settings
Manage Allowed/Blocked lists

Spam Management Settings
Change frequency/timing of your Junk Box Summaries
Download anti-spam applications

To manage your quarantined emails, use your standard username and password to login here:
http://mtrose.corp.example.com

Junk blocking by SonicWALL, Inc.

b Close the window.

8 In the **Miscellaneous Settings** section, choose how email junkbox summary notifications are viewed from the **Enable "single click" view of messages** options:

- **Off**
- **View messages only (user can preview messages without having to type their username/passwords.)** (default)
- **Full access (clicking any link in a Junk Box Summary grants full access to the particular user's settings)**

9 To allow your users to authenticate to unjunk email messages, select the **Enable Authentication to Unjunk checkbox**. This option is not selected by default.

10 To limit junk box summaries notifications to users in LDAP, select the **Only send Junk Box Summary emails to users in LDAP** checkbox.

11 To enable authentication of non-LDAP users, click the **To enable authentication of non ldap users Click here** link. The **Anti-Spam > Users** page displays; for more information about managing users, see Managing the Junk Summary.

12 In the **Other Settings** section, choose how the summary is to be sent by selecting an option from **Email address from which summary is sent**:

- **Send summary from recipient's own email address** (default)
- **Send summary from this email address**

    a) Enter an email address in the field

13 In the **Name from which summary is sent** field, enter the name to be displayed in the user's email for the summary emails. The default name is **Admin Junk Summary**.

14 In the **Email subject** field, enter the subject line for the Junk Box Summary email. The default is **Summary of junk emails blocked**.

15 The **URL for user view** field is filled in automatically based on your server configuration. It is the basis for all the links in the Junk Box Summary email. If this setting is configured, each user Junk Box Summary emails listing that user's received email threats are sent.

Junk Box Summary emails contain URLs to:

- View quarantined emails.
- Unjunk quarantined emails; users unjunk items in the Junk Box summary email by clicking links in the email.
- Log in to the Junk Box.

IMPORTANT: If you change this URL, to ensure connectivity, test the link if you make any changes by clicking the **Test Connectivity** 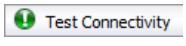 button. If the test fails, ensure the URL is correct.

16 Click the **Apply Changes** button.

Reverting to Defaults

You can revert all custom settings to default settings at any time.

To revert to default settings:

1 Click the **Revert** button.

31

Configuring the Junk Box View

NOTE: Anti-Spam is not supported on the SuperMassive series or the NS*a* 9250-9650 firewalls.

- Anti-Spam > Junk Box
- About the Junk Box Tabs
- Searching the Messages
- Managing Messages in the Junk Store

Anti-Spam > Junk Box

On the **INVESTIGATE | Logs > Anti-Spam > Junkbox** page, you can view, search, and manage all email messages that are currently in the Junk Store on the Exchange or SMTP server.

NOTE: This page is only available if the Junk Store is installed.

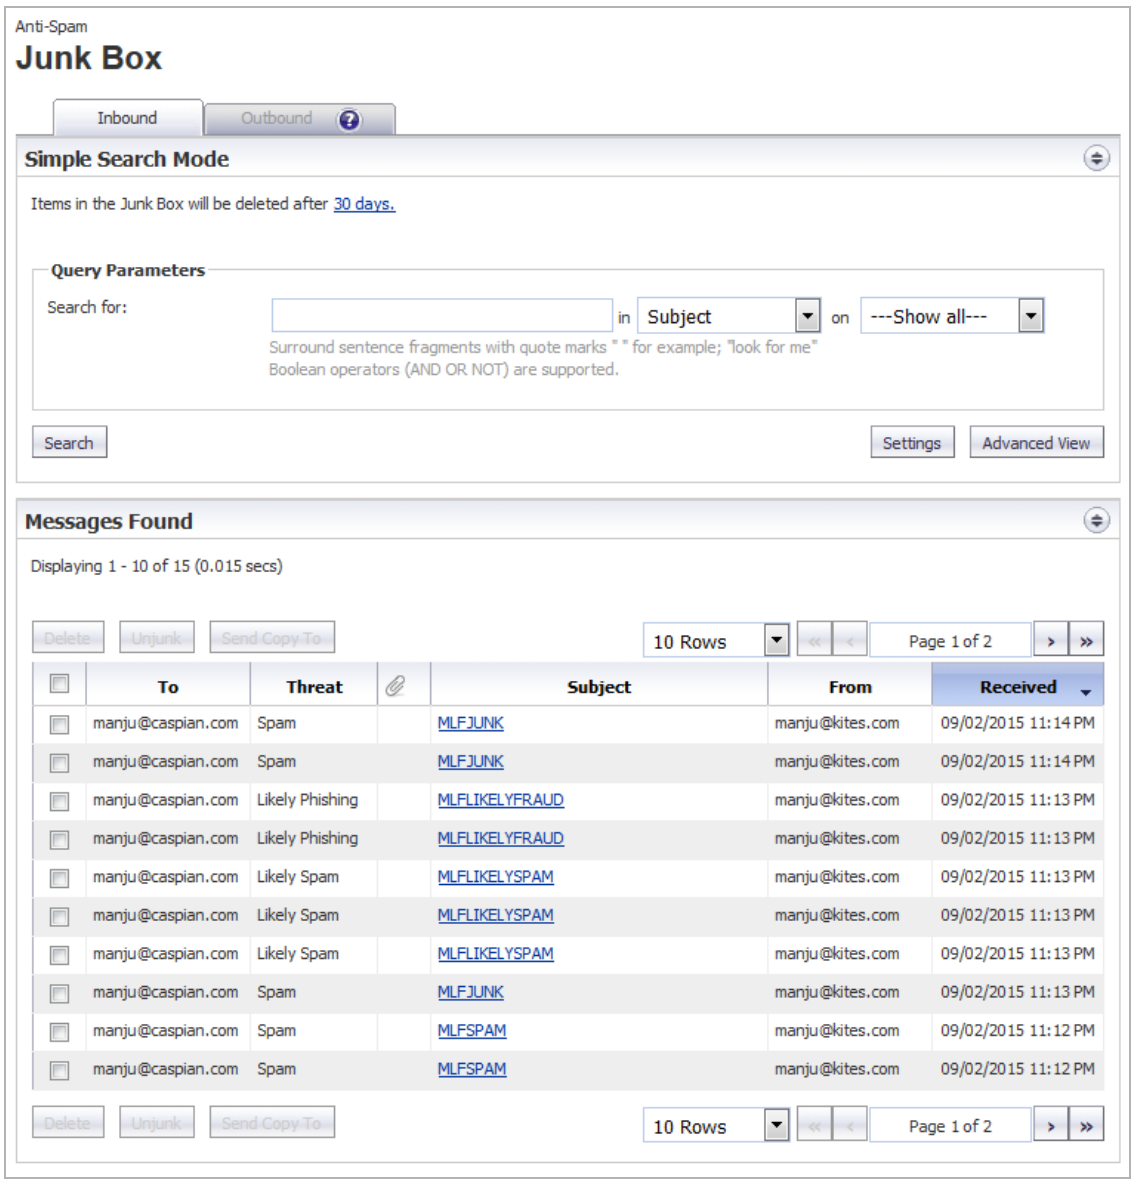

About the Junk Box Tabs

The **INVESTIGATE | Logs | Anti-Spam Junkbox** page contains two tabs:

- **Inbound**, which lists only inbound messages
- **Outbound**, which lists only outbound messages

NOTE: If you cannot view the **Outbound** view, you must upgrade your Junk Store license. If you click on the **Question Mark** icon, this message is displayed:



The function and display of the two tabs are the same. Each view contains two sections:

- **Simple/Advanced Search Mode**
- **Messages Found**

You can collapse or expand either section by clicking its **Expand/Collapse** icon.

In the **Simple Search Mode** section are two links to other pages:

- To change the duration junk mail is held before deletion, click the link at the end of **Items in the Junk Box will be deleted after** at the top of the section.
- To display the **Anti-Spam > Junkbox Settings** page, click the **Settings** button at the bottom of the section.

Information Displayed in the Messages Found Table

The **Messages Found** table displays this information about the quarantined messages:

Information about quarantined messages

This column	Contains or indicates
Checkbox icon	Checkbox for each item in the table. Clicking the **Checkbox** icon in the heading selects all items in the table.
To	Recipient's email address.
Threat	Type of threat the email poses; for more information about threat categories, see Email Threat Category Settings: Options in Configuring Email Threat Categories.
Paperclip icon	Email has attachments.
Subject	Subject line of the email.
From	Sender's email address.
Received	Date the email was sent.

Use the buttons at the top and bottom of the **Messages Found** table to perform the following Junk Store management tasks (see Message Table Buttons) on the **INVESTIGATE | Logs | Anti-Spam Junkbox** page:

Message Table Buttons

Button	Function
Delete	Permanently delete the selected message(s) from the Junk Store; to delete all messages click the checkbox in the table heading
Unjunk	Remove the selected message(s) from the Junk Store and deliver them to the user(s) to whom they are addressed. The delivery time and date are set by the Exchange server when each message is delivered to the user mailbox.
Send Copy To	Keep the selected message(s) in the Junk Store and send a copy of it (them) to a user.

Searching the Messages

You can perform two types of searches on messages found in the Junk Store:

- Simple; see Performing a Simple Search
- Advanced; see Performing an Advanced Search

Performing a Simple Search

To search the Junk Store:

1 On the **INVESTIGATE | Logs | Anti-Spam Junkbox** page, select either the **Inbound** view or the **Outbound** view.

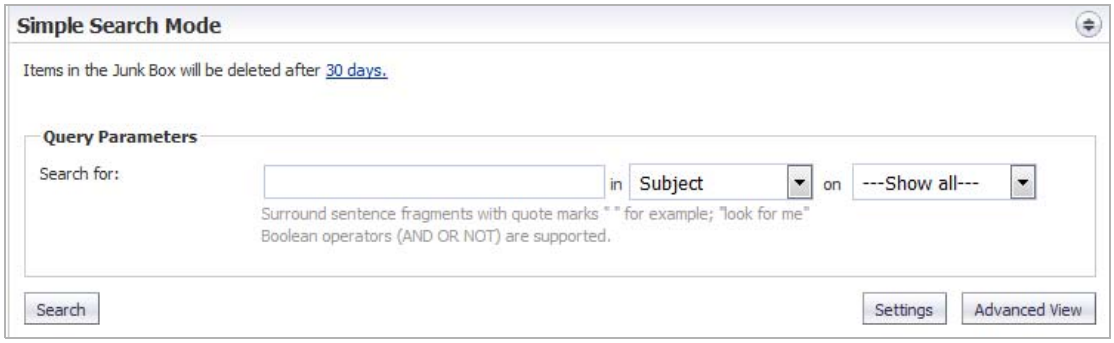

2 Type the text for which to search into the **Search for** field.

Surround sentence fragments with quotation marks (“). Boolean operators (AND, OR, NOT) can be used.

3 Select the desired email field in which to search from the **in** drop-down menu:

- **Subject** (default)
- **From**
- **To**
- **Unique Message ID**

4 From the **on** drop-down menu, select a date to search:

- **---Show all---** (default)
- **Today**

- A particular date; the number of dates vary, depending on the length of time junk messages are held

5 Click the **Search** button to perform the search.

The results are displayed in the **Messages Found** section of the page, and a message is displayed at the top. If the search is successful, the message contains the word, **Success!**, and the entire message is highlighted in green. If a search is not successful, it contains the word, **Warning!**, and the entire message is highlighted in yellow.

6 To return the **Messages Found** table to its original state:

a Delete the data from the **Search for** field.

b Click **Search**.

Performing an Advanced Search

1 On the **INVESTIGATE | Logs > Anti-Spam Junkbox** page, select either the **Inbound** view or the **Outbound** view.

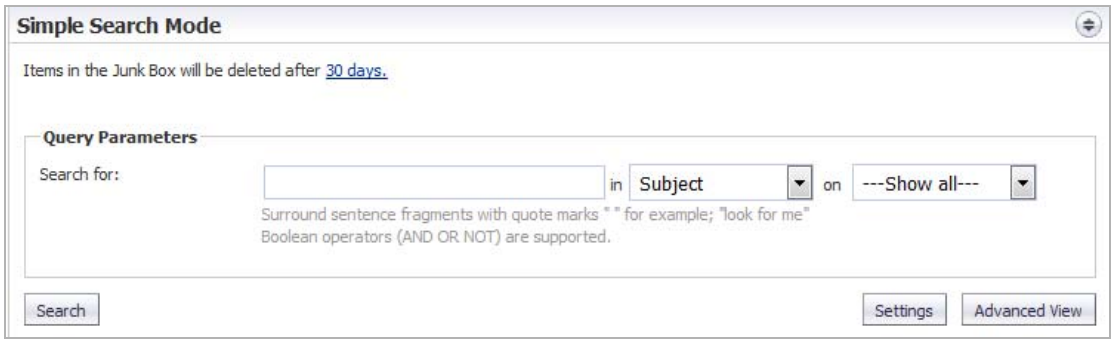

NOTE: To change the settings, click the link in the **Items in the Junk Box will be deleted after *nn* days** to display the **Anti-Spam > Junkbox Settings** page.

2 Click the **Advanced View** button. The **Simple Search Mode** expands to become the **Advanced Search Mode** section.

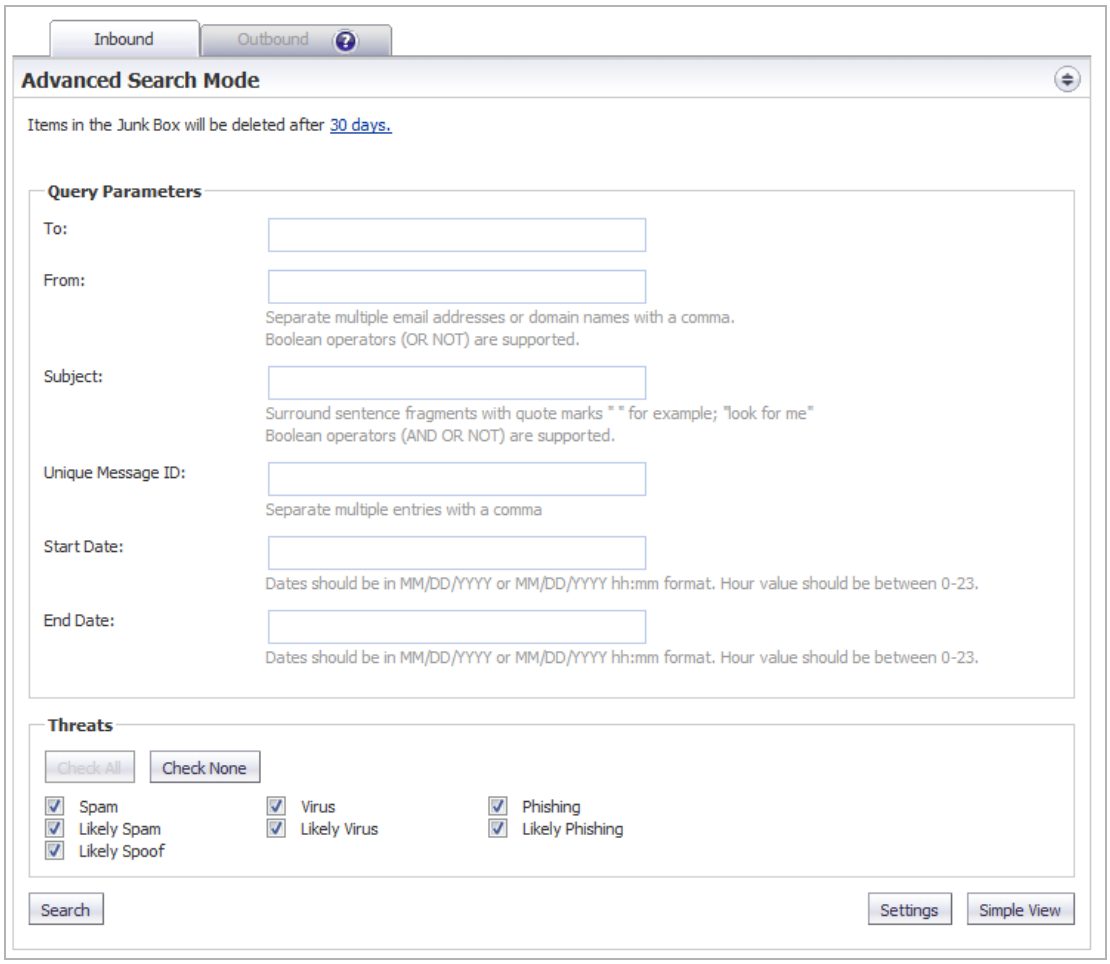

3 In the **Query Parameters** section, enter your search criteria in one or more of the **Query Parameter** fields:

Parameter	Query criteria
To	Recipient's email address.
From	Sender's email address. Separate multiple email addresses or domain names with a comma. Boolean operators OR and NOT are supported
Subject	Subject of the email. Enclose sentence fragments with quotation marks (“). Boolean operators AND, OR, and NOT are supported.
Unique Message ID	Unique message ID. Separate multiple entries with a comma.

Parameter	Query criteria
Start Date	First date to search. Enter dates in either format: • MM/DD/YYYY • MM/DD/YYYY hh:mm (Hour values should be between 0 and 23 [24-hour clock])
End Date	Last date to search. Enter dates in either format: • MM/DD/YYYY • MM/DD/YYYY hh:mm (Hour values should be between 0 and 23 [24-hour clock])

4 In the **Threats** section, specify the threat categories to search for. By default all categories are selected.

Deselect any category you do not want to include in the search by clicking its checkbox. To deselect all categories, click the **Check None** Check None button. All the categories become unchecked, the **Check All** Check All button becomes active, and the **Check None** button becomes dimmed.

Only messages belonging to one of the Email Threat Categories set to **Store in Junk Box** on the **Anti-Spam > Settings** page are included in the Junk Store. All categories, however, are listed on this page, whether any messages of that type are stored in the Junk Store.

NOTE: To change these settings, click the **Settings** button; the **Anti-Spam > Junkbox Settings** page displays.

5 Click the **Search** button to perform the search.

The results are displayed in the **Messages Found** section of the page, and a message is displayed at the top. If the search is successful, the message contains the word, **Success!**, and the entire message is highlighted in green. If a search is not successful, it contains the word, **Warning!**, and the entire message is highlighted in yellow.

6 To return to the **Simple View**, click the **Simple View** button.

7 To return the **Messages Found** table to its original state:

a Delete the data from the **Search for** field.

b Click **Search**.

Managing Messages in the Junk Store

TIP: If you are not searching the Junk Store, click the **Collapse** icon for the **Simple/Advanced Search Mode** section.

You can delete, unjunk, or send a copy of Junk Store messages.

To manage the Junk Store:

1 On the **INVESTIGATE | Logs | Anti-Spam Junkbox** page, scroll to the **Messages Found** table.

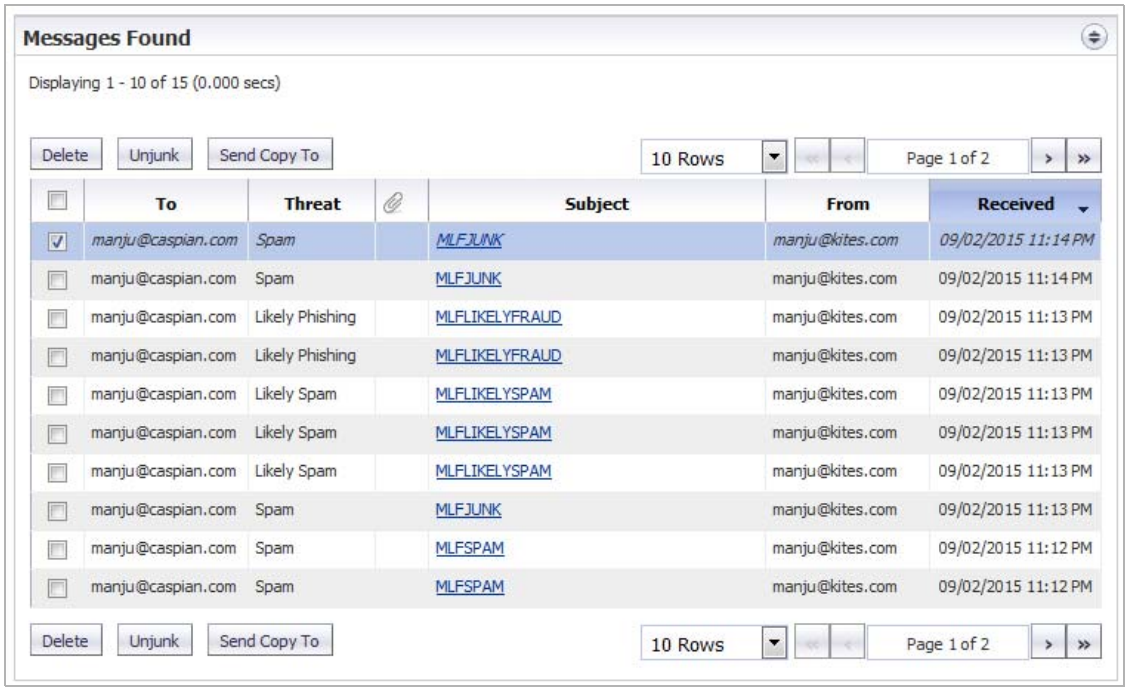

2 Select the checkbox for the message(s) that you want to manage.

TIP: To select all messages, select the checkbox in the table header. All checkboxes are selected.

3 Perform the management task(s):

- To permanently delete the selected messages from the Junk Store, click the **Delete** button.

  NOTE: Messages are deleted automatically after 30 days.

  The selected messages are deleted immediately — there is no confirmation dialog before the deletion. If the deletion is successful, a green notification is displayed at the top of the page. If the deletion fails, the notification is red.

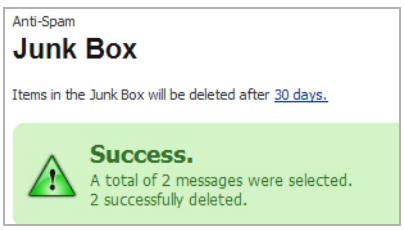

- To remove the selected messages from the Junk Store for delivery to the recipients, click the **Unjunk** button.

The selected messages are unjunked and sent immediately — there is no confirmation dialog before the action. If the action is successful, a green notification is displayed at the top of the page. If the action fails, the notification is red.

- To send a copy of the selected messages to a user, click the **Send Copy To** button. The **Send Copy To** dialog displays.

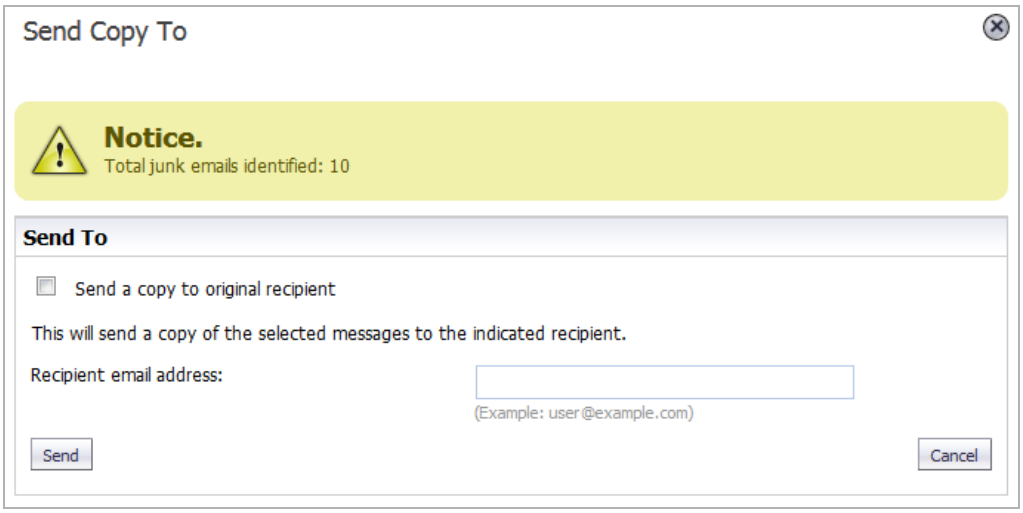

a) Do one of the following:

- Select **Send a copy to original recipient**.
- Type the email address into the **Recipient email address** field.

b) Click the **Send** Send button.

The selected message is sent immediately — there is no confirmation dialog before the action. If the action is successful, a green notification is displayed at the top of the page. If the action fails, the notification is red.

32

Configuring User-Visible Settings

NOTE: Anti-Spam is not supported on the SuperMassive series or the NS*a* 9250-9650 firewalls.

- Anti-Spam > User View Setup
- Configuring User View Setup
- Reverting to Default Settings

Anti-Spam > User View Setup

The **Anti-Spam > User View Setup** page allows you to select and configure which settings are visible for users.

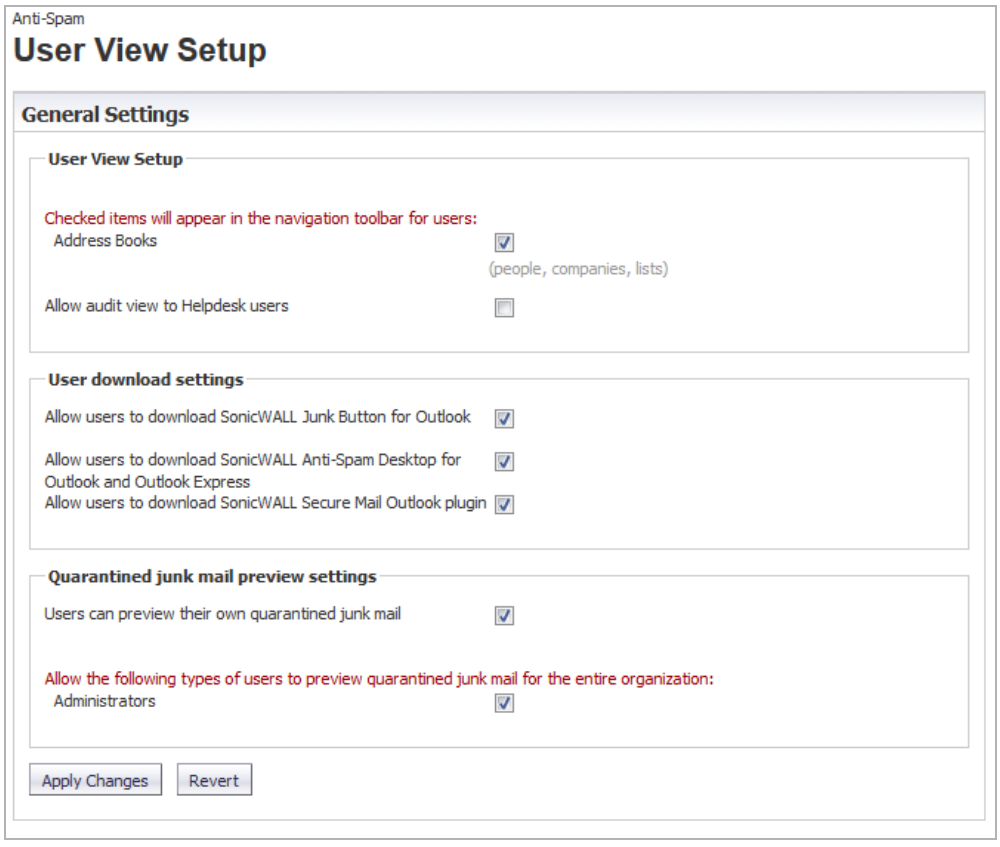

Configuring User View Setup

NOTE: Selected options appear in a user's navigation toolbar.

To configure what the user sees:

1. In the **User View Setup** section, to allow users to see their own Address Book (people, companies, and lists) in the navigation toolbar, select the **Address Books** checkbox. This option is selected by default.
2. To allow Helpdesk to view users' email problems, select the **Allow audit view to Helpdesk users** checkbox. This option is not selected by default.
3. In the **User download settings** section, to allow Outlook users to download the Junk Button, select the **Allow Users to download SonicWall Junk Button for Outlook** checkbox. This option is selected by default.
4. To allow Outlook and Outlook Express users to download the Anti-Spam Desktop, select the **Allow users to download SonicWall Anti-Spam Desktop for Outlook and Outlook Express** checkbox. This option is selected by default.
5. To allow Outlook users to download the Secure Mail plugin, select the **Allow users to download SonicWall Secure Mail Outlook plugin** checkbox. This option is selected by default.
6. In the **Quarantined junk mail preview settings** section, to allow users to preview their quarantined junk mail, select the **Users can preview their own quarantined junk mail** checkbox. This option is selected by default.
7. To allow Administrators to preview all quarantined junk mail for the entire organization, select the **Administrators** checkbox. This option is selected by default.

   NOTE: Administrators have access to preview all quarantined junk mail for the entire organization by default. To change this option, unselect the **Administrators** checkbox.

8. After all necessary changes have been made, click the **Apply Changes** button.

Reverting to Default Settings

You can change all settings back to factory defaults at any time.

To clear any changes made at any time and revert to the default settings:

1. Click the **Revert** button.

33

Configuring Corporate Allowed and Blocked Lists

NOTE: Anti-Spam is not supported on the SuperMassive series or the NS*a* 9250-9650 firewalls.

- Anti-Spam > Address Books on page 331
- About the Tabs on page 332
- Adding Items to the Allowed or Blocked List on page 333
- Deleting Items from the Allowed or Blocked List on page 334
- Importing Address Book Entries on page 334
- Exporting Address Book Entries on page 335
- Searching the Allowed and Blocked Lists on page 335

Anti-Spam > Address Books

The **Anti-Spam > Address Books** page allows you to configure the Allowed and Blocked lists for your organization. The lists are a combination of allowed and blocked senders from the organization's lists and lists provided by the firewall.

NOTE: The **Blocked** view only filters addresses by people, IPs, and companies, while the **Allowed** view filters addresses by people, companies, IPs, and lists.

If your lists are long, you can use a search function to display only desired table entries.

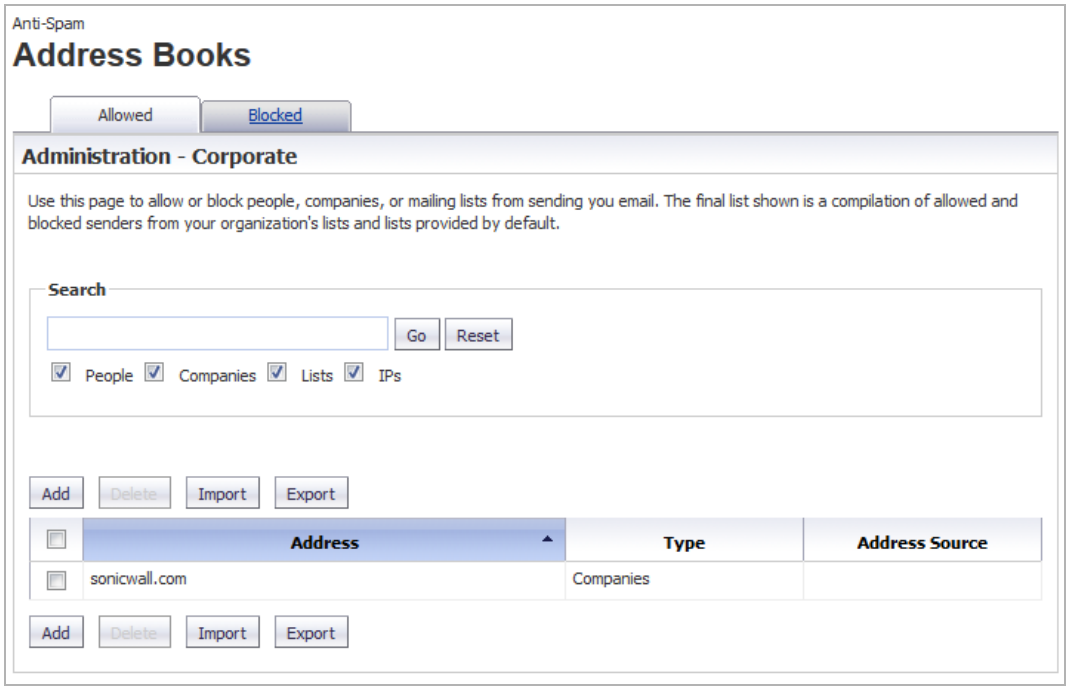

About the Tabs

The two tabs, **Allowed** and **Blocked**, are identical except the search categories for both pages are **People**, **Companies**, and **IPs** while the **Allowed** page also has **Lists**.

Topics:

- Allowed Lists
- Blocked Lists

Allowed Lists

The **Allowed** view enables you to permit people, companies, IP addresses, or lists to send mail to your organization. You can import address books to the Allowed list and export the Corporate Address Book to an Excel spreadsheet or text file.

Blocked Lists

NOTE: Senders added on the Corporate Blocked List by an Administrator are blocked automatically for all users and can only be deleted by an Administrator.

The **Blocked** view allows you to restrict people, companies, and IP addresses from sending mail to your organization. You can import address books to the Blocked list and export the Corporate Address Book to an Excel spreadsheet or text file.

Adding Items to the Allowed or Blocked List

To add an item to the Corporate Allowed/Blocked List:

1 Navigate to the appropriate view on **Anti-Spam > Address Books**.

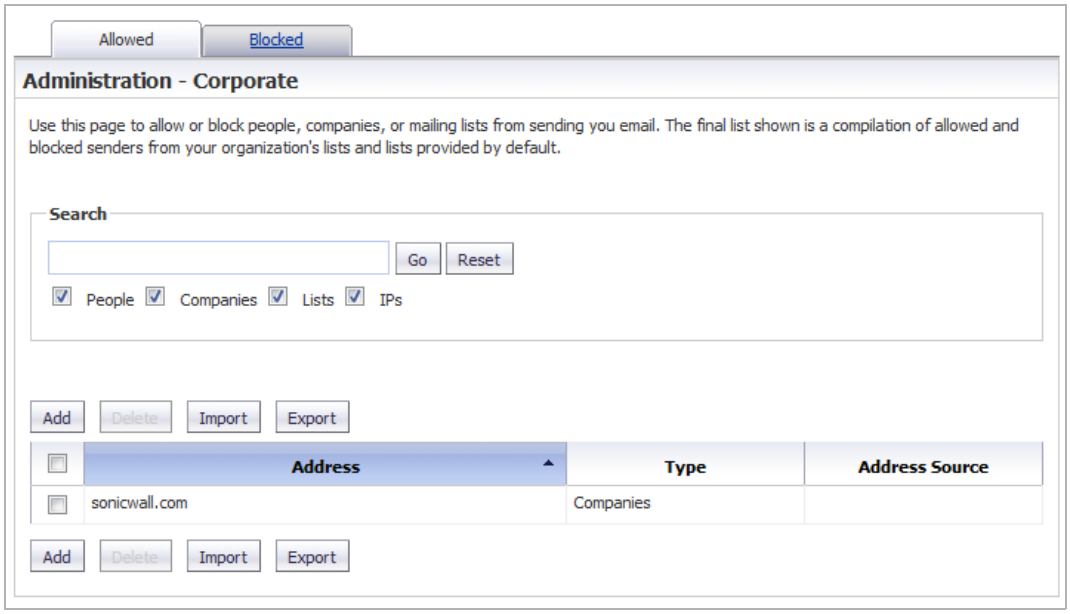

2 Click the **Add** button. The **Add Items Allowed List** dialog displays.

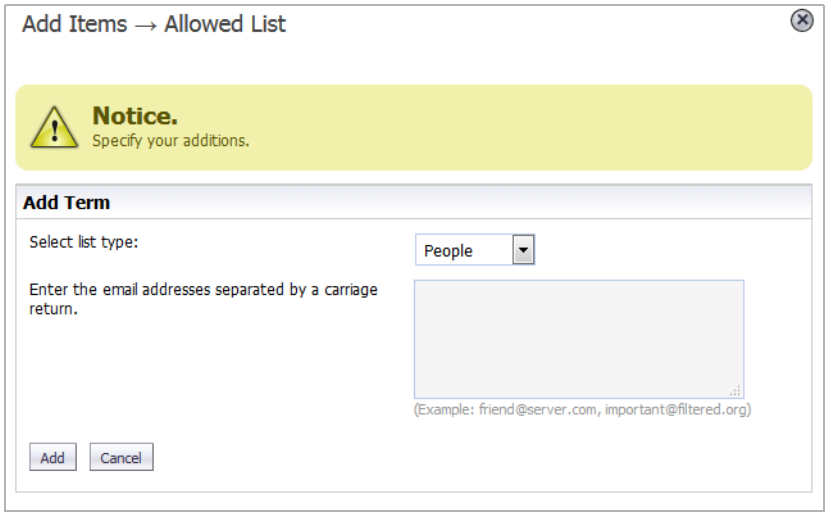

3 Select the type of list user from the **Select list type** drop-down menu:

- **People**
- **Companies**
- **Lists** (available only for the **Allowed** view)
- **IPs**

4 Enter the address(es)/domain(s) in the field. Depending on the list type selected, the field name changes:

- **People** – **Enter IP Addresses separated by a carriage return**
- **Companies** – **Enter the domains separated by a carriage return**
- **Lists** – **Enter the mailing lists separated by a carriage return**
- **IPs** – **Enter IP Addresses separated by a carriage return**

5 Click **Add** to finish. The address(es)/domain(s) are added to the **List** on the **Allowed**/**Blocked** view.

Deleting Items from the Allowed or Blocked List

To delete a sender from the Corporate Allowed/Blocked List:

1 Click the appropriate view.

2 Select the checkbox next to the email address(es) you wish to delete. The **Delete** button becomes active.

3 Click the **Delete** button. A success message appears confirming the deletion.

TIP: To delete all entries, click the checkbox in the table header.

Importing Address Book Entries

You can import entries from one or more address books.

To import address book entries:

1 Click the appropriate view.

2 Click the **Import** button. The **Import AddressBook** dialog displays.

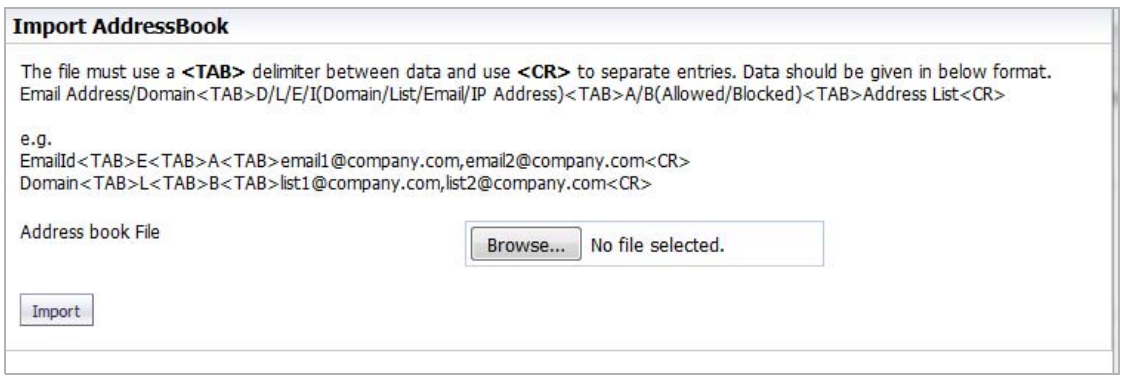

3 Click the **Browse** button. The Windows **File Upload** dialog displays.

4 Select the file to upload. It must be in this format:

<TAB>D/L/E/I<TAB>A/B<TAB>*Address List*<CR>

where

D/L/E/I – Domain/List/Email/IP Address

A/B – Allowed/Blocked

Address List – Address book entries separated by commas

and email addresses, domains, IP addresses, and lists are separated with a carriage return.

For example:

```
<TAB>E<TAB>A<TAB>email1@company.com,email2@company.com<CR>
<TAB>L<TAB>B<TAB>list1@company.com,list2@company.com<CR>
```

5 Click **Open**.

6 Click **Import**.

Exporting Address Book Entries

You can export entries to an Excel spreadsheet or text file.

To export address book entries:

1 On the appropriate view, click the **Export** button. The Windows **Opening *filename*** dialog displays.

2 Select either:

- **Open with Microsoft Excel (default)**
- **Save file**

3 Click **OK**.

Searching the Allowed and Blocked Lists

A search field is available to quickly find Allowed and Blocked entries in the **Allowed** and **Blocked** tables. You can access this field from either the **Allowed** view or the **Blocked** view.

To search the Allowed or Blocked lists:

1 Click the appropriate view.

2 Go to the **Search** section.

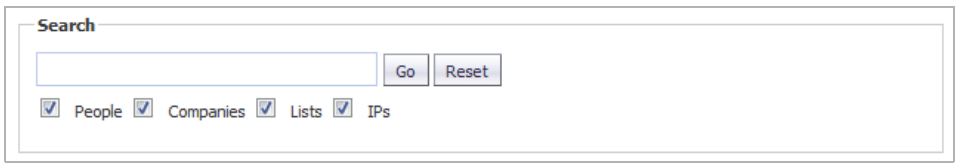

3 Enter an address or domain in the **Search** field. Enter multiple entries separated by a comma.

4 Optionally, you can filter the search between the **Type** of addresses (**People**, **Companies**, **IPs**, or **Lists** [Allowed list only]) by selecting the checkboxes below the search bar; by default, all are selected.

5 Click the **Go** button to begin the search. The results are shown in the **List** table.

To clear the search field:

1 Click the **Reset** button.

34

Managing Users

NOTE: Anti-Spam is not supported on the SuperMassive series or the NS*a* 9250-9650 firewalls.

- Anti-Spam > Manage Users on page 337
- Updating the User Table on page 338
- Enabling Non-LDAP User Authentication on page 338
- Viewing Users on page 339
- Adding Users on page 340
- Signing In as a User on page 342

Anti-Spam > Manage Users

The **Anti-Spam > Manage Users** page allows you to add, remove, and manage all users, on both the Global and LDAP servers. For more information regarding LDAP configuration, refer to Managing Users.

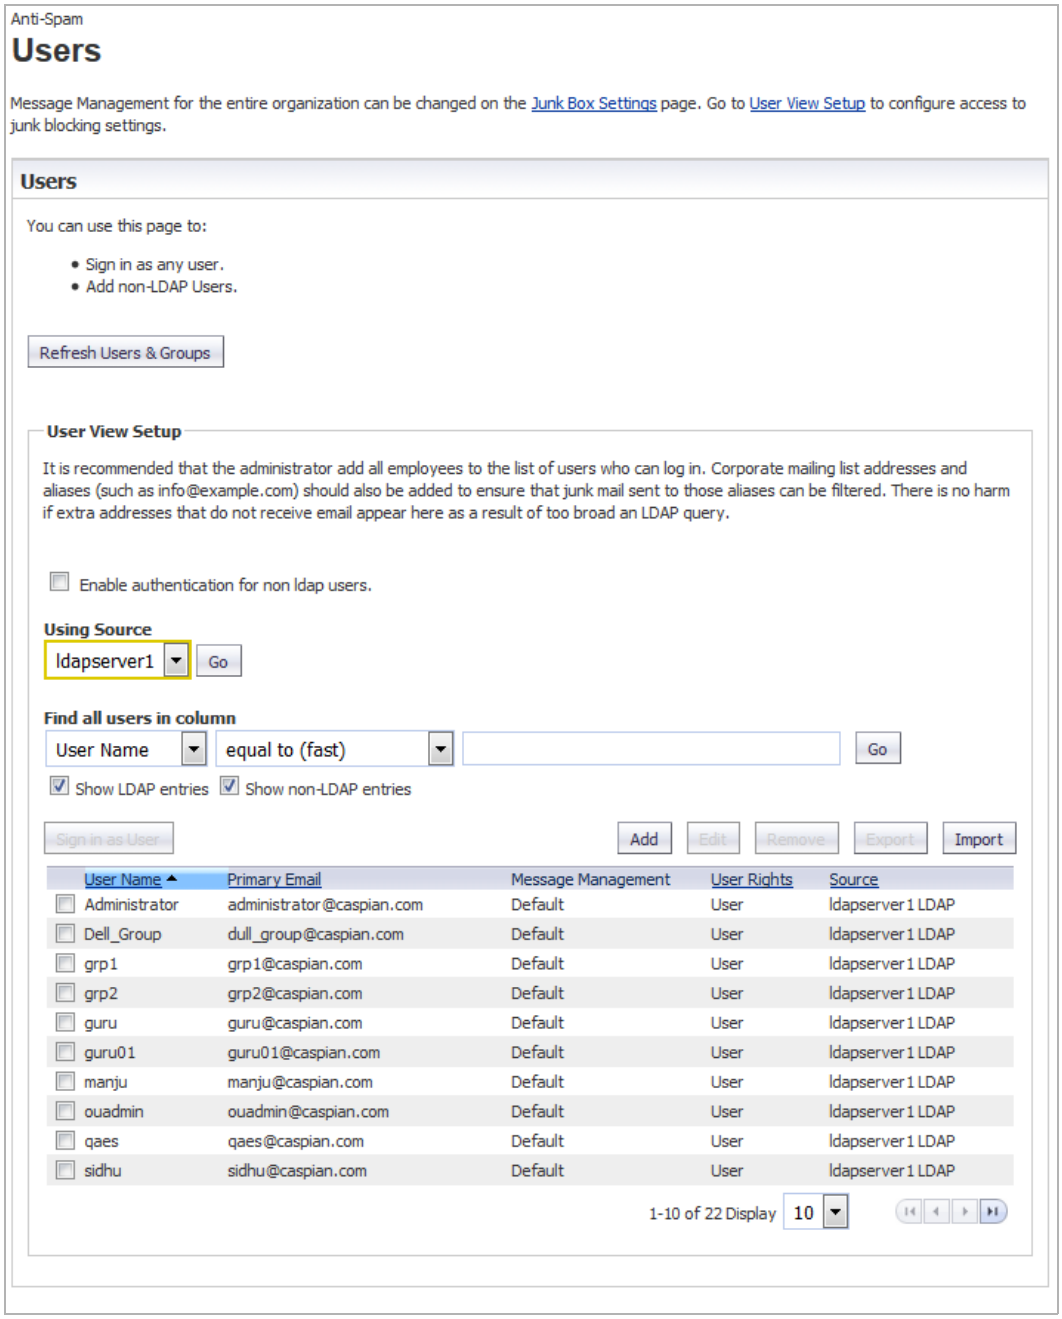

The **User** table displays this information:

Column	Description
User Name	User's user name, which may not be part of the primary email address.
Primary Email	Email address of the user.
Message Management	Displays whether the user adheres to the settings on the **Anti-Spam > Junk Box Summary** page or has modified them: • **Default** – All administrator's settings are used • **Custom** – User has changed one or more settings
User Rights	Is always **User** as user rights cannot be modified in CASS.
Source	Displays the user's server name.

Updating the User Table

To update the list of users in the User Table:

1 Navigate to the **Users** section of **Anti-Spam > Manage Users**.

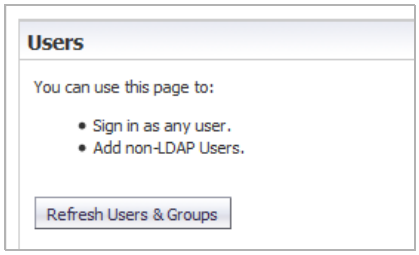

2 Click the **Refresh Users & Groups** 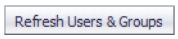 button.

Enabling Non-LDAP User Authentication

Authentication for non-LDAP users must be enabled.

To enable authentication for non-LDAP users:

1 Scroll to the **User View Setup** section of **Anti-Spam > Manage Users**.

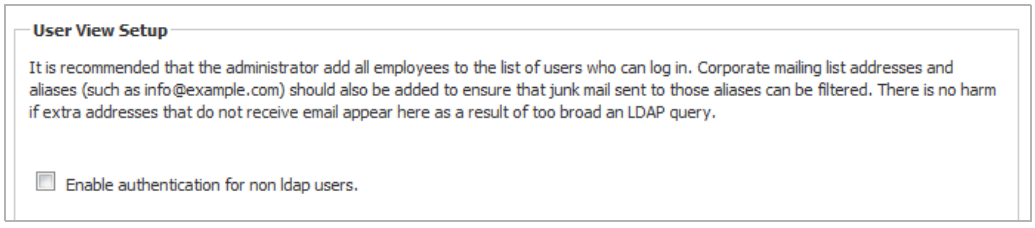

2 Select the **Enable authentication for non ldap users** checkbox. A cautionary message displays.

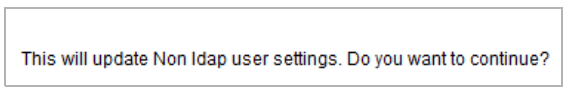

3 Click **OK**.

Viewing Users

The **User Table** displays all the users who can log in. You can filter the users to only those you want to see at the moment by:

- Selecting user type: Selecting the Type of User to View on page 339
- Selecting a source (server); see Selecting a Server's Users to View on page 339
- Specifying a particular user; see Finding a User on page 340

Selecting the Type of User to View

You can see all users, just LDAP users, or just non-LDAP users.

To select the type of user to display:

1 Scroll to the **Find All users in column** section of **Anti-Spam > Manage Users**.



2 Select which type of user:

- Only LDAP – Select the **Show LDAP entries** checkbox; this is the default if your system has only LDAP users.
- Only non-LDAP – Select the **Show non-LDAP entries** checkbox; this is the default if your system has only non-LDAP users.
- Both LDAP and non-LDAP – Select both checkboxes; this is the default if your system has both types of users.

Selecting a Server's Users to View

You can limit the **User** table to display only those users from a particular server.

To select a source (server):

1 Go to the filter section of **User View Setup**.



2 From the **Using Source** drop-down menu, select which server, or source, to view:

- **GLOBAL** (default) – A Global server is always available
- LDAP server name – If one or more LDAP servers have been added, all server names are listed.

3 Click the **Go** button.

Finding a User

You can restrict the view to just one user.

To find a user:

1 Go to the filter section of the **User View Setup** section of **Anti-Spam > Manage Users**.



2 From the **Find all users in column** drop-down menus and field, enter the selection criteria:

   a From the first drop-down menu, select:

      - **User Name**
      - **Primary Email**

   b Filter the search by these conditions from the second drop-down menu:

      - **equal to (fast)** (default)
      - **starting with (medium)**
      - **containing (slow)**

   c Enter the user's information in the field.

3 Click **GO**. The **User** table displays only those emails that meet the specified criteria, and a message displays at the top of the page.



To restore the User table display:

1 Remove the search criterion from the **Find all users in column** field.

2 Click **Go**.

Adding Users

You can add users to the list of users who can log in:

- Manually; see Adding Users Manually to the User Table on page 341
- By importing them; see Importing Users to the User Table on page 341

NOTE: It is recommended that you add all employees to the list of users who can log in. Corporate mailing list addresses and aliases (such as `info@example.com`) should also be added to ensure that junk mail sent to those aliases can be filtered. There is no harm if extra addresses that do not receive email appear here as a result of too broad an LDAP query.

Adding Users Manually to the User Table

To add a user to the Global or LDAP Server:

1 Click the **Add** button above the **User Table**. The **Add User** dialog displays.

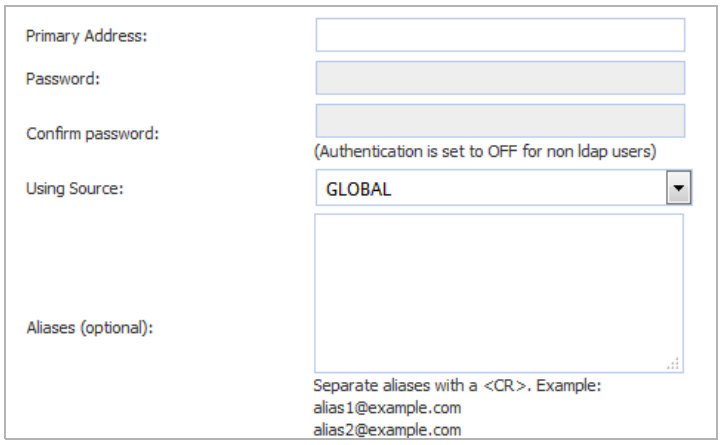

2 Enter the primary address of the user in the **Primary Address** field.

3 If the user is an LDAP user, enter the user's password in the **Password** and **Confirm User** fields.

4 Select which server the user belongs to from the **Using Source** drop-down menu.

5 Optionally, enter any Alias(es) of the user in the **Aliases** field. Separate each entry with a carriage return (<CR>).

6 Click **Add** to finish adding a user.

Importing Users to the User Table

To import a list of users from a file:

1 Click the **Import** button above the **User Table**. The **Import Users** dialog displays.

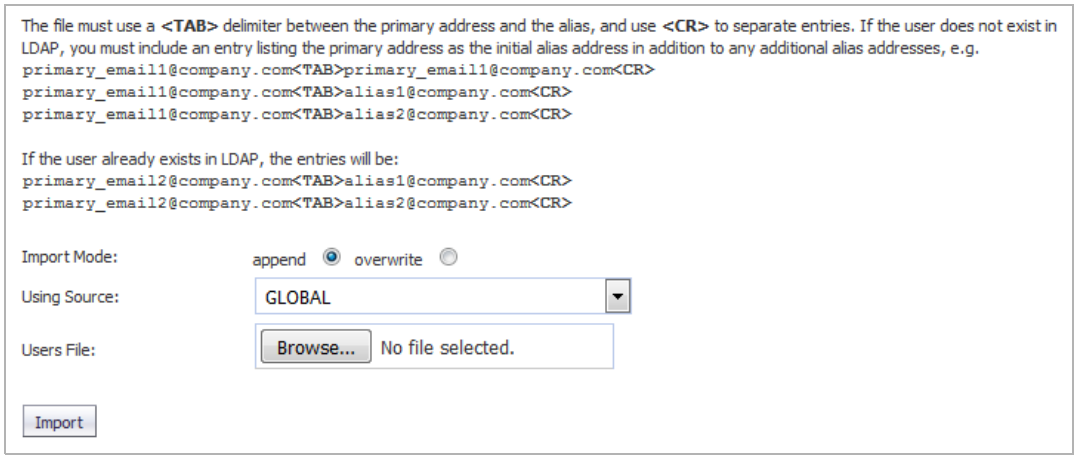

2 Select how the imported file is to be treated by selecting an **Import Mode**:

- **append** – Adds the users to the end of the file containing the list of approved users.
- **overwrite** – Replaces the existing users with the imported users.

3 Specify the server to be used as a source:

- **GLOBAL**
- LDAP server name

4 Click the **Browse** button. The Windows **File Upload** dialog displays.

5 Select the file to upload. It must be in this format, with a tab <TAB> delimiter between the primary address and the alias and a carriage return <CR> delimiter to separate entries:

```
primary_email1@company.com<TAB>primary_email1@company.com<CR>
```

For example:

```
primary_email1@company.com<TAB>primary_email@company.com<CR>
primary_email1@company.com<TAB>alias1@company.com<CR>
primary_email1@company.com<TAB>alias2@company.com<CR>
```

If the user already exists in LDAP, the entries would be:

```
primary_email2@company.com<TAB>alias1@company.com<CR>
primary_email2@company.com<TAB>alias2@company.com<CR>
```

6 Click **Open**.

7 Click **Import**.

Signing In as a User

You can sign in to a user's account to see their Email Security **INVESTIGATE | Logs > Anti-Spam Junkbox**.

To sign in as a user:

1 Navigate to the **User** table of **Anti-Spam > Users**.

2 Select the checkbox of the user you want to sign in as. The **Sign in as User** 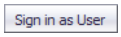 button becomes active.

3 Click the **Sign in as User** button. A separate window displays the Email Security **Anti-Spam > Junk Box** page for that user.

4 To return to the **Anti-Spam > Manage Users** page, click the **Logout** icon on the Email Security page.

35

Configuring the LDAP Server

NOTE: Anti-Spam is not supported on the SuperMassive series or the NS*a* 9250-9650 firewalls.

- Anti-Spam > LDAP Configuration on page 343
- Available LDAP Servers on page 344
- Adding an LDAP Server on page 345
- Configuring LDAP Queries on page 349
- Adding LDAP Mappings on page 351
- Configuring Global LDAP Settings on page 353
- Editing an LDAP Server Configuration on page 354
- Deleting an LDAP Server on page 355

Anti-Spam > LDAP Configuration

The **Anti-Spam > LDAP Configuration** page allows you to configure various settings specific to LDAP servers.

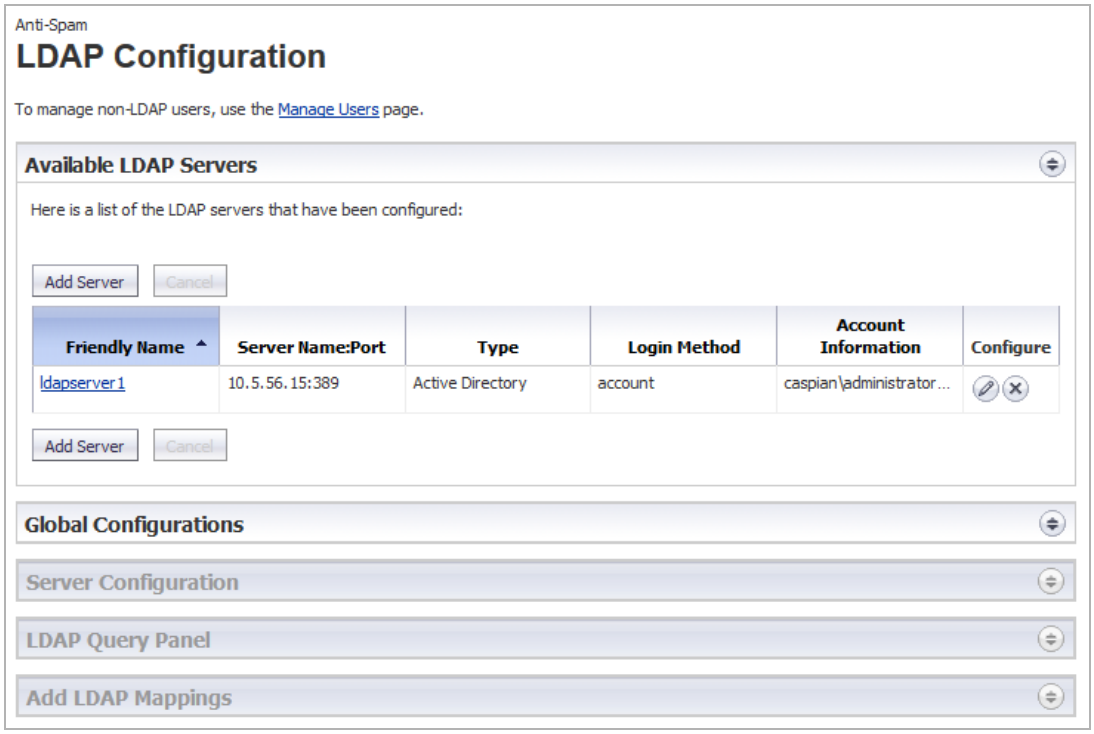

NOTE: All panels can be displayed or hidden by clicking the **Expand/Collapse** icon.

Available LDAP Servers

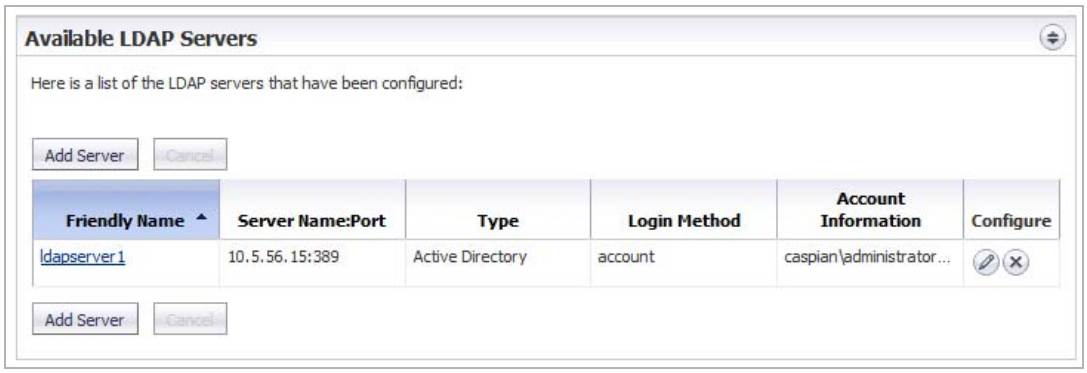

This section displays information about any LDAP Servers configured on the firewall:

- **Friendly Name** – Displays the friendly name of the server. Clicking the link displays the **Server Configuration**, **LDAP Query Panel**, and **Add LDAP Mappings** sections.
- **Server Name:Port** – Displays the IP address and port of the server.
- **Type** – Displays the type of server, such as Active Directory or OpenLDAP.
- **Login Method**
- **Account Information** – Displays
- **Configure** – Contains **Edit** and **Delete** icons.

Adding an LDAP Server

Configure a new LDAP server to enable per-user access and management.

IMPORTANT: Anti-Spam uses your existing Active Directory or LDAP server to authenticate end users as they log in to their personal Junk Boxes. The **Anti-Spam > LDAP Configuration** page must be correctly filled out to return the complete list of users who are allowed to log in to their Junk Box. If a user does not appear in this list, their email is filtered, but they can not log in to their personal junk box.

Correctly filling out the LDAP configuration requires completing the **Server Configuration** panel, **LDAP Query Panel**, and the **Add LDAP Mappings** panel.

To add an LDAP server:

1 In the **Available LDAP Servers** section, click the **Add Server** Add Server button. The **Server Configuration** section expands:

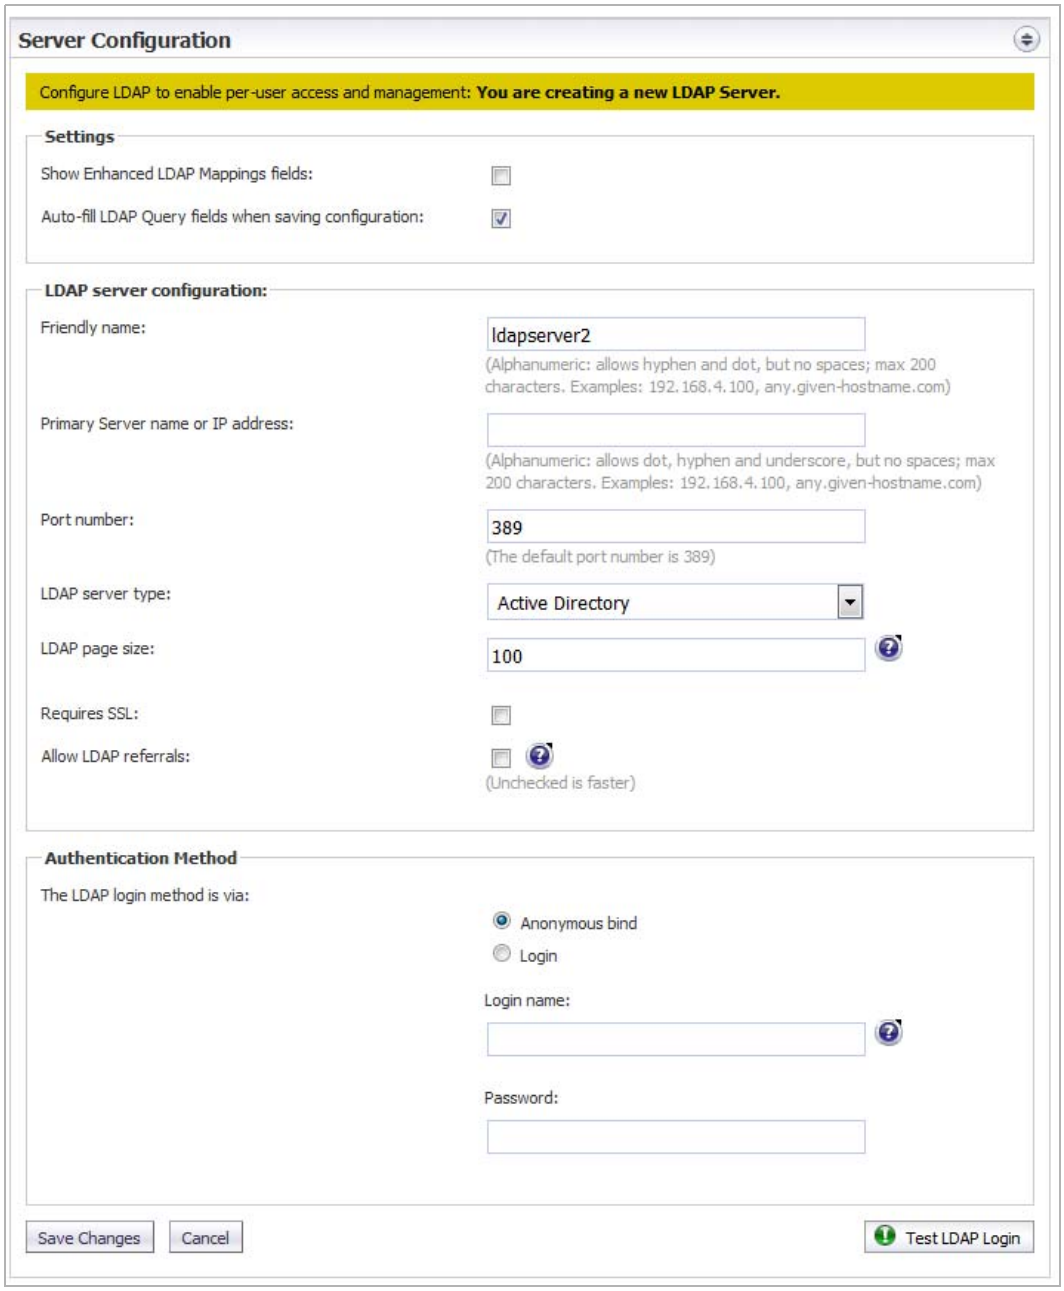

2 Optionally, in the **Settings** section, enable the **Show Enhanced LDAP Mappings fields** checkbox. When this option is enabled, fields for a secondary server display in red in the **LDAP server configuration** section.

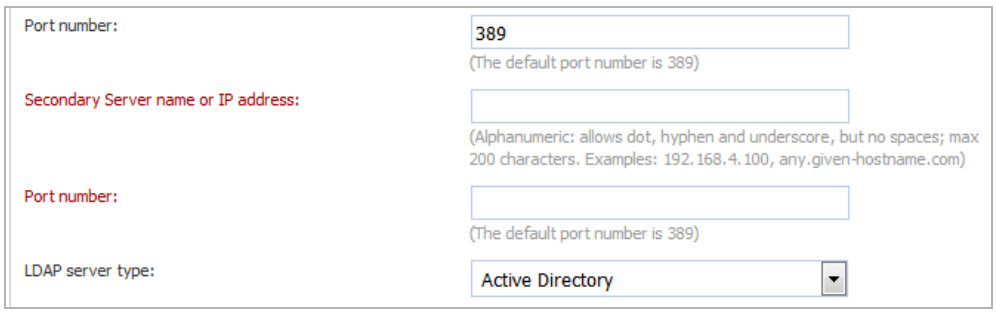

3 To have the fields in the **LDAP Query Panel** completed automatically, ensure the **Auto-fill LDAP Query fields when saving configuration** checkbox is selected. This option is selected by default.

4 In the **LDAP server configuration** section, configure the new LDAP server's settings:

TIP: The primary and secondary names and IP addresses can be up to 200 alphanumeric characters including a hyphen (-) and period (.), but no spaces. Examples:

192.168.4.100
host-name123.com

- **Friendly Name**—Enter a friendly name for the LDAP server. The default name is `ldapserver`*n*, where *n* is a sequential number.
- **Primary Server name or IP address**—The server name or IP address of the LDAP Server.
- **Port Number**—The port number of the LDAP Server. The default port number is **389**.
- **Secondary Server name or IP address**—The server name or IP address of the secondary LDAP Server.

  NOTE: The **Secondary Server name or IP address** and **Port number** options, in red, display only if you selected **Show Enhanced LDAP Mapping fields in the Settings** section.

- **Port Number**—The port number of the secondary LDAP Server. The default port number is **389**.
- **LDAP Server Type**—Select from the drop-down menu:
  - **Active Directory**
  - **Lotus Domino**
  - **Exchange 5.5**
  - **Sun ONE iPlanet**
  - **Other**
- **LDAP Page Size**—Enter the maximum page size to be queried on the LDAP Server. The default is **100**.

CAUTION: Many LDAP servers, including Active Directory, have a setting that specifies the maximum page size to be queried. If the LDAP Page Size setting exceeds that maximum page size, performance problems may occur on both the LDAP server and on . In the rare circumstances that this needs to be adjusted, consult SonicWall Technical Support.

- **Requires SSL**—To have the LDAP Server require SSL, select this checkbox. This option is not selected by default.
- **Allow LDAP Referrals**—Select this option if you have multiple LDAP servers, each of which may have different information. When LDAP referral is enabled, one LDAP server can delegate parts of

a login request for information to other LDAP servers that have more information. This delegation is called a referral and occurs when an administrator or user logs in. A referred login request can be very slow, taking 20 seconds or more. This setting is not selected by default.

NOTE: To speed log ins for administrators and users, disable this option if you have:

- Only one LDAP server.
- Two or more LDAP servers that all share the same information.

TIP: It is safe to disable referrals and then test whether any users are blocked from logging in. No data or settings are lost.

5 From the **Authentication Method** section, configure the LDAP login method for users:

- **Anonymous bind** (default) – Many LDAP servers are configured to provide the list of users to anyone who asks. This is called *Anonymous Bind*.

  TIP: Select this option first, then test it; see Step 8.

- **Login** – If the **Anonymous bind** option failed, select this option. You then need to provide a username and password to get LDAP to return the list of users.

6 If you selected:

- **Anonymous bind**, go to Step 8.
- **Login**, go to Step 7.

7 Specify the **Login name** and **Password**.

Login name is the credential used to allow a user access to the LDAP resource. Each type of LDAP server has a format for a log in name. Use the format appropriate for your server.

TIP: To see examples of the different formats, click the **Question Mark** icon by the **Login name** field.

8 To test the settings you just configured, click the **Test LDAP Login** Test LDAP Login button. The **Test Results** message displays:

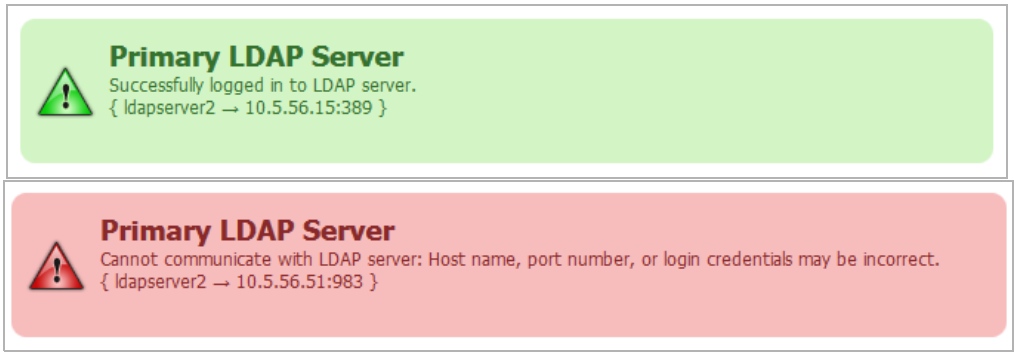

9 Click **Save Changes** to finish adding an LDAP Server. The **LDAP Query Panel** and **Add LDAP Mappings** panel display.

Configuring LDAP Queries

TIP: If you selected the **Auto-fill LDAP Query when saving configuration** option in the **Settings** section, the **LDAP Query Panel** fills with default values automatically.

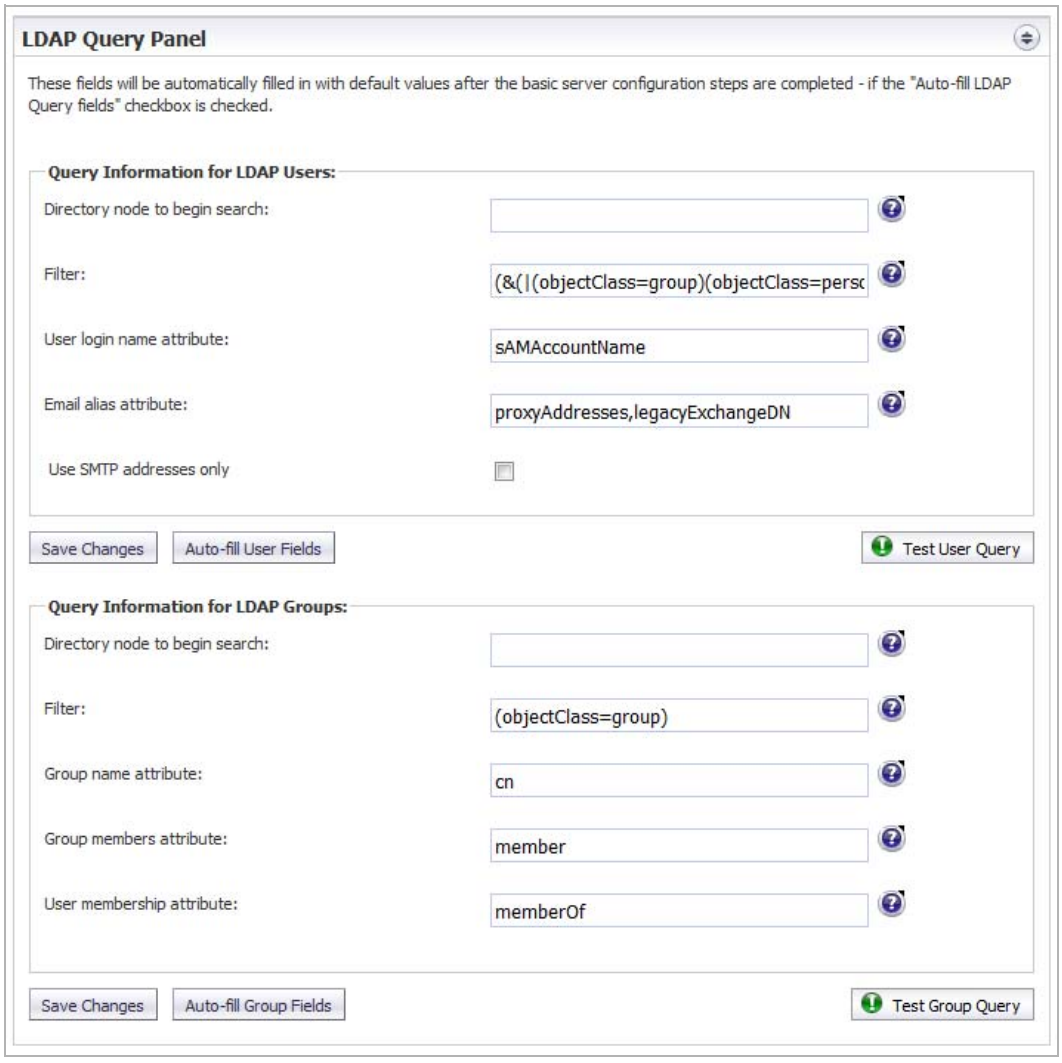

To successfully allow users to login to their Junk Box:

TIP: To examine your LDAP tree in its entirety to get a comprehensive look at your LDAP structure and its various attributes and object classes, run the free program, Softerra LDAP Browser 2.5, available at:

http://www.ldapbrowser.com/download/index.php

On a Windows PC, download the program. When it is running, to determine the best query for your network, browse to a user on the network and examine their attributes.

1 In the **LDAP Query Panel**, go to the **Query Information for LDAP Users** section.

TIP: If you did not specify **Auto-fill LDAP Query fields when saving configuration** in the **Settings** section, you can click the **Auto-fill User Fields** Auto-fill User Fields button to do so.

2 To use the optional Groups functionality, in the **Directory Node to Begin Search** field, specify a full LDAP directory path that points towards a node (directory inside LDAP) containing the information for all groups in the directory. This path narrows the search for LDAP groups to a reasonable size.

The information contained in LDAP is organized into a directory tree much like an ordinary file system. Each directory is specified as a `name=value` pair, where:

- `name` is commonly:

DC (domain component)	OU (organizational unit)
DN (distinguished name)	O (organization)

- **value** is commonly one segment of a fully specified hostname (for example, the word `companyxyz` in `sales.companyxyz.com`).

To specify a particular node in LDAP you use a comma-separated list. To specify multiple nodes to search in, use the ampersand (&) character between full paths.

For example, if the hostname of a particular machine inside `companyxyz` was `computer27.sales.companyxyz.com`, the LDAP path might be:

```
DC=computer27,DC=sales,DC=companyxyz,DC=com
```

TIP: To see examples for the various directory types, click the **Question Mark** icon next to the **Directory Node to Begin Search** field

3 Enter an LDAP filter in the standard LDAP filter syntax in the **Filter** field.

Anti-Spam must be instructed on how to find and identify users and mailing lists. By specifically stating the Object Class and mail attribute in the **Filter** field, non-primary email accounts (such as printers and computers) are not included during an LDAP query. Focusing on primary user accounts speeds up the query.

The **Filter** field contains an example syntax:

```
(&(|(objectClass=group)(objectClass=person)(objectClass=publicFolder))
(mail=*))
```

All LDAP filters are grouped in parenthesis, and the filter itself has a pair of parentheses surrounding the whole string. The very next character from the left is an ampersand (&). The LDAP filter syntax is prefix notation, which means this filter only returns the logical AND of three sub-filters, each grouped in parentheses. Other operators include a pipe (|) for OR and an exclamation point (!) for NOT.

TIP: To see examples for the various directory types, click the **Question Mark** icon next to the **Filter** field

4 Specify the text attribute a user uses fora login name in the **User Login Name Attribute** field. The generally accepted attribute for this field is **`sAMAccountName`**, which is the default. This attribute should work for Microsoft Windows, as well as all other environments.

IMPORTANT: This field works in conjunction and needs to agree with the **Filter** field. If you change `sAMAccountName`, you must change it in both the **Filter** field and the **User Login Name Attribute** field.

TIP: To see examples for the various directory types, click the **Question Mark** icon next to the **User Login Name Attribute** field

5 Specify the email address, employee ID, phone number, or other alias attributes that link a single user to his or her junk box in the **Email Alias Attribute** field.

At many companies, an end user has multiple email accounts that all map to one true email account. For example, `JohnS@example.com` and `John.Smith@example.com` might both be valid email addresses for John Smith's InBox. Anti-Spam supports this by allowing an end user to have one junk email box that groups all email from their various email addresses.

The generally accepted single attribute for this field is **proxyAddresses**. All other attributes must be separated by a comma. For example:

- `proxyAddresses,legacyExchangeDN`
- `proxyAddresses,EmployeeID,PhoneNumber`

TIP: In Microsoft Windows environments, the single attribute, **proxyAddresses**, is often sufficient.

To see examples for the various directory types, click the **Question Mark** icon next to the **Email Alias Attribute** field

6 Optionally, test to see if your settings work, click **Test User Query** Test User Query button under the **Query Information for LDAP Users** section.

7 Save the changes by clicking **Save Changes** under the **Query Information for LDAP Users** section.

8 Go to the **Query Information for LDAP Groups** section.

TIP: If you did not specify **Auto-fill LDAP Query fields when saving configuration** in the **Settings** section, you can click the **Auto-fill Group Fields** Auto-fill Group Fields button to do so.

9 To use the optional Groups functionality, in the **Directory Node to Begin Search** field, specify a full LDAP directory path that points towards a node (directory inside LDAP) containing the information for all groups in the directory. This narrows the search for LDAP groups to a reasonable size. For further information about this setting, see Step 2.

10 To instruct Anti-Spam on how to find and identify users and mailing lists, enter an LDAP filter in the standard LDAP filter syntax in the **Filter** field. The field contains an example syntax. For further information about this setting, see Step 3.

11 Specify the attribute of the group that corresponds to Group names in the **Group name attribute** field

12 A common way to specify a group is a mailing list. In the mailing list entry in LDAP, there is one particular field that specifies the members of the list. Enter that information in the **Group members attribute** field.

13 In some LDAP configurations, there is an attribute, inside each user's entry in LDAP, that lists the groups or mailing lists of which this user is a member. Specify that attribute in the **User membership attribute** field.

14 Optionally, test to see if your settings work, click the **Test User Query** Test User Query button under the **Query Information for LDAP Groups** section.

15 Save the changes by clicking **Save Changes** under the **Query Information for LDAP Groups** section.

Adding LDAP Mappings

If you are using a Microsoft Windows environment, you need to specify the NetBIOS domain name in the **Add LDAP Mappings** panel.

NOTE: The NetBIOS domain name is sometimes called the pre-Windows 2000 domain name.

To add LDAP mapping:

1 Determine your domain name(s).

a Login to your domain controller.

b Navigate to **Start** > **All Programs** > **Administrative Tools** > **Active Directory Domains and Trusts**.

c Highlight your domain from the **Active Directory Domains and Trusts** dialog.

d Click **Action**.

e Click **Properties**. The domain name(s) appear on the domain's **Properties** dialog on the **General** view.

f Record the domain name(s).

2 Navigate to the **Add LDAP Mappings** panel of **Anti-Spam > LDAP Configuration**.

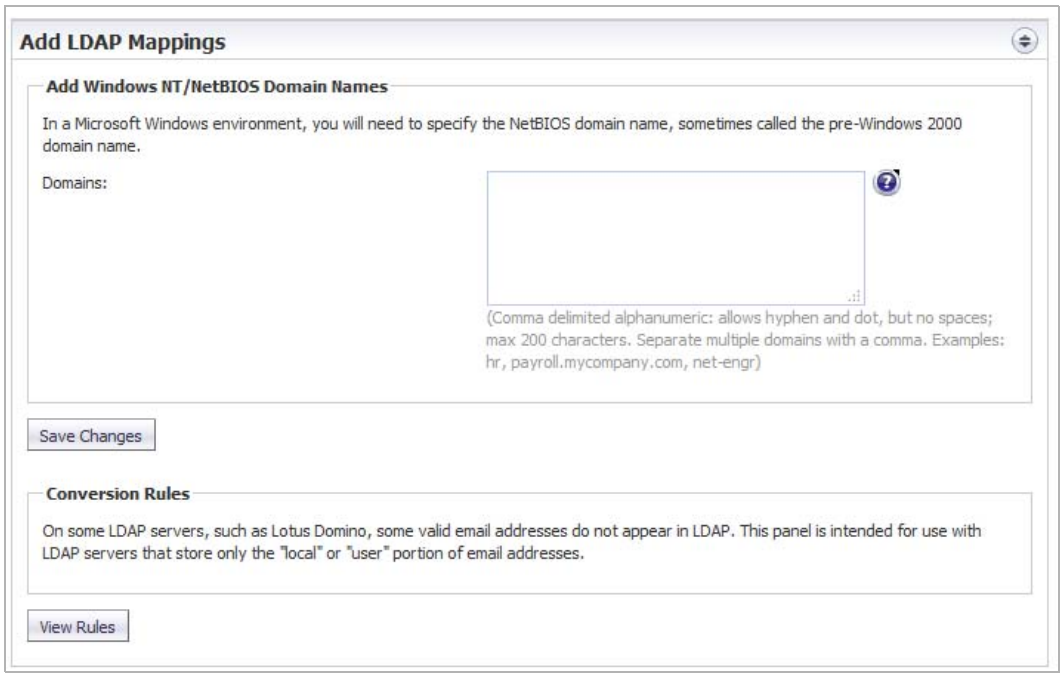

3 Add the NetBIOS domain name(s) to the **Domains** field. Add a maximum of 200 alphanumeric characters. Separate multiple domains with a comma. Hyphens (-) and periods (.) are allowed.

4 Click **Save Changes**.

5 On certain LDAP servers, such as Lotus Domino, some valid email addresses do not appear in the LDAP. The **Conversion Rules** section changes the way the SonicWall Email Security appliance interprets certain email addresses to provide a way to map the email address to the LDAP Server.

If you:

- Have one of these servers, go to Step 6.
- Do not have one of these servers, you have finished configuring LDAP.

6 To map these addresses, click on the **View Rules** View Rules button. The **LDAP Mapping** dialog displays.

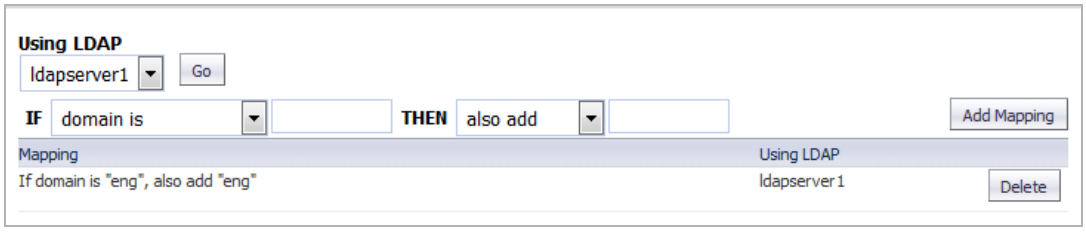

7 Select the LDAP Server you are using from the drop-down menu.

8 Click **Go**.

9 Optionally, add a mapping:

a From the **IF**/**THEN** drop-down menus and fields, select:

- **domain is**—Adds additional mappings from one domain to another; in the field, specify a domain to be mapped
  - **replace with**—Replaces the domain with the one specified

Example: **IF domain is** `engr.corp.com` **THEN replace with** `corp.com`, then email addressed to `anybody@engr.corp.com` is sent to `anybody@corp.com`

- **also add**—Adds the second domain to the list of valid domains

  Example: **IF domain is** `corp.com` **THEN also add** engr.`corp.com`, then if `corp.com` is found in the list of valid LDAP domains, `engr.corp.com` is added to the list

- **left side character is**—Adds character substitution mappings; in the field, specify a character to be substituted

  - **replace with**—Replaces any character specified to the left of the at sign (@) in the email address with the new character

    Example: **IF left side character is** `_` **THEN replace with** `-`, then email addressed to `Jane_Doe@corp.com` is sent to `Jane-Doe@corp`.com

  - **also add—**Adds a second email address to the list of valid email addresses

    Example: **If left side character is** `_` **THEN also add** `-`, then email addressed to either `Jane_Doe@corp.com` or `Jane-Doe@corp.com` is a valid email address

b Click the **Add Mapping** Add Mapping button to finish adding the conversion rules.

NOTE: To delete a mapping, click the **Delete** button for that mapping.

Configuring Global LDAP Settings

Global LDAP settings apply universally across all LDAP server configurations.

To configure global settings:

1 Navigate to the **Global Configurations** panel in **Anti-Spam > LDAP Configuration**.

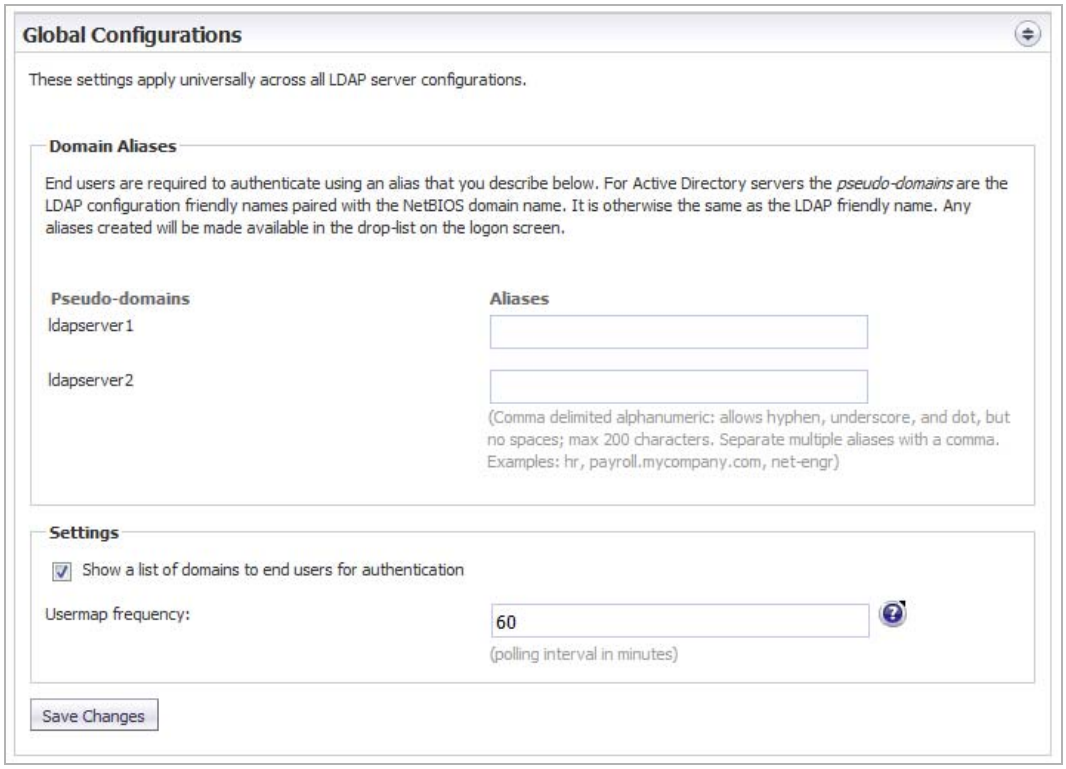

2 In the **Domain Aliases** section, enter one or more aliases for one or more servers for a maximum of 200 alphanumeric characters for each server. Separate multiple aliases with a comma. Hyphens (-), underscores (_), but not spaces, are allowed.

   End users must authenticate using an alias configured here. For Active Directory servers, the pseudo-domains are the LDAP friendly names paired with the NetBIOS domain name. Any aliases are available for authentication in the drop-down menu on the log-on screen if that option is selected in the **Settings** section.

3 To allow the end user to see a list of domains and aliases when logging on, in the **Settings** section, select **Show a list of domains to end users for authentication**. This setting is selected by default.

4 Specify the number of minutes between refreshes of the list of users on the system in the **Usermap Frequency** field.

   This setting applies to the list of aliases and lists of members of groups. In most cases, increase this setting only to lower the load on the LDAP server. Depending on your other settings, fetching the user list once every 24 hours (1440 minutes) is acceptable and results in less load on the LDAP server.

   NOTE: Usermap frequency does not affect a user's ability to log on, which is a real-time reflection of the LDAP directory

5 Click **Save Changes**.

Editing an LDAP Server Configuration

Editing an LDAP server configuration requires the same settings as adding a server.

To configure an LDAP server:

1 From the list of available LDAP servers, click the **Edit** icon. These sections expand for editing:

- **Server Configuration** – see Adding an LDAP Server
- **LDAP Query Panel** – see Configuring LDAP Queries
- **Add LDAP Mappings** – see Adding LDAP Mappings

Deleting an LDAP Server

To delete an LDAP server:

1 Click the **Delete** icon for the server to be deleted. A warning message appears:

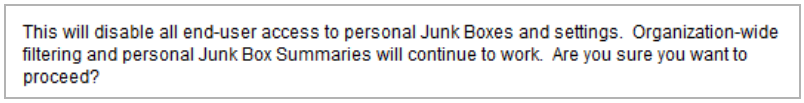

2 Click **OK**. A success message appears at the top of the **Anti-Spam > LDAP Configuration** page.

36

Downloading Anti-Spam Desktop Buttons

NOTE: Anti-Spam is not supported on the SuperMassive series or the NS*a* 9250-9650 firewalls.

- Anti-Spam > Downloads

Anti-Spam > Downloads

The **Anti-Spam > Download Anti-Spam Tools** page allows you to download and install one of SonicWall's latest spam-blocking buttons on your desktop.

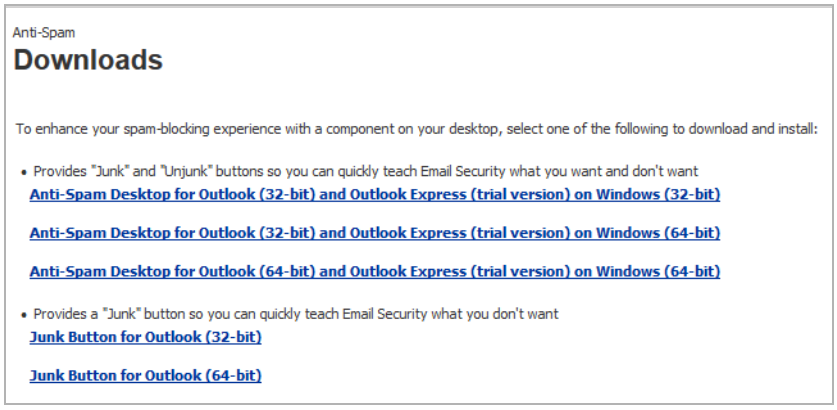

By clicking on a link, you can download these buttons to your desktop:

- Junk and Unjunk buttons to teach Email Security what you want and don't want easily and quickly; select one:
  - **Anti-Spam Desktop for Outlook (32-bit) and Outlook Express (trial version) on Windows (32-bit)**
  - **Anti-Spam Desktop for Outlook (32-bit) and Outlook Express (trial version) on Windows (64-bit)**
  - **Anti-Spam Desktop for Outlook (64-bit) and Outlook Express (trial version) on Windows (64-bit)**
- Junk button to teach Email Security what you want easily and quickly; select one:
  - **Junk Button for Outlook (32-bit)**
  - **Junk Button for Outlook (64-bit)**

Part 5

Security Config | Appendix

- SonicWall Support

A

SonicWall Support

Technical support is available to customers who have purchased SonicWall products with a valid maintenance contract.

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://www.sonicwall.com/support.

The Support Portal enables you to:

- View knowledge base articles and technical documentation
- View and participate in the Community forum discussions at 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, visit https://www.sonicwall.com/support/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.

SonicWall SonicOS 6.5 Security Configuration
Updated - May 2020
Software Version - 6.5.4
232-001811-04 Rev D

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.

End User Product Agreement

To view the SonicWall End User Product Agreement, go to: https://www.sonicwall.com/en-us/legal/license-agreements.

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 USD 25.00 payable to "SonicWall Inc.", to:

General Public License Source Code Request
SonicWall Inc. Attn: Jennifer Anderson
1033 McCarthy Blvd
Milpitas, CA 95035